Received: 6 November 2023 | Accepted: 20 June 2024

DOI: 10.1111/jfir.12423

ORIGINAL ARTICLE

The Journal of Financial Research

The taxonomy of tail risk

Evarist Stoja[1] | **Arnold Polanski**[2] | **Linh H. Nguyen**[3]

[1]University of Bristol Business School, 15-19 Tyndall's Park Road, Bristol BS8 1PQ, UK

[2]School of Economics, Norwich Research Park, University of East Anglia, Norwich NR4 7TJ, UK

[3]University of Nottingham Business School, Jubilee Campus, Wollaton, Nottingham NG8 1BB, UK

Correspondence
Evarist Stoja, University of Bristol Business School, 15-19 Tyndall's Park Road, Bristol, BS8 1PQ, UK.
Email: E.Stoja@bristol.ac.uk

Abstract

We use tail events at different levels of severity to define an asset's tail risk and to decompose the latter into a systematic and an idiosyncratic component. The systematic component captures an asset's tendency to experience joint tail losses with the market and generalizes a classic tail dependence coefficient. However, the idiosyncratic component consists of two parts: idiosyncratic tail risk that leads to asset-specific tail losses and tail risk cushioning that dampens the tail losses emanating from the market. Tail risk cushioning is a novel concept that arises naturally in our framework, is consistent with the previous two and completes the taxonomy of tail risk. We examine the performance of our tail risk decomposition on a large dataset, confirming some previous results on tail risk and uncovering new theoretical and empirical findings.

JEL CLASSIFICATION
C14, G11, G12

1 | INTRODUCTION

In this article, we propose novel measures of systematic and idiosyncratic tail risks (STR and ITR, respectively). These measures arise organically from a decomposition of the tail risk of asset returns and are mutually consistent. In addition to these types of tail risk, we introduce a novel concept, tail risk cushioning (TRC)—the tendency of an asset to dampen tail risk emanating from the systematic factor—and propose a measure that encapsulates it. Tail risk cushioning arises naturally in our framework, is fully consistent with the previous two types, and completes the taxonomy of tail risk. We relate the systematic, idiosyncratic and cushioning component to other measures of tail risk and examine their impact on asset returns. Specifically, we measure how much of the (expected) return of an asset can be attributed to each of these components. Following the literature (see, for example, Bali, Cakici, and Whitelaw (2014)), we apply the Fama and MacBeth (1973) methodology and use a large cross section of stock returns and the Fama-French systematic factors to estimate the significance and magnitude of the premia earned by exposure to the systematic, idiosyncratic and cushioning components of tail risk.

We define tail risk as the probability of a (joint) exceedance of certain thresholds. The thresholds are defined by Value-at-Risk (VaR) which shows how much an investor is likely to lose with a given probability over a given horizon. VaR has been extensively embraced by regulators and practitioners in financial markets under the Basel II and III frameworks as the basis of risk measurement for the purpose of ensuring regulatory capital adequacy, risk management and strategic planning. In extensive empirical exercises, we find a significant positive risk premium associated with the systematic component of tail risk as well as a significant negative premium for tail risk cushioning. However, we find that exposure of a portfolio to idiosyncratic tail risk earns a negative risk premium, contradicting the theory but extending the findings of Ang et al. (2006), among others, on the negative relation between expected stock returns and idiosyncratic volatility to idiosyncratic tail risk.

Our findings are qualitatively similar, although statistically more significant, if instead of tail risk, one examines the impact of (the components of) downside risk—defined as the tendency of an asset to generate losses—on expected returns. Interestingly, similar observations hold if instead one examines the premia of the tail and downside risk measures at different threshold levels α for three months-ahead returns. Indeed, while the coefficients seem to have increased across all three risk measures, the statistical significance has reduced slightly though generally still well above the usual significance levels. Digging deeper, we find considerable nuances, as is usually the case in the empirical asset pricing literature. For example, examining the impact of tail risk in the period immediately following a "financial disaster" and calmer periods, we find that the impact of STR on returns is much stronger in the years subsequent to a market crash. This finding is similar to that of Chabi-Yo et al. (2018) and supports the model of Gennaioli et al. (2015). However, whereas TRC is not significant in periods following market crashes, it appears to be relevant in the remaining years. Further, ITR has a negative impact, though marginal at best, on expected returns in both the post-market crash and the remaining periods, consistent with the findings of Bali et al. (2014).

The paper is structured as follows. In Section 2, we discuss the context within the literature, lay out the theoretical framework, and discuss an array of properties of these new measures of tail risk. In Section 3, we present our empirical results, while Section 4 summarizes the paper. Appendix A contains the proofs of our theoreti-cal results, while Appendices B and C contain initial data analyses and robustness analyses, respectively.

2 | THEORETICAL FRAMEWORK

2.1 | The Context within the Literature

Classic finance theory argues that diversification can eliminate idiosyncratic but not systematic risk, and therefore only exposure to the latter earns a risk premium. Exposure to the former is not rewarded with a premium since it can and should be diversified away (see, for example, Chen and Sears (1984); Statman (1987)). Indeed, diversification is often referred to as the only free lunch in finance. However, many examples, (e.g., the financial crisis 2007–2009, the eurozone crisis 2010–2011, or the crash due to the COVID-19 pandemic) illustrate that, in practice, diversification can abjectly fail to protect investors from tail events. According to an old saying in investment circles, the only things that go up in a crisis are correlations. Assets which before a crisis had low or even negative correlations leading to a well-diversified portfolio become highly interdependent during the crisis, therefore compounding losses instead of offsetting them. These considerations have led to a different approach to investing that can be summarized in two words: concentrated portfolio.[1]

By construction, concentrated portfolios carry significant idiosyncratic risk. Assuming under-diversification, various theories predict a positive relationship between the idiosyncratic risk and the expected stock returns in the cross section (see, for example, Levy (1978); Malkiel and Xu (2004)). Under-diversified investors will demand a return compensation for bearing idiosyncratic risk. In stark contrast however, Ang et al. (2006) find that, in the cross-section, high idiosyncratic volatility in one month predicts abysmally low average returns in the next month, a finding which they call "a substantive puzzle".[2]

Volatility measured as standard deviation or variance of portfolio returns may not be an adequate risk measure, especially for undiversified portfolios. For example, ample empirical evidence shows that individuals value losses and gains differently, usually assigning greater weight to losses (Kahneman and Tversky (1979); Barberis (2013)). Downside risk is particularly important when asset returns are asymmetrically distributed and investors are averse to disasters. Menezes et al. (1980) argue that investors tend to avoid positions that may lead to large losses even though they may have low probability. Rietz (1988) and Barro (2006) show that tail risks are important to explain some of the asset pricing puzzles.

The literature lacks a consensus on a unique definition for the concept of “systematic tail risk” associated with an asset. For instance, some studies base their analysis on statistical moments (e.g., Ang et al. (2006); Conrad et al. (2013)), while others employ co-moments (e.g., Harvey and Siddique (2000); Dittmar (2002), Fran, cois, Heck, François et al. (2022)). However, research relying on moment- and co-moment-based risk metrics offers only indirect insights into how tail risk influences asset pricing. Direct evidence on the role of tail risk remains inconclusive.

Several studies investigating the impact of VaR on expected returns find a positive correlation (e.g., Bali, Demirtas, and Levy (2008); Bali and Cakici (2004)). However, these studies do not differentiate between the systematic and idiosyncratic components of VaR. More recently, Atilgan, Bali, Demirtas, and Gunaydin (2020) find a robust negative effect of tail risk, proxied by VaR, on expected returns. They attribute this effect to behavioral biases, which suggests the significance of idiosyncratic tail risk in asset pricing. However, other studies find insignificant or negative results when examining both idiosyncratic and systematic tail risk, while they do find some supporting evidence for a hybrid tail risk measure (e.g., Bali et al. (2014)).

Chabi-Yo et al. (2018) apply the classic tail dependence coefficient of Sibuya (1960) as a measure of systematic tail risk and find that the risk premium corresponding to this measure is substantial. On the other hand, van Oordt and Zhou (2016) rely on Arzac and Bawa's (1977) asset pricing model and introduce the concept of tail beta, calculated as the product of a tail dependence coefficient and the relative tail risk. They find that systematic tail risk, proxied by tail beta, is linked to future stock returns but does not earn a significantly positive risk premium. Stoja et al. (2023) suggest that negative common features in systematic tail risk measures may be the reason behind the contradictory findings in these two studies. Interestingly, van Oordt and Zhou (2016) find that stocks with high (low) tail betas have high (low) tail dependence with the market, as intuition would suggest, but also high (low) idiosyncratic risk (see their Table 1). This implies a positive correlation between tail dependence and idiosyncratic risk. However, Chabi-Yo et al. (2018) find that idiosyncratic risk correlates negatively with tail dependence (see their Table 2). Because in these models,

[1]For evidence on concentrated portfolios, see Polkovnichenko (2005) who shows that the median number of stocks in household portfolios is two at several points between 1989 and 1998 and increases to three in 2001. Similarly, Goetzmann and Kumar (2008) find that during 1991–1996, the median number of stocks in a portfolio of individual investors is three. See also the interview with the manager of Henderson European Focus Fund explaining that this approach to portfolio construction is in direct response to the demand by clients that he defends as making economic sense at the following link https://www.newstatesman.com/politics/2019/05/a-truer-active-more-idiosyncratic-portfolio-2.

[2]For a recent discussion of the idiosyncratic volatility puzzle and the related empirical studies see Chichernea and Slezak (2013); Stambaugh et al. (2015); Hou and Loh (2016).

idiosyncratic and systematic (tail) risk are not necessarily mutually consistent, it is not clear then what drives these sharply conflicting results or how to reconcile them.

Unlike systematic tail risk, idiosyncratic tail risk has attracted much less attention. Huang, Liu, Rhee, and Wu (2012) model idiosyncratic tail risk with a two-step procedure. In the first step, stock returns are regressed on systematic risk factors. Then, in the second step, idiosyncratic tail risk is estimated as the tail index of the regression residuals. This approach, standard in the literature in the context of idiosyncratic risk, deserves careful consideration in the context of idiosyncratic tail risk. We discuss this issue in detail in the next section.

2.2 | Systematic Tail Risk Defines its Idiosyncratic Counterpart

Two prominent articles that study systematic tail risk are Chabi-Yo et al. (2018) and van Oordt and Zhou (2016). Both their systematic tail risk measures rely, to different degrees, on the classic tail dependence coefficient proposed by Sibuya (1960) (see their Equations (1) and (7)). Essentially, they measure the systematic tail risk of an asset by joint occurrences of extreme events - occurrences when both the market and the asset exceed some thresholds, in both cases their VaRs. Following this approach, idiosyncratic tail risk can then be defined by the outcomes in which the asset is in distress (i.e., exceeds its VaR) while the market is not.

An approach that relies on a threshold such as VaR to estimate idiosyncratic tail risk contrasts sharply with other measures. For example, Huang et al. (2012) use a two-step procedure to estimate idiosyncratic tail risk. In the first step, stock returns are regressed on systematic risk factors. Then, in the second step, idiosyncratic tail risk is estimated as the tail index of the residuals of the regression. In the context of the Sibuya-based tail risk approach, the two-step procedure can lead to misclassification of tail events. Some events may be "double counted" as both systematic and idiosyncratic, and other events may be included in the calculation of idiosyncratic tail risk, when in fact they should not. A careful consideration of these issues and a systematic categorization of tail events is important given their paramount importance for stock returns and since most tail events are idiosyncratic (see, for example, Bali et al. (2014).

To illustrate, suppose that at some severity level α, the market (systematic risk factor) has a VaR of negative five percent, but on a particular day generates a return of 10 percent. Suppose further that an asset has the same VaR of negative five percent, a (tail) beta of one but on the same day generates a return of two percent. Thus, the residual term is negative eight percent, which is large but does not result in a tail event for the asset because its VaR has not been breached. However, the two–stage procedure would classify this as an idiosyncratic tail event. Suppose that on another day, the market generates a return of negative six percent and the asset generates a return of −15 percent. While the residual term of negative nine percent is large, it would appear incorrect to classify this as an idiosyncratic tail event. One could argue that since the market has already breached its VaR, this event should count as a systematic tail event. Indeed, this observation is the essence of the classic tail dependence coefficient of Sibuya (1960) (see also Joe (1997)) which forms the basis of many systematic tail risk measures, including those of van Oordt and Zhou (2016) and Chabi-Yo et al. (2018) as well as systemic risk measures like CoVaR of Adrian and Brunnermeier (2016).

In our framework, for any given level of an asset's total tail risk, systematic tail risk accounts for some part of tail risk, while the idiosyncratic component (composed of idiosyncratic tail risk and tail risk cushioning) accounts for the remaining part. Thus, for a given level of tail risk, a stock with high systematic tail risk will tend to have a low idiosyncratic component and vice versa. Therefore, this approach is the direct analogue of the total volatility decomposition of a stock into systematic and idiosyncratic volatility in the Single-Index Model (SIM; see, for example, Sharpe (1963)). This is an important feature of the model and an advantage relative to other frameworks, in which it is not clear how systematic and idiosyncratic tail risk relate to each other. Of course, this approach is not without its own issues. Going back to the example, suppose the market on a particular day generates a return of negative four percent while the asset has a return of negative six percent. Since the asset's VaR has been breached, this would contribute to idiosyncratic tail risk even though the residual term is negative two percent, many times smaller in absolute value than the negative eight percent term which previously did not contribute to idiosyncratic tail risk. However, this issue ensues from the decision of an investor as to what level of return constitutes a severe tail event (i.e., the rather arbitrary but unavoidable decision as to where exactly the threshold is, the point demarcating moderate losses from tail event losses, that the VaR condition imposes). Any model, including that of van Oordt and Zhou (2016) or Chabi-Yo et al. (2018), that relies on cut-off points for the definition of tails would be subject to this issue.[3]

2.3 | A Tree-Model of Asset Returns

This section presents a simple model of asset returns and shows how it leads directly to our decomposition of tail risk. Assume that a SIM holds and the excess return r_i of stock i is approximately equal to (tail) beta β_i times the market's excess return r_m, where the latter exceeds its threshold with the time-independent probability f.

We assume $\beta_i \geq 0$ for the sake of consistency with the literature and consider two regimes. In the first regime, which occurs with probability p_i, $\beta_i > 0$ and the error term is distributed with a "moderate" dispersion (denoted ϵ_i). Hence, stock i's excess return does not deviate significantly from the prediction of the SIM. More specifically, stock i does (not) exceed its threshold whenever the market does (not). In the second regime, which materializes with the complementary probability $1 - p_i$, $\beta_i = 0$ and the error is distributed with a "large" dispersion (E_i). In this regime, asset i exceeds its threshold independently of the market with probability q_i due to a large negative error term materializing (E_i^-) or, with probability $1 - q_i$, does not exceed it due to a moderate or large positive error term (E_i^+).

For simplicity, all probabilities are assumed to be time-independent. However, this assumption is not essential—the setting can be generalized easily to allow for time-dependent probabilities.

Figure 1 depicts the event tree which illustrates the different paths that lead to the mutually-exclusive and collectively-exhaustive outcomes (i.e., joint tail events represented by the final nodes).

The outcomes in Figure 1 correspond then to the following four tails. In tail $T_\varnothing$, no threshold exceedance has occurred; in tail $T_{\{m\}}$, the market has exceeded its threshold but not the asset; in tail $T_{\{i\}}$ the asset has exceeded its threshold but not the market. Finally, in tail $T_{\{i,m\}}$, both have exceeded their respective thresholds. Figure 2 depicts these outcomes.

The four areas in Figure 2 correspond to the four possible outcomes (i.e., joint tail events) in Figure 1 that materialize due to three binary events: 1) the realization of the market (systematic factor) return (whether it is above or below a given threshold), 2) the occurrence of the first or the second regime and, in the latter case, 3) the realization of the idiosyncratic shock (whether it is above or below a given threshold). Below we show that the parameter values of f, p_i and q_i in the event tree can be uniquely calculated from the observed data on the tail events.

2.4 | The Taxonomy of Tail Risk

When the thresholds in the model above delineate extreme (or tail) events, we can interpret the areas in Figure 2 as follows: the region $T_\varnothing$ corresponds then to the day-to-day moderate losses as well as gains. There is an extensive literature that examines various asset pricing predictions in this region. In fact, the majority of asset pricing studies relate to this area. More recently and, in particular, since the financial crisis of 2007–2009, there has been rather

[3]In a wider context, Supper et al. (2020) explicitly caution that "… several key results from the literature (e.g., Chabi-Yo et al. (2018) […]) need to be treated with care" (page 14) as the dependent structure could be misestimated. As a way to alleviate the impact that this choice may have on the results, we employ a wide range of cut-off points.

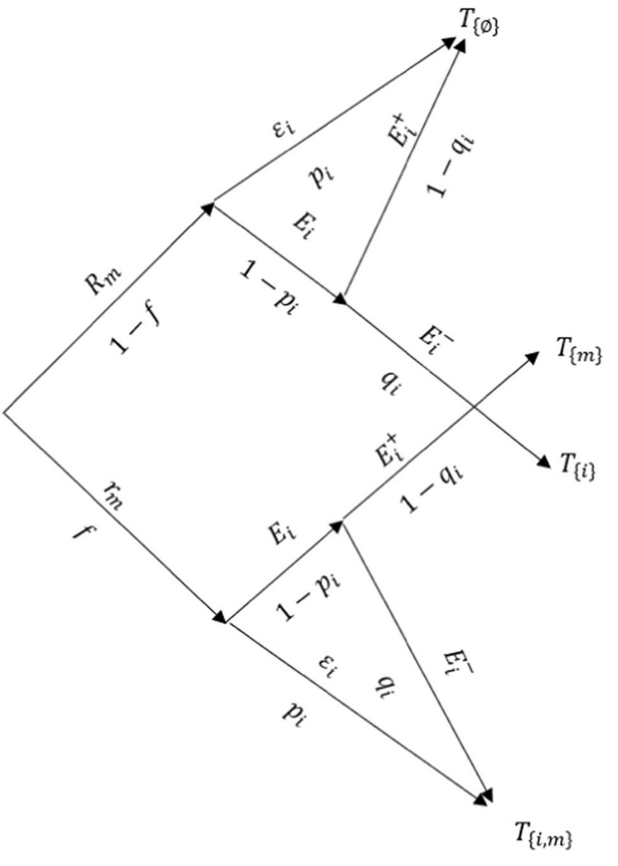

FIGURE 1 The Evolution of Stock Returns. According to our tree model, market excess returns fall below (r_m) or above (R_m) a given threshold with probabilities f and $1 - f$, respectively. Excess returns of asset i follow a SIM with a non-negative β_i according to one out of two possible regimes. In the first regime, $\beta_i > 0$ and the error term is distributed with a "moderate" dispersion (ϵ_i). In this regime, which occurs with probability p_i, stock i's does (not) exceed its threshold whenever the market does (not). In the second regime, which materializes with the complementary probability $1 - p_i$, $\beta_i = 0$ and the error is distributed with a "large" dispersion (E_i). In this case, asset i exceeds its threshold independently of the market with probability q_i due to a large negative error term materializing (E_i^-) or, with probability $1 - q_i$, does not exceed it due to a moderate or large positive error term (E_i^+).

intense interest in the asset pricing implications of the joint tail $T_{\{i,m\}}$ (e.g., Baruník and Nevrla (2023), Bollerslev et al. (2022), Chabi-Yo et al. (2018), and van Oordt and Zhou (2016)). This tail proxies the systematic tail risk which, theoretically, should have important implications for asset pricing despite the mixed empirical findings discussed above.

Similarly important, but to date overlooked, are the remaining two tails. In tail $T_{\{i\}}$ the stock i return exceeds its respective threshold whenever the market return does not. Therefore, this tail captures idiosyncratic tail risk of stock i. The tendency of an asset to exceed its threshold when the market does not is an undesirable property and hence, investors can only be induced to hold this asset if they are compensated with an adequate risk premium. The corresponding theoretical result is rigorously stated in Subsection 2.6. below and proved in Appendix A in the Supplementary Material.

By analogy, tail $T_{\{m\}}$ corresponds to outcomes where the market exceeds its threshold but stock i does not. As a result, this tail captures an important property of stock i: tail risk cushioning (i.e., the tendency of an asset to dampen the losses emanating from the market). Assets that have this property would be in high demand, especially

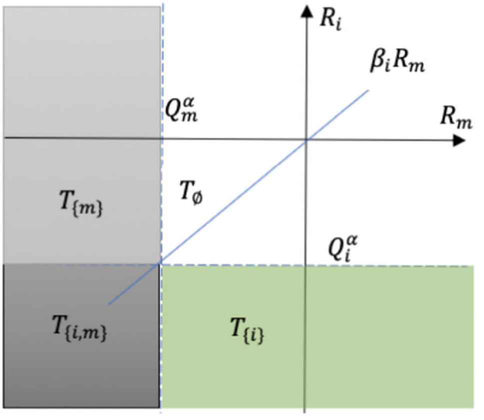

FIGURE 2 The Partition of Outcome Space of Market and Stock Returns. Partition of the two-dimensional outcome space into four joint tails. These joint tails correspond to the final nodes in the event tree depicted in Figure 1: in $T_\emptyset$ no exceedance has occurred (the white area), in $T_{\{m\}}$ the market but not the asset exceeds its threshold (the light grey area), in $T_{\{i\}}$ the asset but not the market exceeds its threshold (the green area), in $T_{\{i,m\}}$ both exceed their respective thresholds (the dark grey area). The dash lines depict the thresholds which in this case correspond to the quantiles $Q_m^\alpha = F_m^{-1}(\alpha)$ and $Q_i^\alpha = F_i^{-1}(\alpha)$.

during periods of market turbulence and would thus, be compensated with lower expected returns. This claim is also rigorously stated in Subsection 2.6. and the proof is given in Appendix A.

Intuitively, stock i's exposure to systematic tail risk can be defined as its tendency (not) to exceed its threshold when the market does (not). Figure 2 illustrates that this situation occurs in the joint tail $T_{\{i,m\}}$ ($T_\emptyset$) where both returns are simultaneously below (above) their respective VaR thresholds. Similarly, idiosyncratic behaviour is displayed whenever stock i diverges strongly from the market (i.e., either stock i or the market exceeds its VaR but not both at the same time). This occurs in joint tail $T_{\{m\}}$ when the market exceeds its threshold but not the asset, and in the joint tail $T_{\{i\}}$ when asset i does exceed its threshold but not the market.

If the prediction of classical finance theory on the reward to risk exposure extends to tail risk, then only exposure to systematic tail risk should earn a risk premium. Idiosyncratic tail risks are supposed to be diversified away and investors would not be compensated with any premia for exposure to such risks. For example, Hwang et al. (2018) find that for portfolios with a small number of stocks, naïve diversification not only outperforms more sophisticated diversification techniques but is also less exposed to tail risk. However, for large portfolios, naïve diversification maintains its superior performance but increases tail risk. Without a clear demarcation and classification of tail risk into systematic and idiosyncratic, it is challenging to understand and interpret these results.

It is important to emphasize that in the foregoing discussion, the meaning of tails can be "expanded" to all outcomes below the median return which effectively modifies tail risk to downside risk (see also Bali et al. (2014)). The three components of risk are still valid and mutually-consistent although they would now represent systematic downside risk, idiosyncratic downside risk, and downside risk cushioning.

2.5 | The Definition of Tail Risk Measures

In this subsection, we formally derive our measures of systematic tail risk, idiosyncratic tail risk, and tail risk cushioning. Define x_0, x_m, x_i and x_{im} as the respective probabilities of the outcomes $T_\emptyset$, $T_{\{m\}}$, $T_{\{i\}}$ and $T_{\{i,m\}}$. In this case, the event tree in Figure 1 leads to a system of linear equations as follows:

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

The last probability $\Pr(T_{\{i,m\}})$, for example, is the sum of the probability $f\cdot p_i$ encapsulating the market exceeding its threshold followed by the asset and the probability $f\cdot(1-p_i)\cdot q_i$ encapsulating the market and the asset exceeding their respective thresholds independently.

In the following discussion, we define the threshold for the asset i equal to $VaR_i^{\alpha_i}$ and for the market equal to $VaR_m^{\alpha_m}$ at the corresponding severity levels α_i and α_m. Then, the probabilities of the tails $\Pr(T_{\{i\}})$ and $\Pr(T_{\{m\}})$ are equal to $x_i = \alpha_i - x_{im}$ and $x_m = \alpha_m - x_{im}$, where $x_{im} = \Pr(T_{\{i,m\}})$. Although not explicitly stated, the probabilities x_i, x_m and x_{im} clearly depend on α_i and α_m.

As the sum of the probabilities of the four collectively-exhaustive and mutually-exclusive outcomes must be one, the following unique solutions for f, q_i and p_i obtain:

$$f = \alpha_m, \tag{1}$$

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

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

The probabilities p_i and q_i are well-defined only if $\alpha_i\alpha_m \le x_{im} < \alpha_m(1+\alpha_i-\alpha_m)$. Note that by construction, $x_{im} \le \alpha_m$ and $x_{im} \le \alpha_i$.

As stock i's excess return r_i closely follows the prediction $\beta_i r_m$ with probability p_i, this probability captures the systematic part of the tail risk of asset i. With the complementary probability $1-p_i$, asset i, independently of the market, either exceeds its threshold or it does not. The former event occurs with probability q_i and captures the idiosyncratic tail risk, while the latter occurs with the complementary probability $1-q_i$ and captures the tail risk cushioning of asset i. Formally, we define:

Systematic Tail Risk (STR):

$$STR_i \equiv STR_i\left(\alpha_i, \alpha_m\right) \equiv p_i = \frac{x_{im} - \alpha_i\alpha_m}{\alpha_m - \alpha_m^2} \tag{4}$$

Idiosyncratic Tail Risk (ITR):

$$ITR_i \equiv ITR_i(\alpha_i, \alpha_m) \equiv (1-p_i)q_i = \frac{x_i}{1-\alpha_m} = \Pr(T_{\{i\}}|T_{\{i\}} \cup T_{\varnothing}) \tag{5}$$

Tail Risk Cushioning (TRC):

$$TRC_i \equiv TRC_i(\alpha_i, \alpha_m) \equiv (1-p_i)(1-q_i) = \frac{x_m}{\alpha_m} = \Pr(T_{\{m\}}|T_{\{m\}} \cup T_{\{i,m\}}). \tag{6}$$

In the context of our model, the Systematic Tail Risk p_i can be interpreted as a coefficient of tail dependence, with values bounded between 0 and 1, that captures joint VaR exceedances by asset i's returns and the market returns. In particular, when $\alpha_m = \alpha_i$ and $p_i = 1$, then the market exceeding its VaR leads always to stock i exceeding its VaR. However, if $p_i = 0$ then VaR exceedances by the market and asset i are independent.

On the other hand, the complementary probability $1 - p_i$ can be decomposed into two parts. The first part, Idiosyncratic Tail Risk, corresponds to asset *i* exceeding its threshold independently of the market and the second part, Tail Risk Cushioning, corresponds to asset *i* not exceeding its threshold independently. We show below that the systematic tail risk measure p_i is similar to that of Chabi-Yo et al. (2018). However, the other two measures are novel in the literature.

It is straightforward to see that measures (4)-(6) are valid and mutually-consistent for any level of alpha below the median (α_m, $\alpha_i \le 50$ percent) although their meaning now generalizes to systematic, idiosyncratic and downside risk cushioning.[4] We return to this important point in the empirical exercises in Section 3. In that section, we also use the fact that STR, ITR, and TRC are computed from probabilities, which allows for their non-parametric estimation and avoids the pitfalls associated with the parametric estimation of tail measures (see, e.g., Frahm et al. (2005) for theory and evidence on that matter).

2.6 | Properties of the Measures of Tail Risk

The tail risk measures that emerge from the framework laid out above have a surprisingly rich array of properties which we elaborate on in this section. Importantly, our tail risk measures are closely related to long-established coefficients of tail dependence. Specifically, the next result shows that the lower (upper) tail dependence coefficient of Sibuya (1960), usually denoted λ_L (λ_U), is a limit case of $STR_i(\alpha_i, \alpha_m)$ when $\alpha_m = \alpha_i = \alpha$. These classic coefficients are of crucial importance in the Extreme Value Theory (EVT) literature (see, e.g., Joe (1997)) and can also be expressed in terms of copulas (see, e.g. Aghakouchak et al. (2013)), which connects STR to this important dependence framework.

Proposition 1. Let F_i and F_m denote the cumulative distribution functions of asset *i* and the market returns, respectively. Then:

$$\lim_{\alpha\to 0} STR_i\Big(\alpha, \alpha\Big) = \lambda_L \equiv \lim_{\alpha\to 0} \Pr\Big(r_i < F_i^{-1}(\alpha)\Big| r_m < F_m^{-1}(\alpha)\Big),$$

$$\lim_{\alpha\to 1} STR_i\Big(\alpha, \alpha\Big) = \lambda_U \equiv \lim_{\alpha\to 1} \Pr\Big(r_i > F_i^{-1}(\alpha)\Big| r_m > F_m^{-1}(\alpha)\Big),$$

Proof.See Appendix A in the Supplementary Material.

By allowing for any value of the severity level α, STR generalizes the classic coefficients of tail dependence to arbitrary severity levels of extreme events. This feature is paramount in empirical studies which rely on multivariate extreme tails because the limited number of observations in these tails make such studies practically infeasible. Moreover, by allowing for cases where $\alpha_m \neq \alpha_i$, $STR(\alpha_i, \alpha_m)$ provides another flexible feature useful in empirical studies.

We note here that naively generalizing λ_L by computing the conditional probability

$$\lambda_L(\alpha) = \Pr\Big\{r_i \le F_i^{-1}(\alpha)\Big| r_m \le F_m^{-1}(\alpha)\Big\} = x_{im}/\alpha, \tag{7}$$

may result in misleading inferences. In particular, when asset *i* is independent of the market, $x_{im} = \alpha^2$ and then $\lambda_L(\alpha) = \alpha$ implies that their dependence increases in the severity level α, while our measure yields $STR_i = 0$ for any severity level α.

[4]Note that if $\alpha_i = \alpha_m = 0.5$, then these measures would bear some resemblance to Bollerslev et al. (2022) semibetas.

Furthermore, STR is also closely related to another classic measure of tail dependence. Huang (1992) proposes the measure $E[\kappa|\kappa \geq 1]$, which is the expected number of tail events given that at least one has occurred (see also Hartmann et al. (2004)). It is straightforward to show that in the bivariate case $E[\kappa|\kappa \geq 1] = \frac{2}{2-p_i}$. Finally, our STR is closely connected to χ-measure proposed by Coles et al. (1999) and used by Poon et al. (2004) to study stock returns. It is again straightforward to verify that in the bivariate lower tails case

$$\lim_{\alpha \to 0} p_i = \chi \tag{8}$$

where $\chi = \lim_{s \to -\infty} \frac{\Pr(S < s, T < s)}{\Pr(S < s)}$ and S and T are two generic random variables with support on the real line. An equivalent result holds in the upper tails.

Another essential aspect of this framework is that our Idiosyncratic Tail Risk and Tail Risk Cushioning measures generalize, respectively, the mixed lower-upper tail dependence coefficient (λ_{LU}) and the mixed upper-lower tail dependence coefficient (λ_{UL}) - see for example, Joe (1997).

Proposition 2

$$\lim_{\alpha \to 0} ITR_i\left(\alpha, 1-\alpha\right) = \lambda_{LU} \equiv \lim_{\alpha \to 0} \Pr\left(r_i < F_i^{-1}(\alpha) \middle| r_m > F_m^{-1}\left(1-\alpha\right)\right),$$

$$\lim_{\alpha \to 0} TRC_i\left(1-\alpha, \alpha\right) = \lambda_{UL} \equiv \lim_{\alpha \to 0} \Pr\left(r_i > F_i^{-1}\left(1-\alpha\right) \middle| r_m < F_m^{-1}(\alpha)\right).$$

Proof. See Appendix A in the Supplementary Material.

As in the case of STR, the Idiosyncratic Tail Risk and Tail Risk Cushioning generalize the mixed tail dependence coefficients to any severity level of tail events. This is an important feature of these measures, especially in the context of asset returns with positive dependence on the systematic factor, which is typically the overwhelming majority of assets. In that case, it may be practically impossible to estimate the tail dependence coefficients in the mixed tails due to the lack of or exceptionally low number of observations in these tails.

The connection between our tail risk measures and the classical tail dependence coefficients is important, and it can be shown theoretically that they impact the expected excess returns. Chabi-Yo et al. (2018) prove in their Theorem 3 that the expected excess return of a risky asset i is an increasing (decreasing) function of its λ_L (λ_U) with the systematic factor (i.e., the market return). The regularity assumptions on the representative investor's utility function necessary for their result are that the first four derivatives of the utility function have altering signs; that is, investors show non-satiation, they are risk-averse, their absolute risk aversion is decreasing (which is equivalent to investors liking skewness), and they are "temperate" (which is equivalent to investors disliking kurtosis). These assumptions hold for a wide class of possible preferences (e.g., constant relative risk aversion preferences). Our Proposition 1 and their Proposition 3 imply then the following corollary.

Corollary 1. *The expected excess return of risky asset i, $E[R_i] - R_f$, increases in $\lim_{\alpha \to 0} STR_i(\alpha, \alpha)$ and decreases in $\lim_{\alpha \to 1} STR_i(\alpha, \alpha)$.*

We can also show that, under the same regularity assumptions on the representative investor's utility function, ITR and TRC have in the limit a similarly unambiguous impact on expected excess returns of risky assets.

Proposition 3. *The expected excess return of risky asset i, $E[R_i] - R_f$, increases in $\lim_{\alpha \to 0} ITR_i(\alpha, 1-\alpha)$ and decreases in $\lim_{\alpha \to 0} TRC_i(1-\alpha, \alpha)$.*

Proof. See Appendix A in the Supplementary Material.

These results suggest that our measures of tail dependence will impact excess returns not only in the limit as the joint tail probability vanishes, but also for moderate values of α, in particular, when these measures of tail risk become measures of downside risk.[5] In Section 3, we rely on measures (4)-(6) - with the limits of α_i and α_m as specified in Corollary 1 and Proposition 3 - to estimate the impact of the different components of tail and downside risk on stock returns. As your theoretical framework allows for the difference between α_i and α_m, in Appendix C of the Supplementary Material, we present further empirical results examining the impact of tail risk components when investors have different tail risk appetites for the market and individual stocks.

3 | EMPIRICAL ANALYSIS

In Section 2.6, we showed theoretically that STR and ITR should have a positive impact whereas TRC should have a negative impact on expected returns. To examine whether these predictions hold in the data, we empirically test them following standard practice in the asset pricing literature. Firstly, we conduct sorting exercises (i.e., examine whether and how expected returns vary in portfolios sorted according to their tail risk components). This analysis, presented in Appendix B to preserve space, though illustrative of the returns seemingly accruing to portfolios with exposure to a particular tail risk component, is only a first step. Because the returns are not properly risk-adjusted or because a particular tail risk measure may be a proxy of or highly correlated with another risk factor, portfolio sorting is generally followed by a number of empirical exercises that try to adjust expected returns for exposure to other risk factors. This is achieved by adding additional factors generally accepted to influence expected returns as control variables in the Fama and MacBeth (1973) regression. This isolates the part of the excess return that is strictly and exclusively linked to the particular risk factor being examined. To this end, in the following sections, we first discuss the data used in our empirical studies and then present the results of the Fama and MacBeth (1973) regressions. Appendix C contains additional empirical exercises that examine the robustness of these results.

3.1 | Data

In our extensive empirical exercises, as is standard in the literature, we use daily and monthly data for all common stocks in the American Stock Exchange (AMEX), National Association of Securities Dealers Automated Quotations (NASDAQ) and New York Stock Exchange (NYSE) markets. Our data is obtained from the Center for Research in Security Prices (CRSP) and covers the period from January 1968 to December 2021. In the empirical exercises, we follow the standard practice and include only stocks with share codes 10 or 11 and with a minimum of two years of data available in every five years. To calculate the Book-to-Market ratios, we obtain the firm accounting data from the CRSP-Compustat Merge database. This results in a sample of 3,278,028 stock-month observations with the average of 5,059 stocks per month although this number varies between 2,149 and 7,932 stocks in each month during the period we examine. Data on the risk-free rate and on the excess market return for the same period are obtained from Kenneth French's online data library.

The tail risk measures that we propose are computed as follows. At the end of each month, we calculate STR, ITR and TRC for a stock using the previous five years of return observations of the market and the stock. We use the lower tail of the actual empirical distribution of excess returns to calculate a non-parametric measure of VaR following the literature (see, for example, Atilgan et al. (2020)). Specifically, VaR is calculated as

[5]Because we focus on the lower part of the asset returns distribution (i.e. when $\alpha \leq 50$ percent), we do not investigate the case when $\lim_{\alpha \to 1} STR_i(\alpha, \alpha)$ and leave it instead for future research.

the α percentile (where $\alpha \in \{0.05, 0.1, 0.2, 0.3, 0.4, 0.5\}$) of the daily excess returns over the past five years as of the end of month t with the restriction that at least 500 non-missing return observations should exist. Having determined the thresholds defined by VaR_i and VaR_m, we then compute the probabilities x_m, x_i and x_{im} of the respective tails $T_{\{m\}}$, $T_{\{i\}}$ and $T_{\{i,m\}}$ as the number of observations that fall into each tail divided by the total number of observations over that period. With the probabilities x_m, x_i and x_{im}, it is then straightforward to obtain the tail risk measures (4)-(6).

In Appendix B of the Supplementary Material, we provide a detailed discussion of descriptive statistics, persistence analysis of the three tail risk measures and portfolio sorting analysis.

3.2 | Fama and MacBeth (1973) Cross-Sectional Regression

In this section, we investigate the impact of tail risk on expected stock returns controlling for other risk factors using the Fama and MacBeth (1973) cross-sectional regression analysis. Specifically, we estimate the factor betas and other risk measures using time series data in the first step, and then, the relation between returns and these variables is estimated in a second step with a cross sectional regression.

Subrahmanyam (2010) highlights that the number of variables shown to predict stock returns in the cross-section is in excess of 50. Controlling for all of these variables is clearly infeasible, and thus, we focus on the most widely used ones in the literature and those that, intuitively, are most likely to be correlated with our tail risk measures (see also Bali et al. (2014).

In Table 1, we report the results of the Fama and MacBeth (1973) cross-sectional regression of monthly excess returns of all listed US stocks on our tail risk measures and also on other canonical measures. Specifically, the excess returns of each stock relative to the T-bill rate over the following month is regressed on the explanatory variables estimated from historical data over the previous five years. We report the time series average of the coefficients estimated monthly for each variable in nine different models. These coefficients capture the premia per unit of risk and are reported with the respective Newey and West (1987) t-statistics (in parentheses).

The regressors in the Models I to III contain varying sets of canonical risk measures including CAPM beta, book-to-market, size, momentum, volatility, illiquidity, coskewness and cokurtosis (see, e.g., van Oordt and Zhou (2016); Bali et al. (2014) and references therein). Book-to-market is measured as the ratio of the book value from the previous fiscal year adjusted for investment tax credits, deferred taxes and preferred shares divided by the market capitalization at the end of the previous calendar year (see, for example, Fama and French (1993)). Size is calculated as the natural logarithm of market capitalization at the end of the previous month. Momentum is the average of previous year returns excluding the last month (see, e.g., Huang et al. (2012)). Volatility is the standard deviation of daily returns. Illiquidity is proxied by average daily illiquidity in the last year, where the latter is calculated as the ratio of the absolute daily return over daily dollar volume (see Amihud (2002)). Coskewness and cokurtosis are computed as in Ang et al. (2006).

Estimates in Models I-III are consistent with results reported in the literature. At the level of individual stocks, the CAPM beta earns a negative or insignificant risk premium when it is calculated from past daily returns (see, e.g., Bali, Engle, and Murray (2016) for an extensive discussion of this finding). Book-to-market is associated with higher expected return and it is highly significant when we include only CAPM beta and size in the regression. However, with additional risk factors included (Model III), book-to-market becomes only marginally significant. Size affects expected returns negatively and is significant. Momentum, illiquidity and cokurtosis are all statistically significant with the signs of the premia consistent with theoretical predictions. Volatility is significantly associated with lower expected returns, reflecting the volatility feedback and leverage effects (see Black (1976); Campbell and Hentschel (1992) among others). Finally, coskewness is not significant, which is probably due to the high level of measurement noise (see, e.g., Bali et al. (2016)).

Models IV to VI include our proposed measures of tail risk (4)-(6) with the limits of α_i and α_m as specified in Corollary 1 and Proposition 3.

Table 1 shows the results calculated at 10% VaR. We note that STR exhibits the expected positive sign and is highly significant. This suggests that investors are rewarded for bearing the systematic tail risk, which is in line with the theoretical predictions laid out in Corollary 1. Importantly, the inclusion of systematic tail risk in the regression does not substantially alter the significance or the magnitude of other coefficients. We conclude therefore that the systematic tail risk captures a distinct risk that is not present in the other canonical factors.

Similarly, the risk premium associated with TRC is also of the expected negative sign, as suggested by Proposition 3, and is statistically significant. Therefore, this finding supports the theoretical prediction that investors are willing to pay a higher price (i.e., accept lower expected returns) to hold stocks with the ability to cushion large losses generated by the market.

However, we find an unexpected result regarding the idiosyncratic tail risk. The risk premium of ITR is negative although only marginally significant. This suggests that investors are not compensated with higher expected returns when investing in stocks with higher idiosyncratic tail risk. This finding seems at odds with the theoretical prediction of Proposition 3 but mirrors the findings on the idiosyncratic volatility of Ang et al. (2006) among others. These results are consistent after controlling for several potential biases including size, illiquidity, idiosyncratic volatility, realized volatility and momentum as suggested in Bali et al. (2014). We also obtain similar results when including all tail risk measures simultaneously in the cross-sectional regressions. The results of these investigations are presented and discussed in Appendix C of the Supplementary Material.

We also investigate the economic significance of the tail risk measures by calculating the annualized change in future returns corresponding to one standard deviation change in the risk measures (the complete results are available upon request). Specifically, one standard deviation increase in STR raises the annualized expected returns by 2.88 percent. This is comparable to the impact of the LTD measure in Chabi-Yo et al. (2018). The corresponding impacts of ITR and TRC are -0.53 and -0.72 percent, respectively. Thus, STR is statistically more significant and has a higher economic impact on stock returns than the other two tail risk components. Our results are also comparable to the cross-sectional analysis of Chabi-Yo et al. (2018) in terms of the average R-squared of the cross-sectional regressions, where the risk factors collectively can explain 6.1 percent of next month stock returns.

The estimation of the tail risk premium might be affected by potential survivorship bias. This might be the case if the estimation sample consists of a large number of surviving companies.[6] In fact, only a small proportion of the firms in our sample have survived throughout the period. From January 1968 to December 2021, there are 21,565 unique stocks, while only 3802 stocks exists at the end of the sample (December 2021), accounting for 17.6 percent of the stock population. Thus, one can safely argue that the impact of survivorship bias in our results is minimal. To validate this, in models VII to IX of Table 1, we present the cross-sectional regression results for stocks in our sample that did not exist in December 2021. Since the majority of firms in the last several years in the sample existed in December 2021, we only obtain the Fama and MacBeth (1973) cross-sectional regression for the January 1968 - December 2011 period, removing the last 10 years of recent data. The results of models VII to IX confirm the robustness of our results with respect to survivorship bias. We observe the same sign and similar statistical significance for the coefficients associated with STR, ITR, and TRC as in the main results in models IV to VI.

3.3 | Time-varying Crash Fears

Chen, Joslin, and Tran (2012) argue that the risk premium for disaster risk increases substantially after a disaster (see also Gennaioli et al. (2015) who propose a theoretical model where investors overstate the fear of a future

[6]We would like to thank an anonymous reviewer for this insightful suggestion.

TABLE 1 Cross-sectional analysis of tail risk and other canonical risk measures.

	I	II	III	IV	V	VI	VII	VIII	IX
Intercept	0.0111	0.0332	0.0469	0.0528	0.0488	0.0499	0.0659	0.0635	0.0639
	(5.2956)	(2.9828)	(6.9334)	(7.7043)	(7.2872)	(7.3407)	(7.8203)	(7.5831)	(7.6096)
Beta	−0.0025	0.0010	0.0018	0.0000	0.0013	0.0012	−0.0013	0.0000	0.0000
	(−1.5627)	(0.4528)	(0.9015)	(0.0008)	(0.6073)	(0.5642)	(−0.5819)	(−0.0068)	(−0.0200)
Size		−0.0015	−0.0022	−0.0026	−0.0022	−0.0023	−0.0033	−0.0030	−0.0030
		(−2.5985)	(−6.4034)	(−7.2989)	(−6.5696)	(−6.6271)	(−7.5339)	(−7.1745)	(−7.2069)
B/M		0.0021	0.0007	0.0007	0.0007	0.0007	0.0005	0.0006	0.0006
		(3.7508)	(1.3601)	(1.2214)	(1.3721)	(1.3980)	(0.6807)	(0.9698)	(1.0177)
Momentum			0.0066	0.0069	0.0065	0.0066	0.0089	0.0084	0.0085
			(4.2319)	(4.4810)	(4.0979)	(4.1322)	(5.6086)	(4.8944)	(5.0596)
Illiquidity			0.0012	0.0012	0.0012	0.0012	0.0005	0.0005	0.0005
			(4.7059)	(4.5713)	(4.7628)	(4.7432)	(5.1385)	(5.2102)	(5.2876)
Real Vol			−0.1954	−0.1752	−0.1869	−0.1885	−0.2214	−0.2355	−0.2368
			(−3.8886)	(−3.4727)	(−3.7649)	(−3.7638)	(−4.0490)	(−4.2918)	(−4.2952)
Coskewness			−0.0014	0.0013	−0.0016	−0.0007	0.0020	−0.0015	−0.0004
			(−0.5026)	(0.4548)	(−0.5710)	(−0.2326)	(0.6197)	(−0.4569)	(−0.1182)
Cokurtosis			0.0024	0.0015	0.0022	0.0023	0.0022	0.0030	0.0030
			(4.2986)	(2.5619)	(4.0020)	(4.0828)	(3.0100)	(4.3485)	(4.3244)
STR				0.0231			0.0268		
				(7.3930)			(6.1953)		
ITR					−0.0126			−0.0181	
					(−1.5058)			(−1.7068)	
TRC						−0.0194			−0.0229
						(−2.1106)			(−2.0324)
R-squared	0.0168	0.0366	0.0603	0.0611	0.0611	0.0610	0.0622	0.0623	0.0623

Fama and MacBeth (1973) average risk premia of the proposed tail risk measures and of the canonical risk measures calculated at α= 10 percent tail threshold (with the corresponding Newey and West (1987) t-statistics in brackets). STR is computed as $STR(\alpha, \alpha)$; ITR is computed as $ITR(\alpha, 1-\alpha)$; TRC is computed as $TRC(1-\alpha, \alpha)$. In each cross-sectional regression, monthly excess return of a stock is regressed on CAPM beta, book-to-market, size, momentum, volatility, illiquidity, coskewness, cokurtosis, and the proposed tail risk measure. Models I to VI use data from January 1968 to December 2021 of all stocks in the market. Models VII to IX use data from January 1968 to December 2011 of stocks not existing in December 2021.

market crash following the occurrence of a tail event). Therefore, we examine the impact of the realization of a market tail event on the three components of tail risk. To that end, and following the literature (see, for example, Bali et al. (2014)), we divide our data set into two subsamples centered around large tail events: the "Post–market crash" subsample containing five years after a market tail event and the "Remaining years" subsample. However, because of the pandemic, our sample of market crashes contains three more days which occurred in March 2020 and has a total of 13 worst days.[7] As these three days occurred towards the end of our sample, we only have just one more year of data after the crash. The results of this analysis are presented in Table 2.

We find that the impact of STR on returns is much stronger in the years subsequent to a market crash. The impact of STR on returns is almost three times as high in the "Post–market crash" subsample with a coefficient for the impact of STR of 0.034 in contrast to a coefficient of 0.016 for the "Remaining years". This finding is similar to that of Chabi-Yo et al. (2018) and supports the theoretical model of Gennaioli et al. (2015). However, whereas TRC is not significant in periods following market crashes, it appears to be relevant in the remaining years with a statistically significant coefficient of -0.023. This suggests that investors are willing to pay a premium for stocks that cushion potential blows emanating from the market during relatively calmer periods, but rather surprisingly, this effect does not appear to be priced following a market crash. While at face value this looks like a paradox, an explanation for it may be that in highly turbulent periods, risk aversion increases sufficiently to deter under-diversified investors from market participation (see, for example, Zhou (2020) and Meister and Schulze (2022) for evidence that household market participation decreases substantially following a market crash). This, in turn, reduces any impact these investors may have on the pricing of stocks. Finally, ITR has a negative but insignificant impact on expected returns in both, the post-market crash and the remaining periods. This result is consistent with the findings of Bali et al. (2014) so it is not entirely surprising.

3.4 | The Impact of Downside Risk on Expected Returns

In Section 2, we argue that the three components of risk are valid and mutually consistent if the meaning of tails is "expanded" to all outcomes below the median return, effectively modifying the tail risk to downside risk (see, for example, Bali et al. (2014)). Therefore, we now examine the premia of the downside (tail) risk measures for α ranging from five to 50 percent and report the results of the Fama and MacBeth (1973) cross-sectional regression in Table 3.

The findings are consistent with those in the main cross-sectional regression. Specifically, the risk premium of systematic downside (tail) risk is positive and highly significant at every threshold level α defining the downside (tail) risk. Similarly the risk premium of downside (tail) risk cushioning is negative and significant with the only exception at five percent severity level. The puzzling negative risk premium associated with idiosyncratic tail risk can also be observed at almost all levels of idiosyncratic downside risk. Indeed, for α ranging from 50 to 20 percent, there seems to be strongly significant evidence of a negative impact of ITR on expected returns. At these high levels of α, ITR is a proxy for idiosyncratic (semi-) volatility rather than tail risk. In this context, this result is not surprising and entirely in line with the findings of Ang et al. (2006), among others, that stocks with high idiosyncratic volatility have very low average returns even after controlling for exposure to aggregate volatility. Interestingly, the ITR risk premium flips back to the positive sign indicated by Proposition 3 at the five percent severity level of tail risk.[8]

Interestingly, we observe that the impact of tail risk components on expected returns becomes slightly weaker when the tail threshold reduces below 10 percent. This is largely true for all three measures but especially for ITR and TRC, suggesting that the impact of downside risk on expected returns becomes weaker when the overall downside risk becomes tail risk. At face value, this finding suggests that investors care more about downside risk rather than tail risk.

In the above analysis, the tail risk measures are estimated with a window of five years of daily data. To check the robustness of the results obtained with this window size, we repeat the investigation and examine the premia of

[7]The market crash dates are: October 19, 1987, October 26, 1987, August 31, 1998, April 14, 2000, September 29, 2008, October 09, 2008, October 15, 2008, November 20, 2008, December 01, 2008, August 08, 2011, March 09, 2020, March 12, 2020, March 16, 2020.

[8]We did not carry out the investigation for one percent tail threshold since the exceptionally low number of observations made the estimation of the measures infeasible.

TABLE 2 Cross-sectional analysis of time-varying tail risk.

	Post-market crash			Remaining years		
	I	II	III	I	II	III
Intercept	0.061	0.053	0.053	0.047	0.046	0.047
	(5.648)	(5.018)	(5.009)	(5.496)	(5.443)	(5.503)
Beta	0.002	0.004	0.004	−0.001	0.000	−0.001
	(0.372)	(0.835)	(0.851)	(−0.612)	(−0.238)	(−0.359)
Size	−0.003	−0.002	−0.003	−0.002	−0.002	−0.002
	(−5.529)	(−4.667)	(−4.699)	(−5.147)	(−4.852)	(−4.893)
B/M	0.000	0.000	0.000	0.001	0.001	0.001
	(0.075)	(0.315)	(0.348)	(1.450)	(1.458)	(1.468)
Momentum	0.003	0.002	0.002	0.009	0.009	0.009
	(0.991)	(0.711)	(0.727)	(6.699)	(6.630)	(6.614)
Illiquidity	0.002	0.001	0.002	0.001	0.001	0.001
	(2.821)	(2.911)	(2.935)	(4.680)	(4.632)	(4.695)
Real Vol	−0.070	−0.086	−0.092	−0.244	−0.253	−0.252
	(−0.780)	(−0.968)	(−1.040)	(−4.223)	(−4.433)	(−4.321)
Coskewness	0.003	0.001	0.001	0.000	−0.003	−0.002
	(0.768)	(0.265)	(0.238)	(−0.002)	(−0.856)	(−0.471)
Cokurtosis	0.001	0.002	0.002	0.002	0.002	0.002
	(1.428)	(2.626)	(2.505)	(2.161)	(3.134)	(3.250)
STR	0.034			0.016		
	(6.163)			(4.681)		
ITR		−0.012			−0.013	
		(−0.788)			(−1.489)	
TRC			−0.014			−0.023
			(−0.857)			(−2.365)
R-squared	0.0569	0.0568	0.0568	0.0638	0.0639	0.0638

This table shows the Fama and MacBeth (1973) average risk premia of canonical risk measures and of the proposed tail risk measures calculated at α = 10 percent tail threshold, along with their corresponding Newey and West (1987) t-statistics (in brackets). These results relate to two subsamples: the "Post–market Crash" subsample containing the five subsequent years after a market tail event and the "Remaining years" subsample. The market tail events are defined as the 13 worst market returns in our sample which occurred on: October 19, 1987, October 26, 1987, August 31, 1998, April 14, 2000, September 29, 2008, October 09, 2008, October 15, 2008, November 20, 2008, December 01, 2008, August 08, 2011, March 09, 2020, March 12, 2020, March 16, 2020. 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 proposed tail risk measure. STR is computed as $STR(\alpha, \alpha)$; ITR is computed as $ITR(\alpha, 1-\alpha)$; TRC is computed as $TRC(1-\alpha, \alpha)$. The sample period is from January 1968 to December 2021.

TABLE 3 Cross-sectional analysis: tail risk premia at different tail thresholds with tail risk measured over a five year horizon.

	STR						ITR						TRC					
alpha	50%	40%	30%	20%	10%	5%	50%	40%	30%	20%	10%	5%	50%	40%	30%	20%	10%	5%
Intercept	0.060	0.059	0.058	0.057	0.053	0.050	0.094	0.083	0.074	0.059	0.049	0.046	0.094	0.083	0.072	0.060	0.050	0.047
	(8.052)	(8.031)	(8.122)	(8.042)	(7.704)	(7.299)	(8.894)	(8.843)	(9.020)	(8.327)	(7.287)	(6.899)	(8.894)	(8.910)	(8.815)	(8.254)	(7.341)	(6.964)
Beta	−0.002	−0.002	−0.002	−0.002	0.000	0.001	−0.002	−0.002	−0.002	0.000	0.001	0.002	−0.002	−0.002	−0.002	−0.001	0.001	0.002
	(−0.781)	(−1.148)	(−1.104)	(−0.781)	(0.001)	(0.579)	(−0.982)	(−1.074)	(−0.821)	(−0.112)	(0.607)	(0.981)	(−0.982)	(−1.077)	(−0.879)	(−0.281)	(0.564)	(0.894)
Size	−0.003	−0.003	−0.003	−0.003	−0.003	−0.002	−0.003	−0.003	−0.003	−0.003	−0.002	−0.002	−0.003	−0.003	−0.003	−0.002	−0.002	−0.002
	(−7.592)	(−7.582)	(−7.722)	(−7.603)	(−7.299)	(−6.833)	(−7.507)	(−7.511)	(−7.566)	(−7.152)	(−6.570)	(−6.378)	(−7.507)	(−7.505)	(−7.382)	(−7.075)	(−6.627)	(−6.457)
B/M	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001
	(1.000)	(0.961)	(0.962)	(1.073)	(1.221)	(1.335)	(0.954)	(0.976)	(1.089)	(1.289)	(1.372)	(1.321)	(0.954)	(1.014)	(1.057)	(1.255)	(1.398)	(1.338)
Momentum	0.007	0.007	0.007	0.007	0.007	0.007	0.007	0.007	0.007	0.007	0.007	0.007	0.007	0.007	0.007	0.007	0.007	0.007
	(4.609)	(4.633)	(4.636)	(4.658)	(4.481)	(4.292)	(4.642)	(4.573)	(4.419)	(4.198)	(4.098)	(4.169)	(4.642)	(4.594)	(4.508)	(4.308)	(4.132)	(4.171)
Illiquidity	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001
	(4.487)	(4.484)	(4.518)	(4.534)	(4.571)	(4.619)	(4.532)	(4.513)	(4.492)	(4.550)	(4.763)	(4.778)	(4.532)	(4.512)	(4.538)	(4.636)	(4.743)	(4.736)
Real Vol	−0.159	−0.150	−0.151	−0.158	−0.175	−0.189	−0.158	−0.149	−0.152	−0.168	−0.187	−0.201	−0.158	−0.149	−0.150	−0.161	−0.189	−0.197
	(−3.048)	(−2.861)	(−2.887)	(−3.036)	(−3.473)	(−3.753)	(−3.037)	(−2.845)	(−2.962)	(−3.332)	(−3.765)	(−4.043)	(−3.037)	(−2.851)	(−2.856)	(−3.116)	(−3.764)	(−3.968)
Coskewness	−0.003	−0.001	0.000	0.001	0.001	0.000	−0.002	−0.002	−0.003	−0.003	−0.002	−0.001	−0.002	−0.002	−0.001	0.000	−0.001	−0.001
	(−0.906)	(−0.483)	(−0.117)	(0.303)	(0.455)	(0.124)	(−0.734)	(−0.842)	(−1.065)	(−0.958)	(−0.571)	(−0.363)	(−0.734)	(−0.554)	(−0.281)	(−0.026)	(−0.233)	(−0.433)
Cokurtosis	0.002	0.002	0.001	0.001	0.001	0.002	0.002	0.002	0.002	0.002	0.002	0.002	0.002	0.002	0.002	0.002	0.002	0.002
	(2.986)	(2.846)	(2.621)	(2.381)	(2.562)	(3.359)	(2.936)	(3.029)	(2.974)	(3.500)	(4.002)	(4.270)	(2.936)	(2.850)	(3.114)	(3.461)	(4.083)	(4.205)
Tail risk	0.034	0.037	0.037	0.034	0.023	0.011	−0.066	−0.060	−0.058	−0.039	−0.013	0.013	−0.066	−0.062	−0.056	−0.046	−0.019	−0.004
	(6.997)	(7.352)	(7.602)	(7.298)	(7.393)	(3.938)	(−7.221)	(−6.877)	(−6.871)	(−4.552)	(−1.506)	(1.652)	(−7.221)	(−7.386)	(−6.582)	(−4.805)	(−2.111)	(−0.551)

This table shows the Fama and MacBeth (1973) average risk premia of canonical risk measures and of the proposed tail risk component measures calculated at five to 50 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 proposed tail risk measure. STR is computed as $STR(\alpha, \alpha)$; ITR is computed as $ITR(\alpha, 1-\alpha)$; TRC is computed as $TRC(1-\alpha, \alpha)$. The sample period is from January 1968 to December 2021.

TABLE 4 Cross-sectional analysis: tail risk premia at different tail thresholds with tail risk measured over a two year horizon.

	STR						ITR						TRC					
alpha	50%	40%	30%	20%	10%	5%	50%	40%	30%	20%	10%	5%	50%	40%	30%	20%	10%	5%
Intercept	0.053	0.054	0.052	0.051	0.048	0.046	0.074	0.071	0.063	0.053	0.046	0.045	0.074	0.068	0.061	0.053	0.047	0.045
	(8.264)	(8.187)	(8.085)	(7.988)	(7.738)	(7.396)	(8.946)	(8.782)	(8.936)	(8.212)	(7.200)	(7.108)	(8.946)	(8.553)	(8.622)	(8.061)	(7.413)	(7.198)
Beta	−0.002	−0.002	−0.002	−0.001	−0.001	0.000	−0.002	−0.002	−0.002	−0.001	0.000	0.000	−0.002	−0.002	−0.002	−0.001	0.000	0.000
	(−0.976)	(−1.349)	(−1.117)	(−0.799)	(−0.391)	(−0.028)	(−1.118)	(−1.300)	(−1.111)	(−0.592)	(0.048)	(0.224)	(−1.118)	(−1.196)	(−1.063)	(−0.649)	(−0.245)	(0.052)
Size	−0.003	−0.003	−0.003	−0.002	−0.002	−0.002	−0.003	−0.003	−0.003	−0.002	−0.002	−0.002	−0.003	−0.003	−0.002	−0.002	−0.002	−0.002
	(−7.717)	(−7.743)	(−7.615)	(−7.603)	(−7.370)	(−6.968)	(−7.696)	(−7.711)	(−7.715)	(−7.246)	(−6.672)	(−6.612)	(−7.696)	(−7.579)	(−7.487)	(−7.109)	(−6.744)	(−6.644)
B/M	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001
	(1.761)	(1.718)	(1.726)	(1.789)	(1.849)	(1.892)	(1.768)	(1.759)	(1.775)	(1.847)	(1.871)	(1.834)	(1.768)	(1.837)	(1.774)	(1.868)	(1.883)	(1.866)
Momentum	0.007	0.007	0.007	0.007	0.007	0.007	0.007	0.007	0.007	0.007	0.007	0.007	0.007	0.007	0.007	0.007	0.007	0.007
	(4.656)	(4.717)	(4.728)	(4.622)	(4.495)	(4.319)	(4.676)	(4.630)	(4.490)	(4.323)	(4.257)	(4.362)	(4.676)	(4.632)	(4.574)	(4.363)	(4.286)	(4.347)
Illiquidity	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001
	(4.291)	(4.259)	(4.264)	(4.395)	(4.450)	(4.541)	(4.307)	(4.283)	(4.275)	(4.440)	(4.723)	(4.672)	(4.307)	(4.299)	(4.329)	(4.458)	(4.618)	(4.607)
Real Vol	−0.182	−0.170	−0.177	−0.183	−0.189	−0.194	−0.180	−0.171	−0.174	−0.182	−0.197	−0.203	−0.180	−0.172	−0.173	−0.178	−0.189	−0.197
	(−3.415)	(−3.203)	(−3.325)	(−3.473)	(−3.638)	(−3.732)	(−3.376)	(−3.205)	(−3.291)	(−3.479)	(−3.777)	(−3.862)	(−3.376)	(−3.238)	(−3.257)	(−3.347)	(−3.590)	(−3.758)
Coskewness	−0.002	−0.001	0.000	0.000	0.001	0.000	−0.002	−0.002	−0.003	−0.002	−0.002	−0.001	−0.002	−0.001	0.000	0.000	0.000	−0.001
	(−0.823)	(−0.389)	(−0.118)	(0.166)	(0.393)	(−0.042)	(−0.684)	(−0.927)	(−1.197)	(−1.092)	(−0.671)	(−0.408)	(−0.684)	(−0.414)	(−0.154)	(0.021)	(−0.067)	(−0.321)
Cokurtosis	0.002	0.002	0.002	0.002	0.002	0.003	0.002	0.002	0.002	0.003	0.003	0.003	0.002	0.002	0.002	0.003	0.003	0.003
	(3.685)	(3.645)	(3.728)	(3.400)	(3.339)	(3.976)	(3.684)	(3.815)	(3.942)	(4.373)	(4.938)	(4.986)	(3.684)	(3.705)	(4.008)	(4.659)	(4.833)	(4.913)

TABLE 4 (Continued)

alpha	STR 50%	STR 40%	STR 30%	STR 20%	STR 10%	STR 5%	ITR 50%	ITR 40%	ITR 30%	ITR 20%	ITR 10%	ITR 5%	TRC 50%	TRC 40%	TRC 30%	TRC 20%	TRC 10%	TRC 5%
Tail risk	0.021	0.025	0.022	0.018	0.014	0.006	−0.040	−0.041	−0.038	−0.025	−0.003	0.009	−0.040	−0.038	−0.036	−0.028	−0.023	−0.015
	(5.417)	(6.388)	(5.729)	(5.536)	(5.457)	(2.975)	(−5.723)	(−5.818)	(−6.108)	(−4.565)	(−0.658)	(1.492)	(−5.723)	(−5.742)	(−5.590)	(−4.026)	(−3.840)	(−2.037)

This table shows the Fama and MacBeth (1973) average risk premia of canonical risk measures and of the proposed tail risk component measures calculated at five to 50 percent tail thresholds, along with their corresponding Newey and West (1987) t-statistics (in brackets). The tail risk measures STR, ITR and TRC are estimated over a 2-year horizon. 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 proposed tail risk measure. STR is computed as $STR(\alpha, \alpha)$; ITR is computed as $ITR(\alpha, 1-\alpha)$; TRC is computed as $TRC(1-\alpha, \alpha)$. The sample period is January 1968 - December 2021.

the tail risk measures at different tail thresholds but now with the tail risk measures estimated with a window of two years of daily data. The results of the Fama and MacBeth (1973) cross-sectional regressions are reported in Table 4. Although the significance of the results has been somewhat reduced, as illustrated by the slightly lower t-statistics, the results are qualitatively similar to those of Table 3. An exception is the risk premium for TRC which is now significant for all levels of α (recall that in the previous exercise, TRC was significant only for α from 50 to 10 percent). Similar to previous results, STR is significantly and positively related to expected returns for all levels of α examined whereas ITR is significantly but negatively related to expected returns for levels of α between 50 and 20 percent. Interestingly and in line with the finding in the previous table, although negative and significant for most levels of α, the ITR risk premium switches to the positive sign (though statistically insignificant) at the five percent severity level of tail risk.

In a similar exercise, but with a different focus, we modify the above investigation and examine the premia of the tail risk measures at different threshold levels α for three months ahead returns. In each Fama and MacBeth (1973) cross-sectional regression, the three months-ahead (i.e., the cumulative excess return of a stock from $t + 1$ to $t + 3$) is regressed on its time-t tail risk measures as well as other canonical risk measures. The results, presented in Table 5, are largely similar to those in Table 3. The difference is that, while the coefficients seem to have increased and in some cases doubled across all three risk measures, the statistical significance has reduced slightly, though generally still well above the usual significance levels. For example, at α = 50 percent, the impact of STR increases from 0.034 on the one-month ahead expected returns to 0.069 on the three-month ahead expected returns (with the respective Newey and West (1987) t-statistics 6.997 and 5.772). Similarly, these coefficients are 0.011 and 0.024 at α = five percent (with Newey and West (1987) t-statistics of 3.938 and 3.043 respectively). Similar observations can be made for ITR and TRC. This suggests that the impact of these downside risk measures may take more than one month to "show up" in expected returns although the reduced precision makes it harder to detect it deep in the tails.

To preserve space, the results of extensive robustness analyses, which are generally along the lines of the results discussed above (although they provide some more nuance as well as some interesting insights) are provided in Appendix C of the Supplementary Material.

4 | CONCLUSION

There are several studies that examine the relationship between (different measures of) systematic tail risk and expected returns. The impact of idiosyncratic tail risk on stock returns, on the other hand, has attracted much less attention.

In this article, we decompose the tail risk of stock returns into systematic and idiosyncratic parts, with the latter being further decomposed into the tendency of a stock to contribute to or dampen tail risk. These three components of tail risk correspond, respectively, to systematic tail risk, idiosyncratic tail risk, and tail risk cushioning of a stock.

In the theoretical part, we propose a simple model of asset returns and show how it leads directly to our decomposition of tail risk. From this model, we derive closed-form measures for the three aforementioned tail risk components that can be empirically estimated. The explicit formulae for the derived measures allow for their detailed studies, and we prove a number of their properties. In particular, we show that STR generalizes the classic lower (upper) tail dependence coefficient of Sibuya (1960) to any level of severity of extreme events. Moreover, we prove that all our measures have in the limit an unambiguous impact on expected excess returns.

In the empirical part, we extensively investigate the impact of systematic and idiosyncratic components on asset returns. We find, in particular, that our measure of systematic tail risk has a considerable impact on stock returns which confirms the findings reported by Chabi-Yo et al. (2018). Moreover, we find evidence that exposure

TABLE 5 Cross-sectional analysis: tail risk premia three months ahead.

	STR						ITR						TRC					
alpha	50%	40%	30%	20%	10%	5%	50%	40%	30%	20%	10%	5%	50%	40%	30%	20%	10%	5%
Intercept	0.124	0.122	0.121	0.118	0.11	0.103	0.195	0.169	0.151	0.123	0.102	0.095	0.195	0.172	0.148	0.124	0.104	0.098
	(6.083)	(6.012)	(6.080)	(5.976)	(5.603)	(5.285)	(7.113)	(6.911)	(7.076)	(6.351)	(5.372)	(4.977)	(7.113)	(7.17)	(6.795)	(6.202)	(5.378)	(5.034)
Beta	−0.001	−0.002	−0.002	−0.001	0.002	0.005	−0.002	−0.002	−0.001	0.002	0.005	0.007	−0.002	−0.002	−0.001	0.001	0.005	0.006
	(−0.171)	(−0.403)	(−0.369)	(−0.191)	(0.393)	(0.768)	(−0.334)	(−0.337)	(−0.169)	(0.263)	(0.807)	(1.123)	(−0.334)	(−0.4)	(−0.233)	(0.182)	(0.786)	(1.035)
Size	−0.006	−0.006	−0.006	−0.006	−0.005	−0.005	−0.006	−0.006	−0.006	−0.005	−0.004	−0.004	−0.006	−0.006	−0.005	−0.005	−0.004	−0.004
	(−5.544)	(−5.478)	(−5.596)	(−5.468)	(−5.106)	(−4.728)	(−5.442)	(−5.418)	(−5.47)	(−5.106)	(−4.568)	(−4.391)	(−5.442)	(−5.499)	(−5.302)	(−5.017)	(−4.603)	(−4.441)
B/M	0.003	0.003	0.003	0.004	0.004	0.004	0.003	0.003	0.004	0.004	0.004	0.004	0.003	0.003	0.004	0.004	0.004	0.004
	(2.248)	(2.232)	(2.230)	(2.301)	(2.399)	(2.484)	(2.203)	(2.251)	(2.332)	(2.445)	(2.482)	(2.446)	(2.203)	(2.251)	(2.296)	(2.414)	(2.498)	(2.467)
Momentum	0.018	0.018	0.018	0.018	0.017	0.017	0.018	0.018	0.017	0.017	0.017	0.017	0.018	0.018	0.018	0.017	0.017	0.017
	(3.850)	(3.878)	(3.903)	(3.947)	(3.798)	(3.671)	(3.881)	(3.812)	(3.695)	(3.533)	(3.489)	(3.62)	(3.881)	(3.852)	(3.806)	(3.647)	(3.537)	(3.61)
Illiquidity	0.003	0.003	0.003	0.003	0.003	0.003	0.003	0.003	0.003	0.003	0.003	0.003	0.003	0.003	0.003	0.003	0.003	0.003
	(4.886)	(4.910)	(4.920)	(4.918)	(4.939)	(4.972)	(4.925)	(4.908)	(4.936)	(5.005)	(5.055)	(5.138)	(4.925)	(4.923)	(4.954)	(4.999)	(5.044)	(5.039)
Real Vol	−0.394	−0.38	−0.380	−0.386	−0.427	−0.453	−0.392	−0.378	−0.385	−0.41	−0.454	−0.486	−0.392	−0.371	−0.377	−0.399	−0.456	−0.473
	(−2.522)	(−2.408)	(−2.414)	(−2.458)	(−2.828)	(−3.027)	(−2.516)	(−2.401)	(−2.486)	(−2.686)	(−3.055)	(−3.288)	(−2.516)	(−2.361)	(−2.39)	(−2.589)	(−3.066)	(−3.205)
Coskewness	−0.005	−0.002	0.000	0.002	0.003	0.002	−0.003	−0.004	−0.006	−0.005	−0.003	−0.002	−0.003	−0.003	−0.001	0.001	−0.001	−0.002
	(−0.570)	(−0.294)	(0.050)	(0.300)	(0.430)	(0.217)	(−0.43)	(−0.518)	(−0.693)	(−0.613)	(−0.359)	(−0.187)	(−0.43)	(−0.321)	(−0.143)	(0.063)	(−0.067)	(−0.207)
Cokurtosis	0.003	0.003	0.003	0.002	0.002	0.003	0.003	0.003	0.003	0.003	0.004	0.004	0.003	0.003	0.003	0.003	0.004	0.004
	(1.850)	(1.702)	(1.640)	(1.368)	(1.504)	(2.174)	(1.837)	(1.924)	(1.853)	(2.196)	(2.543)	(2.846)	(1.837)	(1.75)	(1.974)	(2.114)	(2.633)	(2.745)
Tail risk	0.069	0.075	0.075	0.071	0.049	0.024	−0.138	−0.119	−0.115	−0.081	−0.026	0.039	−0.138	−0.127	−0.115	−0.093	−0.040	−0.006
	(5.772)	(5.750)	(5.816)	(5.439)	(5.876)	(3.043)	(−5.767)	(−5.314)	(−5.228)	(−3.377)	(−1.153)	(1.839)	(−5.767)	(−5.894)	(−5.018)	(−3.330)	(−1.543)	(−0.276)

This table shows the Fama and MacBeth (1973) average risk premia of canonical risk measures and of the proposed tail risk component measures calculated at five to 50 percent tail thresholds, along with their corresponding Newey and West (1987) t-statistics (in brackets) for expected returns three months ahead. STR is computed as $STR(\alpha, \alpha)$; ITR is computed as $ITR(\alpha, 1-\alpha)$; TRC is computed as $TRC(1-\alpha, \alpha)$. In each cross-sectional regression, three months-ahead 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 proposed tail risk measure. The sample period is from January 1968 to December 2021.

of a stock portfolio to tail risk cushioning earns a significant negative risk premium as suggested by our theoretical results. However, evidence suggests that exposure of a portfolio to idiosyncratic tail risk earns a negative risk premium, extending the findings of Ang et al. (2006), among others, on the negative impact of idiosyncratic volatility on expected stock returns to idiosyncratic tail risk. The components of downside risk have an economically similar although statistically stronger impact on expected returns relative to their tail risk counterparts. Our findings on idiosyncratic tail and downside risk add to the existing wealth of results on idiosyncratic risk that contradict theoretical predictions, an issue which clearly deserves further study.

ACKNOWLEDGMENTS

We would like to thank the Editor, Murali Jagannathan and an anonymous Reviewer for comments and suggestions that have greatly helped to improve the paper. We also thank Turan Bali, Bob Dittmar, Chris Polk, Richard Harris, Frank Windmeijer, Jozef Baruník, Fernando Vega Redondo, Jon Danielsson, Enkelejd Hashorva and the seminar and conference participants at University of Bristol, University of East of Anglia, Cambridge University Isaac Newton Institute, Chinese University of Hong Kong, Systemic Risk Centre at LSE, Norwegian School of Economics (NHH), Charles University, University of Lausanne, Bank of England and European Central Bank.

ORCID

Evarist Stoja http://orcid.org/0000-0003-1082-9974

REFERENCES

Adrian, T. & Brunnermeier, M. K. (2016). CoVaR. *The American Economic Review*, *106*(7), 1705.

AghaKouchak, A. Sellars, S. & Sorooshian, S. (2013). *Methods of tail dependence estimation*. Springer.

Amihud, Y. (2002). Illiquidity and stock returns: cross-section and time-series effects. *Journal of Financial Markets*, *5*(1), 31–56.

Ang, A., Hodrick, R. J., Xing, Y., & Zhang, X. (2006). The cross-section of volatility and expected returns. *The Journal of Finance*, *61*(1), 259–299.

Arzac, E. R., & Bawa, V. S. (1977). Portfolio choice and equilibrium in capital markets with safety-first investors. *Journal of Financial Economics*, *4*(3), 277–288.

Atilgan, Y., Bali, T. G., Demirtas, K. O., & Gunaydin, A. D. (2020). Left-tail momentum: underreaction to bad news, costly arbitrage and equity returns. *Journal of Financial Economics*, *135*(3), 725–753.

Bali, T. G. & Cakici, N. (2004). Value at risk and expected stock returns. *Financial Analysts Journal*, *60*(2), 57–73.

Bali, T. G., Cakici, N., & Whitelaw, R. F. (2014). Hybrid tail risk and expected stock returns: when does the tail wag the dog? *The Review of Asset Pricing Studies*, *4*(2), 206–246.

Bali, T. G., Demirtas, K. O., & Levy, H. (2008). Nonlinear mean reversion in stock prices. *Journal of Banking & Finance*, *32*(5), 767–782.

Bali, T. G., Engle, R. F., & Murray, S. (2016). *Empirical asset pricing: The cross section of stock returns*. John Wiley & Sons.

Barberis, N. C. (2013). Thirty years of prospect theory in economics: A review and assessment. *Journal of Economic Perspectives*, *27*(1), 173–196.

Barro, R. J. (2006). Rare disasters and asset markets in the twentieth century. *The Quarterly Journal of Economics*, *121*(3), 823–866.

Baruník, J. & Nevrla, M. (2023). Quantile spectral beta: A tale of tail risks, investment horizons, and asset prices. *Journal of Financial Econometrics*, *21*(5), 1590–1646.

Black, F. (1976). The pricing of commodity contracts. *Journal of Financial Economics*, *3*(1–2), 167–179.

Bollerslev, T., Patton, A. J., & Quaedvlieg, R. (2022). Realized semibetas: disentangling "good" and "bad" downside risks. *Journal of Financial Economics*, *144*(1), 227–246.

Campbell, J. & Hentschel, L. (1992). No news is good news: an asymmetric model of changing volatility in stock returns. *Journal of Financial Economics*, *31*(3), 281–318.

Chabi-Yo, F., Ruenzi, S., & Weigert, F. (2018). Crash sensitivity and the cross section of expected stock returns. *Journal of Financial and Quantitative Analysis*, *53*(3), 1059–1100.

Chen, H., Joslin, S., & Tran, N.-K. (2012). Rare disasters and risk sharing with heterogeneous beliefs. *The Review of Financial Studies*, *25*(7), 2189–2224.

Chen, K. & Sears, R. S. (1984). How many small firms are enough? *Journal of Financial Research*, *7*(4), 341–349.

Chichernea, D. C. & Slezak, S. L. (2013). Idiosyncratic risk premia and momentum. *Journal of Financial Research*, *36*(3), 389–412.

Coles, S., Heffernan, J., & Tawn, J. (1999). Dependence measures for extreme value analyses. *Extremes*, *2*, 339–365.

Conrad, J., Dittmar, R. F., & Ghysels, E. (2013). Ex ante skewness and expected stock returns. *The Journal of Finance*, *68*(1), 85–124.

Dittmar, R. F. (2002). Nonlinear pricing kernels, kurtosis preference, and evidence from the cross section of equity returns. *The Journal of Finance*, *57*(1), 369–403.

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

Fama, E. F. & MacBeth, J. D. (1973). Risk, return, and equilibrium: empirical tests. *Journal of Political Economy*, *81*(3), 607–636.

Frahm, G., Junker, M., & Schmidt, R. (2005). Estimating the tail-dependence coefficient: properties and pitfalls. *Insurance: mathematics and Economics*, *37*(1), 80–100.

François, P. Heck, S. Hübner, G. & Lejeune, T. (2022). Comoment risk in corporate bond yields and returns. *Journal of Financial Research*, *45*(3), 471–512.

Gennaioli, N., Shleifer, A., & Vishny, R. (2015). Neglected risks: the psychology of financial crises. *American Economic Review*, *105*(5), 310–314.

Goetzmann, W. N. & Kumar, A. (2008). Equity portfolio diversification. *Review of Finance*, *12*(3), 433–463.

Hartmann, P., Straetmans, S., & de Vries, C. G. (2004). Asset market linkages in crisis periods. *Review of Economics and Statistics*, *86*(1), 313–326.

Harvey, C. R. & Siddique, A. (2000). Conditional skewness in asset pricing tests. *The Journal of Finance*, *55*(3), 1263–1295.

Hou, K. & Loh, R. K. (2016). Have we solved the idiosyncratic volatility puzzle? *Journal of Financial Economics*, *121*(1), 167–194.

Huang, W., Liu, Q., Rhee, S. G., & Wu, F. (2012). Extreme downside risk and expected stock returns. *Journal of Banking & Finance*, *36*(5), 1492–1502.

Huang, X. (1992). Statistics of bivariate extreme values. Thesis Publishers Amsterdam.

Hwang, I., Xu, S., & In, F. (2018). Naive versus optimal diversification: tail risk and performance. *European Journal of Operational Research*, *265*(1), 372–388.

Joe, H. (1997). *Multivariate models and multivariate dependence concepts*. CRC Press.

Kahneman, D. & Tversky, A. (1979). Prospect theory: an analysis of decision under risk. *Econometrica*, *47*(2), 263–291.

Levy, H. (1978). Equilibrium in an imperfect market: A constraint on the number of securities in the portfolio. *The American Economic Review*, *68*(4), 643–658.

Malkiel, B. G. & Xu, Y. (2004). Idiosyncratic risk and security returns. *Working Paper*. *Available at SSRN*: https://ssrn.com/abstract=255303

Meister, L. & Schulze, K. (2022). How shocks affect stock market participation. *DIW Roundup: Politik im Fokus*, *142*;13. https://www.diw.de/de/diw_01.c.860745.de/publikationen/roundup/2022_0142/how_shocks_affect_stock_market_participation.html

Menezes, C., Geiss, C., & Tressler, J. (1980). Increasing downside risk. *The American Economic Review*, *70*(5), 921–932.

Newey, W. K. & West, K. D. (1987). Hypothesis testing with efficient method of moments estimation. *International Economic Review*, *28*(3), 777–787.

Polkovnichenko, V. (2005). Household portfolio diversification: A case for rank-dependent preferences. *The Review of Financial Studies*, *18*(4), 1467–1502.

Poon, S.-H., Rockinger, M., & Tawn, J. (2004). Extreme value dependence in financial markets: diagnostics, models, and financial implications. *The Review of Financial Studies*, *17*(2), 581–610.

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

Sharpe, W. F. (1963). A simplified model for portfolio analysis. *Management Science*, *9*(2), 277–293.

Sibuya, M. (1960). Bivariate extreme statistics. *Annals of the Institute of Statistical Mathematics*, *11*(2), 195–210.

Stambaugh, R. F., Yu, J., & Yuan, Y. (2015). Arbitrage asymmetry and the idiosyncratic volatility puzzle. *The Journal of Finance*, *70*(5), 1903–1948.

Statman, M. (1987). How many stocks make a diversified portfolio? *Journal of Financial and Quantitative Analysis*, *22*(3), 353–363.

Stoja, E. Polanski, A. Nguyen, L. H. & Pereverzin, A. (2023). Does systematic tail risk matter? *Journal of International Financial Markets, Institutions and Money*, *82*, 101698.

Subrahmanyam, A. (2010). The cross-section of expected stock returns: what have we learnt from the past twenty-five years of research? *European Financial Management*, *16*(1), 27–42.

Supper, H., Irresberger, F., & Weiß, G. (2020). A comparison of tail dependence estimators. *European Journal of Operational Research*, *284*(2), 728–742.

van Oordt, M. R. & Zhou, C. (2016). Systematic tail risk. *Journal of Financial and Quantitative Analysis*, *51*(2), 685–705.
Zhou, J. (2020). Household stock market participation during the great financial crisis. *The Quarterly Review of Economics and Finance*, *75*, 265–275.

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

How to cite this article: Stoja, E., Polanski, A., & Nguyen, L. H. (2025). The taxonomy of tail risk. *Journal of Financial Research*, *48*, 701–724. https://doi.org/10.1111/jfir.12423