International Review of Financial Analysis 86 (2023) 102531

Contents lists available at ScienceDirect

International Review of Financial Analysis

journal homepage: www.elsevier.com/locate/irfa

Correlation versus co-fractality: Evidence from foreign-exchange-rate variances[☆]

Klaus Grobys[a,b,*]

[a] Finance Research Group, School of Accounting and Finance, University of Vaasa, Wolffintie 34, 65200 Vaasa, Finland
[b] Innovation and Entrepreneurship (InnoLab), University of Vaasa, Wolffintie 34, 65200 Vaasa, Finland

ARTICLE INFO

JEL classification:
C22
G12
G13
G14
O10

Keywords:
Foreign exchange rates
Pareto distributions
Power laws
Second moment
Variance
Variance of variance

ABSTRACT

The concept of correlation appears to be the cornerstone of modern finance as it is applied in almost all finance-related research studies. However, Fama (1963) argued that "if the [population] variance is infinite, other statistical tools (e.g., least-squares regression) which are based on the assumption of finite variance will, at best, be considerably weakened and may in fact give very misleading answers" (p. 421). This study shows variances of foreign exchange rates to be governed by power laws with a tail exponent of $\alpha < 3$, suggesting infinite second moments. We derive a new concept to measure dependencies between power-law processes with this tail exponent, which we term *co-fractality*. We show that risk diversification based on the concept of correlation indeed gives misleading results. Notably, foreign-exchange-rate variances lacking co-fractality in our earlier subsample do not show evidence for co-fractality in our later subsample. We argue that co-fractality, as opposed to correlation, should be used to measure the dependency between processes governed by power laws.

"Invariances make life easier. If you can find market properties that remain constant over time or place, you can build better, more useful models and make sounder financial decisions."
(Benoit Mandelbrot, *The (Mis)Behavior of Markets*, p. 242)

1. Introduction

The concept of correlation appears to be the cornerstone of modern finance, as it is used in almost all finance-related research studies. Notably, the Nobel prize for economic-related research has been awarded about half a dozen times for research drawing conclusions based on the statistical concept of correlation. Consider two assets A and B. Denoting the return-generating processes of assets A and B at time t as y_t^A and y_t^B, and provided $y_t^A \sim N(\mu_A, \sigma_A^2)$ and $y_t^B \sim N(\mu_B, \sigma_B^2)$, the correlation of those assets denoted as $\rho(y_t^A, y_t^B)$ is given by

$$\rho\left(y_t^A, y_t^B\right) = \frac{COV\left(y_t^A, y_t^B\right)}{\sqrt{VAR\left(y_t^A\right)}\sqrt{VAR\left(y_t^B\right)}}. \tag{1}$$

It is important to note that the variance is an integral part of this metric. Indeed, variances, covariances, and correlations are used in virtually all finance studies for point estimation, hypothesis testing, portfolio optimization, risk management, etc. A textbook example is portfolio diversification, in which the lack of correlation is considered important for minimizing overall portfolio risk. Modern portfolio theory in the spirit of Nobel prize laureate Markowitz (1952) tells us that the expectation of portfolio P consisting of assets A and B is given by

$$E\left(y_t^P\right) = w\mu_A + (1-w)\mu_B, \tag{2a}$$

where w denotes the weight in terms of percentage of wealth invested in A and μ_A (μ_B) is the expected return of asset A (B). The portfolio variance, denoted as $VAR(y_t^P)$ is, hence, given by

$$VAR\left(y_t^P\right) = w^2 VAR\left(y_t^A\right) + (1-w)^2 VAR\left(y_t^B\right) + 2w(1-w)COV\left(y_t^A, y_t^B\right). \tag{2b}$$

From Eq. (2b), it becomes evident that for any given $VAR(y_t^A)$ and $VAR(y_t^B)$, portfolio variance $VAR(y_t^P)$ decreases as $COV(y_t^A, y_t^B)$ decreases,

☆ The author is thankful for having received useful feedback from two anonymous reviewers.
* Corresponding author at: Finance Research Group, School of Accounting and Finance, University of Vaasa, Wolffintie 34, 65200 Vaasa, Finland.
E-mail address: klaus.grobys@uwasa.fi.

https://doi.org/10.1016/j.irfa.2023.102531
Received 14 September 2022; Received in revised form 30 October 2022; Accepted 10 January 2023
Available online 13 January 2023

and assuming there are no negatively correlated assets, the portfolio risk is minimized if $\rho(y_t^A, y_t^B) = 0$. Such an example can be found in perhaps all standard finance books used by business schools. Unsurprisingly, Ray Dalio, founder of Bridgewater Associates, which manages the world's largest and perhaps most successful hedge fund, calls the search for uncorrelated assets "the holy grail of investing." In his bestselling book *Principles*, Dalio (2017) highlighted, "Having a few good uncorrelated return streams is better than having just one, and knowing how to combine return streams is even more effective than being able to choose good ones (though of course you have to do both). At the time (and still today), most investment managers did not take advantage of this" (p. 57).

It is important to note that correlation-based methods have been intensively used also for measuring co-dependencies between asset market uncertainties in terms of variances or volatilities, respectively. For instance, Da Fonseca and Zhang (2019) used high-frequency data for major volatility indices to compute the volatility of volatility and argued that the correlation between the volatility and the volatility of volatility is positive, consistent with observations in the volatility option market. Moreover, Andersen, Bollerslev, Diebold, and Labys (2001a, 2001b) used high-frequency data on deutschemark and yen returns against the dollar to construct model-free estimates of daily exchange rate volatility and correlation. Using correlation as metric for measuring co-dependencies, they found high correlation between foreign exchange rate volatilities. Furthermore, Boldanov, Degiannakis, and Filis (2016) used the realized volatilities for both oil and stock returns to study the dynamic interdependencies between volatilities. The authors found that the correlation between the volatilities is subject to time-variation.

Even though the concept of correlation appears to be appealing for both academics and practitioners, there is overwhelming evidence that reality looks very different from what is taught in business schools. Manifestations of reality are hedge fund bankruptcies and replication failure of published research. An example for the former has been provided by Taleb (2010), who described the bankruptcy of the Long-Term Capital Management (LTCM) hedge fund, for which Nobel laureates Robert Merton Jr. and Myron Scholes served as members in the board of directors: "… during the summer of 1998, a combination of large events, triggered by a Russian financial crisis, took place that lay outside their models. It was a Black Swan. LTCM went bust and almost took down the entire financial system with it, as the exposures were massive" (p. 288). An example for the latter comes from the study of Hou, Xue, and Zhang (2020), who showed that 82% of cross-sectional asset-pricing phenomena mostly published in top-notch finance journals fail scientific replication.

Paradoxically, these realities should not come as a surprise—in his seminal study, Mandelbrot (1963) became the first to document that cotton price changes do not exhibit a finite variance. Note that if the variance does not exist, *t*-statistics are not defined either, and as a consequence, the concept of correlation falls to pieces. Interestingly, in his article entitled "Mandelbrot and the stable Paretian hypothesis," Eugene Fama (1963), who happened to be both a Nobel laureate (like Merton Jr. and Scholes) and a doctoral student of Benoit Mandelbrot, commented that:

> … the infinite variance assumption of the stable Paretian model has extreme implications. From a purely statistical standpoint, if the population variance of the distribution of first differences is infinite, the sample variance is probably a meaningless measure of dispersion. Moreover, if the variance is infinite, other statistical tools (e.g., least-squares regression) which are based on the assumption of finite variance will, at best, be considerably weakened and may in fact give very misleading answers. (p. 421).

Hence, prominent scholars in financial economics have been aware of this issue for many years. Unsurprisingly, the seminal studies of Mandelbrot (1963) and Fama (1963) resulted in an enormous amount of follow-up research.[1] The corresponding research hypothesis postulating infinite variance is often referred to as the Lévy-stable hypothesis. In this regard, Lux and Alfarano (2016) highlighted the following:

> The Lévy distributions are characterized by an asymptotic power-law behavior of their tails with an index α (called the characteristic exponent) which implies a complementary cumulative density function of returns (denoted by ret. in the following) which in the tails converges to:

$$\Pr(|ret| > x) \approx x^{-a}.$$

> The Lévy hypothesis restricts the power-law for returns to the admissible range of α ∈ (0,2) which indicates the mentioned non-convergence of the second moment (with α < 1 not even the mean would converge). Empirical estimates based upon the Lévy model typically found α hovering around 1.7. (p. 4).

On the other hand, Lux and Alfarano (2016) argued that other studies raised doubts over the validity of the Lévy hypothesis by questioning the stability-under-aggregation property of these estimates, and the pertinent literature gradually converged to the insight of an exponent significantly larger than 2 and mostly close to 3. Moreover, the authors argued that the approximate cubic form of the power law of returns appears to be accepted as a universal feature of practically all types of financial market, from share markets and futures to foreign exchange (FX) and precious metal markets. This finding implies, in turn, a rejection of the time-honored Lévy hypothesis because α ≈ 3 means that the decay of the outer part of the distribution is faster than allowed by this family of distributions.

While earlier studies modeled power-law functions for asset returns, recent studies by Grobys, Junttila, Kolari, and Sapkota (2021) and Grobys (2021) used realized volatility or realized variances estimated by intraday price ranges. For instance, Grobys et al. (2021) explored the volatility processes of stable cryptocurrencies (e.g., stablecoins) and their potential stochastic interdependencies with Bitcoin volatility. In doing so, the authors used realized daily volatilities to model the probability density functions of five stablecoins exhibiting the largest market capitalizations. Using power laws, the authors found the exponents to be statistically significantly below 3 across all stablecoins, thereby rendering stablecoin volatilities statistically unstable. Relatedly, Grobys (2021) used realized daily variances for five key financial asset markets to test the power-law null hypothesis. His findings indicated first that the power-law null hypothesis cannot be rejected for any of those asset markets and second that the exponents are below 3, implying that the variance of variance does not exist for any of those asset markets. In a robustness check, the power-law null hypothesis was investigated using lower frequented data on a monthly frequency. His findings indicate that the economic magnitude for the power-law exponent of the realized monthly variance of the S&P 500 is statistically the same for monthly data as it is for daily data, which strongly supports Mandelbrot (1963) early findings on the scaling property of cotton price changes.

Findings from recent literature have some serious implications. If the

[1] Relevant studies in this stream of research are, for instance, those of Hall, Brorsen, and Irwin (1989), Lau, Lau, and Lau, Lau, and Wingender (1990), Hill (1975), Jansen and de Vries (1991), Lux (1996), Werner and Upper (2002), Cont, Potters, and Bouchaud (1997), and Gopikrishnan, Meyer, Amaral, and Stanley (1998).

variances of asset variances do not exist, we cannot use correlation as a statistical metric to measure the co-movement between the uncertainty of financial assets because the correlation is undefined.[2] Even if we were interested in examining the correlation for simple asset returns and assuming that $\alpha \approx 3$, the research environment would still not meet the conditions allowing us to use correlation or correlation-based methodologies to evaluate data dependencies. Taleb (2020) conjectured that if the fourth moment is undefined, which is the case if $\alpha < 5$, the stability of the second moment is not ensured. Therefore, we cannot work with the variance even if it exists in the theoretical distribution.

The purpose of this study is manifold. First, we test the power-law null hypothesis for realized FX rate variances using Group of Ten (G10) currencies. Matching the data samples, we use the intersection for all data sets. Hence, the overall sample period used in this study is from May 16, 2006 to November 19, 2021. We test the power-law null hypothesis first for the whole sample and then for two subsamples of equal length. The two subsamples contain nonoverlapping daily data and correspond to an earlier subsample from May 16, 2006 to March 07, 2014 and a later subsample from March 10, 2014 to November 19, 2021. The stability of the estimated power-law exponents is studied in a comparative manner across subsamples. We then propose a new test to examine the co-dependencies in the tails of the variance distributions accounting for nonexistent second moments. Provided the tails of the variance distributions are governed by fractal processes with infinite variances, as indicated by earlier research, we term our metric *co-fractality* as opposed to correlation. Deriving the statistical distribution of the test statistic, we test for co-fractality in the overall sample, as well as in our two independent subsamples—thereby examining whether potential co-fractality has changed over time.

Our study has some important contributions. First and most importantly, we propose a new methodology that allows for testing the co-dependencies of power-laws with infinite variances. This is an important issue because of the following reasons: First of all, power-law behavior of financial assets is not a new finding per se. In reviewing Mandelbrot (1963) early study documenting power-law behavior in cotton price fluctuations, already Fama (1963) argued that if the population variance of the distribution of first differences is infinite, the sample variance is a useless measure of dispersion and that other standard statistical tools such least-squares regression, which are based on the assumption of finite variance, will provide very misleading research results. Second, Grobys (2021) argued that "a hypothetical root cause for the high rate of replication failures in financial economics could be that many researchers correctly use incorrect methods, that is, these methods do not work well, given the very nature of financial markets." The high rate of replication failure of finance research studies published in top-notch finance journals is a well-documented issue, subject to investigation in the studies of Hou et al. (2020), Harvey, Liu, and Zhu (2016) and Serra-Garcia and Gneezy (2021). Given the failure of correlation-based methodologies, a methodological approach addressing the prevalent instability-of-variance-problem in financial markets is in urgent need.

Next, there is a large body of research examining the power-law hypothesis for financial markets. An interesting overview on this literature was provided, for instance, by Lux and Alfarano (2016). In view of this literature, our study is related to those of Grobys (2021) and Grobys and Kolari (2022), who tested the plausibility of the power-law null hypothesis using the goodness-of-fit (GoF) test proposed by Clauset, Shalizi, and Newman (2009). Specifically, Grobys (2021) tested the power-law null hypothesis for realized variances of gold, crude oil, the USD/GBP exchange rate, the S&P 500, and Bitcoin. His findings indicate that the power-law null hypothesis of $\alpha < 3$ cannot be rejected for any asset market variance. Similarly, Grobys and Kolari (2022) tested the power-law null hypothesis for realized variances for G10 currencies using daily and weekly data. Their findings indicate that the power-law null hypothesis of $\alpha < 3$ cannot be rejected for most FX rate variances, with a surprising result of this finding being even stronger when using lower frequented data (e.g., weekly data). This result is surprising because Mandelbrot (2008) highlighted that some economists would argue that "the degree of wildness–the fatness of the tails–appeared to diminish as you looked at returns over longer time periods" (p. 219). The studies of Grobys and Kolari (2022) and Grobys (2021) did not find such evidence for realized variances for G10 currencies or the S&P 500, implying strong fractal behavior as their exponents did not appear to alter as the time scale changed.

Finally, in a critical manner Mandelbrot (2008) highlighted the following:

> When you pick a stock by the conventional [correlation-based] method, you may actually be adding risk rather than reducing it. […] A new approach is needed. Today, building a portfolio by the book is a game of statistics rather than intelligence: You start by assuming the market has correctly priced each stock, and so your task is simply to combine the particular stocks in your portfolio in such a way as to meet your investment goals. (Mandelbrot, 2008, p. 266).

Every endeavor begins with a first step and the concept of co-fractality could serve as a more robust tool for managing risks in wild market environments as opposed to correlation-based methods. Hence, from a broader perspective, the present study extends the literature on risk assessment and risk management. There is a large body of finance research related to risk assessment and risk management. Some relevant studies on risk management related to asset management are Alexander (2009), Palomba and Riccetti (2012), or Kaplanski and Levy (2015). Moreover, Aven (2016) provided an interesting review on the relevant literature, pointing out that Taleb (2010) made the "black swan" metaphor well-known and that his work has inspired many authors to examine tail risks. This study extends this wide strand of literature by examining co-dependencies among fractal processes that may exhibit infinite variances.

Our results show that the realized variances for G10 currencies are governed by Paretian tails with power-law exponents varying between $\alpha = 2.25$ and $\alpha = 2.78$. The part of the distributions governed by power laws varies between 4.75% and 31.61%. Clauset et al.'s (2009) goodness of fit test cannot reject the power-law null hypothesis for at least seven out of nine realized variances for G10 currencies. This result strongly supports the findings of Grobys and Kolari (2022), even for a restricted sample where some of the extreme events are neglected. Sample-split tests show that the hypothesis of $\alpha < 3$ cannot be rejected for any realized variance, irrespective of which realized FX rate variance is considered. Strikingly, for four out of nine realized FX rate variances, the estimated α for the later subsample falls into the 95% confidence interval for the corresponding estimated α for the earlier subsample. We interpret this finding as evidence for parameter stability.

Next, our results employing the whole data sample indicate that the variance for the AUD/USD exchange rate is co-fractal in its strong form with the variances for the CAD/USD and SEK/USD exchange rates, whereas the variance for the CHF/USD exchange rate is only co-fractal in its strong form with the variance for the EUR/USD exchange rate. Moreover, the co-fractality matrix shows that only 9 out of 36 variance pairs do not exhibit co-fractality in its weak form, implying that there is not much room for risk diversification. Interestingly, in the earlier subsample, only 7 out of 36 variance pairs do not exhibit co-fractality in its weak form. Strikingly, those seven do not exhibit co-fractality in its weak form in the later subsample either, showing that the absence of co-fractality, as measured in terms of its weak form, appears to be time-

[2] Uncertainty of financial assets can be measured by the volatility or the variance, where the former is the square root of the latter. Modeling variances processes directly, as Grobys et al. (2021) and Grobys (2021), it becomes evident from Eq. (1) that the correlation between variances is undefined if the variances of variances do not exist.

persistent, whereas the presence of co-fractality does not necessarily imply that extreme events in variance pairs coincide in later time periods. Finally, examining the correlations, we find that the absence of correlation in one sample is a poor predictor of noncorrelation in later samples. This should not come as a surprise, given that correlations are to a high degree sample-specific in this research environment (e.g., financial market data).

This study is organized as follows: The next section presents the data. The third section presents the methodology. The fourth section describes the results, whereas the fifth section provides a detailed discussion. The last section concludes.

2. Data

Publicly available daily series for the AUD/USD, CAD/USD, CHF/USD, EUR/USD, GBP/USD, JPY/USD, NOK/USD, NZD/USD, and SEK/USD exchange rates were downloaded from finance.yahoo.com. Because data for the AUD/USD exchange rate are only publicly available from May 16, 2006 onward, we restrict the data samples to cover only data from May 16, 2006 to November 19, 2021. We only use the intersection; that is, we only account for daily data where all FX rates were quoted on the same day, leaving us with 4014 daily observations.[3]

Following Grobys (2021), the annualized daily realized variances for each FX market i are computed using the following Parkinson (1980) estimator:

$$\sigma_{i,t}^2 = T\frac{1}{4ln(2)}\left(ln\left(H_{i,t}\right) - ln\left(L_{i,t}\right)\right)^2, \tag{3}$$

where $H_{i,\,t}$ and $L_{i,\,t}$ denote the highest and lowest price for FX market i on day t, respectively, $\sigma_{i,\,t}^2$ denotes FX market i's corresponding realized annualized variance, and $T = 250$. Table 1 reports the descriptive statistics. As shown in Table 1, kurtosis values vary between 140.57 for the CAD/USD exchange rate and 4014.00 for the EUR/USD exchange rate, strongly suggesting fat tails.

3. Methodology

3.1. Power laws

Following Grobys (2021) and Grobys and Kolari (2022), we model realized FX rate variances using the following probability function, respectively, power-law function:

$$p(x) = Cx^{-\alpha}, \tag{4}$$

where $C = (\alpha - 1)x_{MIN}^{\alpha-1}$ with $\alpha \in \{\mathbb{R}_+|\alpha > 1\}$, x denotes the respective realized FX rate variance provided $x \in \{\mathbb{R}_+|x_{MIN} \le x < \infty\}$, x_{MIN} is the minimum value of realized FX rate variance governed by the power-law process, and α is the magnitude of the corresponding tail exponent.[4] Taleb (2020) argued that tail exponent α of a power-law function captures via extrapolation the low-probability deviation not seen in the data, playing a disproportionately large role in determining the mean. Using power laws as a methodological approach is also in line with the argument of Lux and Alfarano (2016), who stated that "recent research shows that these models provide a very versatile, yet simple framework to model asset returns in a parsimonious way, and they have been found to perform at least as good as the time-honored GARCH models in terms of forecasting volatility, and often outperform the later to some extent" (p. 5). Next, it can be shown that the conditional expectation of the variance, defined in this context as $E[x|x > x_{MIN}]$, is given by

$$E[x|x > x_{MIN}] = \int_{x_{MIN}}^{\infty} x\,p(x)dx = \frac{(\alpha - 1)}{(\alpha - 2)}x_{MIN} \tag{5}$$

and that the second moment, $E[X^2]$, or the variance of the variance, is defined as

$$E\left[X^2|x > x_{MIN}\right] = \int_{x_{MIN}}^{\infty} x^2\,p(x)dx = \frac{(\alpha - 1)}{(\alpha - 3)}x_{MIN}^2. \tag{6}$$

Higher moments of order k are analogously defined as

$$E\left[X^k|x > x_{MIN}\right] = \frac{(\alpha - 1)}{(\alpha - 1 - k)}x_{MIN}^k. \tag{7}$$

From Eqs. (3) and (4), we see that the theoretical conditional mean for the respective realized variance only exists for $\alpha > 2$, whereas the theoretical conditional variance of variance only exists for $\alpha > 3$.

3.2. Maximum likelihood estimation and goodness of fit test

Following White, Enquist, and Green (2008) and Clauset et al. (2009), who concluded maximum likelihood estimation (MLE) as the most accurate procedure for estimating power-law exponents, we estimate tail exponents as

$$\widehat{\alpha} = 1 + N\left(\sum_{i=1}^{N} ln\left(\frac{x_i}{x_{MIN}}\right)\right)^{-1}, \tag{8}$$

where $\widehat{\alpha}$ denotes the MLE estimator, N is the number of observations exceeding x_{MIN}, and other notations are as previously defined. Clauset et al. (2009) noted that determining the corresponding values for α and x_{MIN} is important for accurately estimating probability density functions. From Eq. (8), we see that the MLE estimator depends on the chosen x_{MIN}, and hence, there are different possible MLE estimators from which to select the most accurate. In this regard, Clauset et al. (2009) documented it as common practice to employ the $\widehat{\alpha}/x_{MIN}$ plot and select the value for x_{MIN} beyond which $\widehat{\alpha}$ is stable. Because this procedure is somewhat subjective and can be sensitive to noise or fluctuation in the tail of the distribution, the authors proposed a goodness of fit (GoF) test based on minimizing distance D between the power-law function and the empirical data. First, the Kolmogorov–Smirnov (KS) distance is the maximum distance between the cumulative density functions (CDFs) of the data and the fitted power-law model as defined by

$$D = MAX_{x \ge x_{MIN}}|S(x) - P(x)|, \tag{9}$$

where $S(x)$ is the CDF of the data for the observation with a value of at least x_{MIN}, and $P(x)$ is the CDF for the power-law model that best fits the data in the region $x \ge x_{MIN}$. Estimate $\widehat{x}_{MIN}$ is then the value of x_{MIN} that minimizes D. Using the parameter vector $(\widehat{\alpha}, \widehat{x}_{MIN})$ that optimizes D, Clauset et al.'s (2009) GoF test generates a p-value that quantifies the plausibility of the power-law null hypothesis. Specifically, this test compares D with distance measurements for comparable synthetic data sets drawn from the hypothesized model. The p-value is then defined as the fraction of synthetic distances that are larger than the empirical distance. If we wish to use a significance level of 5%, the power-law null hypothesis is not rejected for p-values exceeding 5% because the difference between the empirical data and the model can be attributed to statistical fluctuations alone. The implementation of this test was detailed by Clauset et al. (2009, p. 675–678).

[3] Comparing Table 1 here with Table 1 of Grobys and Kolari (2022), we see that our matching procedure results in losing 23 observations. Moreover, the maximums for GBP/USD, JPY/USD, and NZD/USD change from 73.5617 (2012-01-27), 0.9140 (1998-10-08), and 14.3882 (2012-01-27) as reported by Grobys and Kolari (2022) to 0.8985 (2016-06-24), 0.4658 (2008-10-24), and 14.3093 (2012−02−01) as reported in the current research.

[4] This study follows the notation by Clauset et al. (2009). To simplify notation, index i denoting the respective realized variance of the individual FX rate is dropped.

Table 1
Descriptive statistics.

Exchange rate	AUD/USD	CAD/USD	CHF/USD	EUR/USD	GBP/USD	JPY/USD	NOK/USD	NZD/USD	SEK/USD
Mean	0.1538	0.0070	0.0104	0.1100	0.0074	0.0076	0.0156	0.0168	0.0132
Median	0.0060	0.0035	0.0039	0.0018	0.0039	0.0035	0.0075	0.0066	0.0065
Maximum	282.2602	0.3105	9.9818	428.3998	0.8985	0.4658	2.4978	14.3093	1.2027
Date of maximum	2006-12-25	2008-10-29	2015-01-15	2008-03-17	2016-06-24	2008-10-24	2020-03-20	2012-02-01	2020-03-19
Minimum	1.0884E-04	1.2608E-05	1.4467E-07	8.8948E-06	1.4375E-05	8.9947E-06	8.6650E-05	3.3508E-05	1.2262E-04
Std.Dev.	6.2894	0.0127	0.1590	6.7617	1.0905	0.0178	0.0583	0.2272	0.0298
Skewness	44.7806	8.8695	61.5019	63.3560	28.1440	13.4117	28.8911	62.0796	19.3949
Kurtosis	2004.3740	140.5692	3852.7490	4013.9950	1218.9120	274.4099	1053.1140	3905.2760	665.2244
Start of the sample	2006-05-16	2006-05-16	2006-05-16	2006-05-16	2006-05-16	2006-05-16	2006-05-16	2006-05-16	2006-05-16
End of the sample	2021-11-19	2021-11-19	2021-11-19	2021-11-19	2021-11-19	2021-11-19	2021-11-19	2021-11-19	2021-11-19
Total observations	4014	4014	4014	4014	4014	4014	4014	4014	4014

This table reports the descriptive statistics for the annualized daily realized variance for AUD/USD, CAD/USD, CHF/USD, EUR/USD, GBP/USD, JPY/USD, NOK/USD, NZD/USD, and SEK/USD. The annualized daily realized variances for each FX market i are in line with those of Parkinson (1980), computed as

$$\sigma_{i,t}^2 = T\frac{1}{4ln(2)}\left(ln\left(H_{i,t}\right) - ln\left(L_{i,t}\right)\right)^2,$$

where $H_{i,\,t}$ and $L_{i,\,t}$ denote the highest and lowest price for FX market i on day t, $\sigma_{i,\,t}^2$ denotes FX market i's corresponding realized annualized variance, and $T = 250$. Publicly available daily data on FX rates against the USD were retrieved from finance.yahoo.com.

3.3. Defining the concept of co-fractality: strong co-fractality versus weak co-fractality

Clauset et al. (2009) argued that "in practice, few empirical phenomena obey power laws for all values of x. More often the power law applies only for values greater than some minimum x_{MIN}. In such cases we say that the *tail* of the distribution follows a power law" (p. 662). Indeed, Grobys (2021) and Grobys and Kolari (2022) showed that 7.04% to 27.52% (12.70% to 36.27%) of daily (weekly) observations of realized FX rate variances are governed by power-law functions. Moreover, both studies found that the null hypothesis of $\alpha < 3$ cannot be rejected for virtually all realized FX rate variances, irrespective of which data frequency is considered. From Eq. (6), we know that the variance does not exist for $\alpha < 3$, and hence, the correlation between pairs of realized variances is not defined either; see Eq. (1). This also makes intuitive sense for the following reasons. First of all, and as noted by Taleb (2020), even if the theoretical mean of a variance exists, we are not allowed to use it because we do not observe it in finite samples. Second, measuring the correlation between variances requires the variances of the variances, not the means of the variances. If $\alpha < 3$, the variance of variance is different in every sample we consider, and we do not observe the true value of the variance's theoretical mean. For these reasons, using correlations or correlation-based metrics is not only useless, but gives inevitably misleading results, as already pointed out by Fama (1963). However, we know that extreme events occur in the tails of the distributions of the FX variances. Specifically, extreme events may occur for $x \geq x_{MIN}$, but we do not know "how extreme" the realizations in future samples look as the variance of variance does not converge.

However, we can interpret the part of the distributions governed by power laws as "power-law regimes" and can define binary variables that have values equal to 1 whenever the corresponding distribution is in a regime governed by a power law; otherwise, the values are equal to 0. For instance, if an FX rate variance has M observations in the power-law regime and N observations in total, $M/(N\text{-}M)$ is then the part of the distribution governed by the power-law process and $(1–M/(N\text{-}M))$ is the part governed by a supposedly thin-tailed distribution. Because $(1–M/(N\text{-}M)) + M/(N\text{-}M) = 1$, we can interpret these figures as empirical probabilities for these two distinct regimes. Based on these elaborations, we can define the concept of co-fractality as follows.

Definition 1. Co-fractality in its strong form. Let us define two financial assets A and B whose variances consist of two different distributional components. Then, the distributions of the realizations of the variances at time t, denoted as x_t^A and x_t^B, are given by ${}_t^A \sim \chi^2(1)$ with probability p_A, $x_t^A \sim PL(\alpha_A)$ with probability $(1 - p_A)$,

$x_t^B \sim \chi^2(1)$ with probability p_B, and. $x_t^B \sim PL(\alpha_B)$ with probability $(1 - p_B)$, where p_A and p_B denote the probabilities that x_t^A and x_t^B, respectively, are realizations of the $\chi^2(1)$, and $PL(\alpha_A)$ and $PL(\alpha_B)$ denote power laws with exponents α_A and α_B, provided $\alpha_A > 1$ and $\alpha_B > 1$. Defining the Tx1 binary vectors $\mathbf{x}^{A*}$and $\mathbf{x}^{B*}$ that have values of 1 if $x_t^A \geq x_{MIN}^A$ or $x_t^B \geq x_{MIN}^B$, and values of 0 if $x_t^A < x_{MIN}^A$ or $x_t^B < x_{MIN}^B$, the processes x_t^A and x_t^B are strongly co-fractal if the following condition is satisfied: $\lambda^s = \frac{\mathbf{x}^{A*'}\mathbf{x}^{B*}}{MAX(\mathbf{x}^{A*'}\mathbf{1}, \mathbf{x}^{B*'}\mathbf{1})} > 0.5$, where $\mathbf{1}$ is a Tx1 vector of ones and λ^s defines the co-fractality coefficient (CFC) measuring strong co-fractality.

Comment 1: From Definition 1 it becomes evident that $\lambda^s = 1$ can only be satisfied if both processes x_t^A and x_t^B exhibit the same number of observations governed by a power law—that is, $(1 - p_A) = (1 - p_B)$—and those realizations are generated at the same points in time t. If $(1 - p_A) \neq (1 - p_B)$, co-fractality in its strong form cannot exist. For instance, if $(1 - p_A) > (1 - p_B)$, more realizations of x_t^A are governed by a power law than of x_t^B. Even if all realizations of x_t^B for which $x_t^B \sim PL(\alpha_B)$ is satisfied coincide with realizations x_t^A for which $x_t^A \sim PL(\alpha_A)$ is satisfied, x_t^A and x_t^B cannot be co-fractal in its strong form.

Comment 2: Note that Grobys (2021) argues that when standard assumptions of finance theory hold, the variance of standardized financial returns must follow a $\chi^2(1)$-distribution. As a consequence, we assume that one part of the variance distribution is in line with standard finance theory, whereas the other part of the distribution obeys a power-law process which is in line with Mandelbrot (1963) (and subsequent research studies confirming Mandelbrot's finding that financial data are governed by power-law processes).

Definition 2. Co-fractality in its weak form.

Let us define two financial assets A and B whose variances consist of two different distributional components. Then, the distributions of the realizations of the variances at time t, denoted as x_t^A and x_t^B, are given by $x_t^A \sim \chi^2(1)$ with probability p_A, $x_t^A \sim PL(\alpha_A)$ with probability $(1 - p_A)$, $x_t^B \sim \chi^2(1)$ with probability p_B, $x_t^B \sim PL(\alpha_B)$ with probability $(1 - p_B)$, where p_A and p_B denote the probabilities that $x_t^A$and x_t^B, respectively, are realizations of the $\chi^2(1)$, whereas $PL(\alpha_A)$ and $PL(\alpha_B)$ denote a power law with exponents α_A and α_B, provided $\alpha_A > 1$ and $\alpha_B > 1$. Defining the Tx1 binary vectors $\mathbf{x}^{A*}$and $\mathbf{x}^{B*}$ that have values of 1 if $x_t^A \geq x_{MIN}^A$ or $x_t^B \geq x_{MIN}^B$, and values of 0 if $x_t^A < x_{MIN}^A$ or $x_t^B < x_{MIN}^B$, the processes x_t^A and x_t^B are weakly co-fractal if the following condition is satisfied:

$$\lambda^{w} = \frac{\mathbf{x}^{A^{*\prime}}\mathbf{x}^{B^{*}}}{MIN(\mathbf{x}^{A^{*\prime}}\mathbf{1}, \mathbf{x}^{B^{*\prime}}\mathbf{1})} > 0.5$$

, where $\mathbf{1}$ is a Tx1 vector of ones and λ^{w} defines the co-fractality coefficient (CFC) measuring weak co-fractality. Comment: Unlike co-fractality in its strong form, from Definition 2 it becomes evident that $\lambda^{W} = 1$ can also be satisfied if $(1 - p_A) \neq (1 - p_B)$. For instance, if $(1 - p_A) > (1 - p_B)$, more realizations of x_t^A are governed by a power law than of x_t^B. However, if all realizations of x_t^B for which $x_t^B \sim PL(\alpha_B)$ is satisfied coincide with realizations of x_t^A for which $x_t^A \sim PL(\alpha_A)$ is satisfied, x_t^A and x_t^B are co-fractal in its weak form. From Definitions 1 and 2 it becomes evident that strong co-fractality subsumes it weak form.

Next, the question arises—how can we test the significance of co-fractality? Let us consider first the case of strong co-fractality. We can interpret the fraction $\frac{x^{A^{*\prime}}x^{B^{*}}}{MAX(x^{A^{*\prime}}1, x^{B^{*\prime}}1)}$, which measures the coincides of two random variables in the power-law regime in relative terms, as stemming from a Bernoulli-distributed random variable with $MAX(\boldsymbol{x}^{A*\prime}\mathbf{1}, \boldsymbol{x}^{B*\prime}\mathbf{1})$ drawings. Next, let us assume that $\frac{x^{A^{*\prime}}x^{B^{*}}}{MAX(x^{A^{*\prime}}1, x^{B^{*\prime}}1)} = \widehat{p}_s$ and $MAX(\boldsymbol{x}^{A*\prime}\mathbf{1}, \boldsymbol{x}^{B*\prime}\mathbf{1}) = n_s$. Then, if $\widehat{p}_s < 0.5$, we cannot argue that the variances of assets A and B are co-fractal in its strong form because <50% of the power-law observations generated in the variance process of one asset coincide with power-law observations in the variance process of the other asset. In other words, if $\widehat{p}_s < 0.5$ for our Bernoulli-distributed random variable "contemporary coincides in power law regimes", coincides in power law regimes can be regarded a *matter of chance*. In a Bernoulli experiment such as "tossing a coin", a coin is regarded as *unfair* if the relative fraction of one side is statistically significantly larger than 0.5. The notion of an unfair coin in a Bernoulli experiment is in the current study regarded as *co-dependency* in terms of contemporary coincides. This is indeed a very simple, yet very intuitive and powerful definition of co-dependency. Therefore, the variances of assets A and B are co-fractal in its strong form *if and only if* $\widehat{p}_s > 0.5$. Furthermore, testing for statistical significance, we can define test statistic t_s as

$$t_s = \frac{\sqrt{n_s}(\widehat{p}_s - p)}{\sqrt{\widehat{p}_s(1 - \widehat{p}_s)}},$$

where $p = 0.5$. Again, the rationale is here that we are interested in analyzing whether or not coincides in power-law regimes can be regarded a matter of chance. Hence, if $\widehat{p}_s$ is statistically significantly larger than $p = 0.5$, coincides in power-law regimes cannot be regarded a matter of chance.

By the central limit theorem (CLT), as $p = E(x_i)$ if $x_i \sim Bernoulli$, $\widehat{p}_s = \sum_i x_i / n_s$ and assuming random sampling,

$$\sqrt{n_s}(\widehat{p}_s - p) \underset{d}{\rightarrow} N(0, VAR(x_i)\Big).$$

But $VAR(x_i) = p(1 - p)$ for such a Bernoulli-distributed random variable, so the test statistic converges in distribution to

$$\sqrt{n_s}\,\frac{(\widehat{p}_s - p)}{\sqrt{p(1 - p)}} \underset{d}{\rightarrow} N(0, 1).$$

By the law of large numbers (LLN), $\widehat{p}_s \underset{d}{\rightarrow} p$; hence, we can use $\widehat{p}_s$ to replace the true $p(1 - p)$ because a consistent estimator of this quantity does not alter the asymptotic distribution. Analogously, evaluating co-fractality in its weak form, we assume that $\frac{x^{A^{*\prime}}x^{B^{*}}}{MIN(x^{A^{*\prime}}1, x^{B^{*\prime}}1)} = \widehat{p}_w$ and $MIN(\boldsymbol{x}^{A*\prime}\mathbf{1}, \boldsymbol{x}^{B*\prime}\mathbf{1}) = n_w$. Then, if $\widehat{p}_w < 0.5$, we cannot argue that the variances of assets A and B are co-fractal in its weak form because <50% of the power-law observations generated in the variance process with *fewer* observations coincide with power-law observations of the variance process with *more* observations in the power-law regime. Having defined $\widehat{p}_w$ and n_w, the implementation of the corresponding test statistic is analogously as derived above.

Figs. 1a – d visualize the concept of co-fractality. The figures show two distributions with data-generating processes evolving form $t = 0$ to $t = T$. In the areas that are marked as grey, the distributions are governed by distribution-specific power-law processes. The intersection between $t = 1$ and $t = 2$ marks the time span where both distributions are in power-law regimes. Fig. 1a visualizes co-fractality in its strong form. We see that the intersection of observations governed by power-law processes covers >0.5 of the observations in power-law regimes for both distributions. Hence, coincident realizations in power law regimes cannot be regarded as matter of chance (e.g., "throwing a fair coin" in a Bernoulli experiment). Next, Fig. 1b visualizes co-fractality in its weak form. We see that the intersection of observations governed by power-law processes covers >0.5 of the observations in the power-law regime for one distribution (i.e., the distribution with less realizations in the power-law regime). Finally, Fig. 1c and d visualize the absence of co-fractality. We see that irrespective of which pair of distributions is considered, the intersection of observations governed by power-law processes covers <0.5 of the observations in power-law regimes for both distributions. Therefore, coincident realizations in power-law regimes can be regarded as a matter of chance.[5]

4. Results

4.1. Evidence for co-fractality among foreign-exchange-rate variances

Table 2 reports the results from estimated power-law models for our variances for G10 currencies using the whole data set from May 16, 2006 to November 19, 2021. We observe from Table 2 that the exponents for all realized FX variances are at least two standard deviations below 3. The realized variance for AUD/USD (EUR/USD) exhibits the lowest (largest) economic magnitude, corresponding to 2.25 (2.78). Moreover, the realized variance for AUD/USD (NOK/USD) is governed to the largest (lowest) extent by a power-law process because 31.61% (4.76%) of the overall realized variance distribution is governed by a Paretian tail. Clauset et al.'s (2009) GoF test shows that we cannot reject the power-law null hypothesis for seven out of nine realized FX variances. The GoF test rejects the power-law null hypothesis only for the realized CAD/USD and SEK/USD variances.

Our results are in line with those of Grobys and Kolari (2022), who reported very similar findings even though the vast majority of their realized FX rate variance samples covers more data by a substantial margin. Whereas the authors also found the power-law null hypothesis to be rejected for the realized SEK/USD variance, they did not reject the power-law null hypothesis for the realized CAD/USD variance using a longer sample period from September 17, 2003 to November 19, 2021. Hence, the discrepancy in the results is a matter of data used in the models. This is, however, in line with the results of Taleb (2010), who argued that a long time series is needed for some fractal processes to reveal their properties. A key characteristic of fractal processes is the occurrence of low-probability events that have a high impact. Per definition, it may take a long time until such events occur and, hence, become visible in the data. That means that at the same time, we cannot rule out that the realized SEK/USD variance is not governed by a power law. For instance, the maximum realized variance in the SEK/USD, observed on March 19, 2020, corresponds to a 45-sigma event. Mandelbrot (2008) similarly pointed to a 22-sigma event: "… on October 19, 1987, the worst day of trading in at least a century, the index fell 29.2 percent. The probability of that happening, based on the standard reckoning of financial theorists, was less than one in 10^{10} – odds so small they have no meaning. It is a number outside the scale of nature. You

[5] Note that the Figures 1a – d illustrate this issue in a simplified manner because in empirical time series data, power-law regimes may occur in reoccurring clusters which is also referred to as *multifractality*.

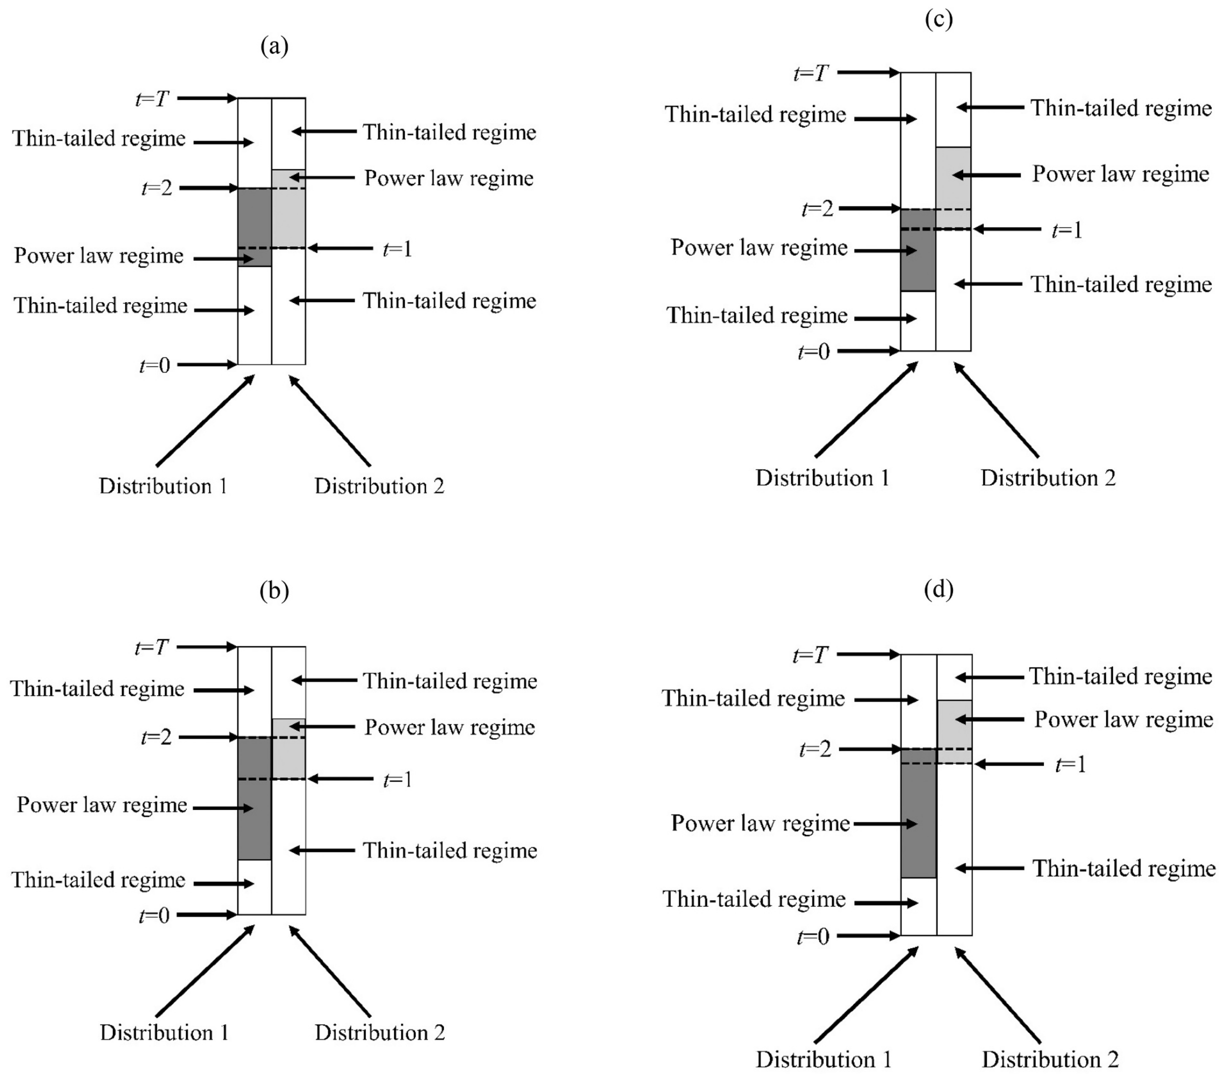

Fig. 1. Visualizing the concept of co-fractality.
Figs. 1a – d visualize the concept of co-fractality. The figures show two distributions with data-generating processes evolving form $t = 0$ to $t = T$. In the areas that are marked as grey, the distributions are governed by distribution-specific power-law processes. The intersection between $t = 1$ and $t = 2$ marks the time span where both distributions are in power-law regimes. Fig. 1a visualizes co-fractality in its strong form. Fig. 1b visualizes co-fractality in its weak form. Fig. 1c and d visualize the absence of co-fractality.

could span the powers of ten from the smallest subatomic particle to the breadth of the measurable universe – and still never meet such a number" (p. 4). Using the same logic, if the realized SEK/USD variance were governed by some standard distribution (e.g., $\chi^2(1)$-distribution), this event would have never happened. Only power laws allow for this kind of low-probability event that has an enormous impact.

Next, Panel A of Table 3 reports the total number of observations in the power-law regime for each realized FX variance and the overall sample. We see that the realized variance for AUS/USD (NOK/USD) has the highest (lowest) number of realized variance observations governed by a power law. Panel B of Table 3 reports the number of coinciding observations. We see that 777 variance realizations exceed $x_{\min} = 0.0097$ for realized AUS/USD variance and $x_{\min} = 0.0066$ for realized CAD/USD variance at the same point in time. The co-fractality matrix in Panel C of Table 3 shows that a fraction corresponding to 0.61 of the 1268 observations of the realized AUS/USD variance governed by a power law coincide with observations in the realized CAD/USD variance that are time-congruently in the power-law regime. Because realized AUS/USD variance is, according to Panel C of Table 3, strongly co-fractal with the realized variances for CAD/USD and SEK/USD, combining these currencies means that the currency portfolio may be exposed to extreme events at the same time and, as a consequence, may deliver poor risk diversification when it may be needed most.

Further, Panel D of Table 3 reports the co-fractality matrix in its weak form. From our definitions, it may be clear that strong co-fractality subsumes co-fractality in its weak form. Elaborating on our previous example, we see here that a fraction corresponding to 0.69 of the 1127 observations for the realized CAD/USD variance governed by a power law coincide with observations in the realized AUS/USD variance obeying a power law time-congruently. Weak-form co-fractality relates the coinciding observations between two variances to the data set exhibiting fewer observations. Because fewer observations imply less exposure to extreme events, we define this metric consequently as the weak form of co-fractality. While only three realized variance pairs exhibit co-fractality in its strong form, we see from Panel D of Table 3 that only 9 out of 36 realized variance pairs do not exhibit co-fractality in its weak form, implying that there is not much room for risk diversification.

Next, we split the overall subsample into two subsamples of equal length. The first (earlier) subsample is from May 16, 2006 to 2014 March

Table 2
Estimated power-law functions.

Distribution	$\widehat{\alpha}$	95% CI	$\left\lvert\frac{(3-\widehat{\alpha})}{\widehat{\sigma}}\right\rvert$	$\widehat{x}_{MIN}$	GoF test (*p*-value)	N	N_{PL}
AUD/USD	2.2515	[2.1830; 2.3204]	21.30	0.0097	0.9030	4014	31.61%
CAD/USD	2.3893	[2.3082; 2.4704]	14.76	0.0066	0.0000	4014	28.10%
CHF/USD	2.7785	[2.6102; 2.9468]	2.58	0.0159	0.6380	4014	10.69%
EUR/USD	2.7807	[2.6090; 2.9524]	2.50	0.0072	0.8340	4014	10.29%
GBP/USD	2.6085	[2.4638; 2.7532]	5.30	0.0113	0.4630	4014	11.83%
JPY/USD	2.5735	[2.4394; 2.7076]	6.23	0.0131	0.4900	4014	13.18%
NOK/USD	2.7414	[2.4942; 2.9886]	2.05	0.0408	0.3640	4014	4.75%
NZD/USD	2.5614	[2.4006; 2.7222]	5.34	0.0294	0.6710	4014	9.02%
SEK/USD	2.3455	[2.2671; 2.4239]	16.36	0.0119	0.0000	4014	28.18%

This table reports the estimates for power-law models $p(x) = (\alpha - 1)x_{MIN}^{\alpha-1}x^{-\alpha}$ using MLE. Tail exponent α is estimated as

$$\widehat{\alpha} = 1 + N\left(\sum_{i=1}^{N} ln\left(\frac{x_i}{x_{MIN}}\right)\right)^{-1},$$

where $\widehat{\alpha}$ denotes the MLE estimator, and N denotes the number of observations, provided $x_i \geq x_{MIN}$. In this model, estimate $\widehat{x}_{MIN}$ is assessed via the KS statistic D, which is the maximum distance between the CDFs of the data and the fitted model:

$$D = MAX_{x \geq x_{MIN}} |S(x) - P(x)|,$$

where $S(x)$ is the CDF of the data for the observation with a value at least x_{MIN}, and $P(x)$ is the CDF for the power-law model that best fits the data in the region $x \geq x_{MIN}$. The estimate of the $\widehat{x}_{MIN}$ is the value of x_{MIN} that minimizes D. Clauset et al.'s (2009) GoF test generates a *p*-value that quantifies the plausibility of the hypothesis. This test compares D with distance measurements for comparable synthetic data sets drawn from the hypothesized model, and the *p*-value is defined as the fraction of synthetic distances that are larger than the empirical distance. Given a significance level of 5%, the power-law null hypothesis is not rejected because the difference between the empirical data and the model can be attributed to statistical fluctuations alone. The implementation of this test was detailed by Clauset et al. (2009, p. 675–678). N_{PL} denotes the percentage of sample observations governed by a power-law process, whereas CI denotes the confidence interval for $\widehat{\alpha}$. The sample period used in this table is from 2006 - 05-16 to 2021-11-19.

07, 2014, whereas the second (later) subsample is from March 10, 2014 to November 19, 2014. For both subsamples, we again estimate power-law models and employ Clauset et al.'s (2009) GoF test to test successively the power-law null hypothesis across all nine realized FX rate variances. The results are reported in Tables 4 and 5. We observe from Table 4 that for eight out of nine power-law exponents, the estimates for α are below 3; moreover, for all realized variances, the 95% confidence intervals for the power-law exponents indicate that $\alpha \leq 2.7949 < 3$ cannot be rejected for any realized FX rate variance. Furthermore, in line with the results for the overall sample, for two realized FX rate variances, the GoF rejects the power-law null hypothesis. From Table 5 we observe that for seven out of nine power-law exponents, the estimates for α are below 3, and for all realized variances, the 95% confidence intervals for the power-law exponents indicate that $\alpha \leq 2.9291 < 3$ cannot be rejected for any realized FX rate variance. We interpret this finding as strong evidence for stability of the power-law behavior over time. These results nicely complement those of Grobys and Kolari (2022) and Grobys (2021), who found power-law exponents to remain virtually the same with time-scale changes from daily to weekly to monthly data.

Moreover, the point estimates for α for realized variances for CAD/USD, CHF/USD, EUR/USD, and JPY/USD conditional on the later subsample fall into the 95% confidence interval for the estimated power-law exponents for the earlier subsample. This is indeed an intriguing result because, as pointed out by Taleb (2010), a long time series is actually able to reveal the properties of power-law processes. Furthermore, in the later subsample, for none of the realized FX rate variances is the GoF able to reject the power-law null hypothesis. Especially, the realized SEK/USD variance exhibits a *p*-value of 0.9990, implying that there is no doubt about the power-law property in the second sample.

However, the empirical fact that realized FX rate variances are governed by power laws in both subsamples does not necessarily imply that they coincide in both samples in the same manner. To explore whether the dynamics of co-fractality are subject to change over time, we test for co-fractality in both subsamples. The results for the first subsample are reported in Table 6, and the results for the second subsample are reported in Table 7. From Panel C of Table 6 we observe that in the first subsample, only the realized variance pair AUS/USD-SEK/USD exhibits statistically significant co-fractality in its strong form. Moreover, from Panel D we see that only 7 out of 36 realized variance pairs do not exhibit co-fractality in its weak form: CAD/USD-CHF/USD, CAD/USD-EUR/USD, CAD/USD-GBP/USD, JPY/USD-CHF, JPY/USD-EUR/USD, JPY/USD-GBP/USD, CHF/USD-NZD/USD.

Notably, from Panels C and D of Table 7 we see that those seven realized FX rate variance pairs do not exhibit co-fractality in the later subsample either. Surprisingly, we also note that only 10 out of 36 realized FX rate variance pairs exhibit co-fractality in its weak form in the later subsample, implying that periods where extreme events are generated have become less connected in the later subsample as opposed to the first subsample. Overall, the results from Tables 6 and 7 suggest that while a lack of co-fractality between realized FX rate variance pairs in earlier subsamples may predict absence of co-fractality also seen in subsequent subsamples, the presence of co-fractality is not as persistent as its absence. Indeed, from 29 realized FX rate variance pairs exhibiting weak co-fractality in the earlier sample, only one-third exhibit co-fractality in its weak form also in the later subsample.

4.2. Co-fractality versus correlation

Table A.1 in the appendix reports the correlation matrices for realized FX rate variances for the overall sample, as well as for both subsamples. We see that the concept of co-fractality is very different from the concept of correlation. While Panel B of Table A.1 shows that the realized CAD/USD variance is statistically significantly positively correlated with realized CHF/USD variance, this pair of realized variances is not exposed to co-fractality. That means that while the thin-tailed part of those realized variance distributions (e.g., x values for which $x < x_{\min}$) co-moves, the tails of those distributions (e.g., x values for which $x \geq x_{\min}$) do not. Comparing Panels B and C, we see that an investor trying to gain portfolio diversification benefits by relying on the concept of correlation may be misled. As an example, the first subsample

Table 3
Co-fractality matrix for the whole sample.

Currency	AUD/USD	CAD/USD	CHF/USD	EUR/USD	GBP/USD	JPY/USD	NOK/USD	NZD/USD	SEK/USD
Panel A. Total observations governed by a power law.									
#	1268	1127	428	412	591	528	253	362	1126
Panel B. Matrix of coinciding observations.									
AUD/USD	1268	777	308	332	410	374	234	340	728
CAD/USD		1127	298	322	404	337	221	308	664
CHF/USD			428	258	198	202	139	167	348
EUR/USD				412	230	189	172	197	388
GBP/USD					591	232	179	230	408
JPY/USD						528	134	204	338
NOK/USD							253	165	232
NZD/USD								362	300
SEK/USD									1126
Panel C. Co-fractality matrix in terms of its strong form.									
AUD/USD	1	0.61***	0.24	0.26	0.32	0.29	0.18	0.27	0.57***
CAD/USD		1	0.26	0.29	0.36	0.30	0.20	0.27	0.52
CHF/USD			1	0.60***	0.34	0.38	0.32	0.39	0.31
EUR/USD				1	0.39	0.36	0.42	0.48	0.34
GBP/USD					1	0.39	0.30	0.39	0.36
JPY/USD						1	0.25	0.39	0.30
NOK/USD							1	0.46	0.21
NZD/USD								1	0.27
SEK/USD									1
Panel D. Co-fractality matrix in terms of its weak form.									
AUS/USD	1	0.69***	0.72***	0.81***	0.69***	0.71***	0.92***	0.94***	0.65***
CAD/USD		1	0.70***	0.78***	0.68***	0.64***	0.87***	0.85***	0.59***
CHF/USD			1	0.63***	0.46	0.47	0.55	0.46	0.81***
EUR/USD				1	0.56**	0.46	0.68***	0.54	0.94***
GBP/USD					1	0.44	0.71***	0.64***	0.69***
JPY/USD						1	0.53	0.56	0.64***
NOK/USD							1	0.65***	0.92***
NZD/USD								1	0.83***
SEK/USD									1

This table reports the co-fractality matrix for realized FX rate variances.
** Statistically significant at a 5% level.
*** Statistically significant at a 1% level.

shows that the realized AUS/USD variance appears to be uncorrelated to any other realized FX rate variance, whereas the second subsample shows that the realized AUS/USD variance is statistically significantly correlated with seven out of eight realized FX rate variances. Moreover, from Panel D of Table 7, we know that the realized AUS/USD variance is statistically significantly co-fractal in its weak form with at least two other realized FX rate variances (e.g., realized CAD/USD and NZD/USD variances). The correlation matrix shows that 16 out of 36 realized FX rate variance pairs do not exhibit statistically significant correlation in the first subsample, whereas 15 of those 16 realized FX rate variance pairs show statistically significantly positive correlation in the later subsample. Hence, relying on the concept of correlation for risk diversification is a poor choice because correlation is sample-specific for processes governed by power laws with tail exponent $\alpha < 3$.

5. Discussion

5.1. Power law or log-normal distribution?

The finding of a power law with a relatively low tail index for some measure of realized volatility is in strong contrast to earlier literature documenting that realized asset variances are typically very close to a log-normal distribution (e.g., Andersen, Bollerslev, Diebold, & Ebens, 2001; Andersen, Bollerslev, Diebold, & Labys, 2001a, 2001b). However, Renò and Rizza (2003), who study the unconditional volatility distribution of the Italian futures market, conclude that the standard assumption of log-normal unconditional volatility has to be rejected and a much better description is provided by a Pareto distribution. Recent studies have confirmed Renò and Rizza (2003) study for realized foreign exchange rate variances, realized variances of commodities (e.g., gold and oil), and the realized variances for the S&P 500 (Grobys, 2021; Grobys & Kolari, 2022). To clarify whether Pareto distributions are more appropriate than the log-normal, we apply Bayes' rule as proposed in Taleb (2020) to examine the plausibility of the lognormal distribution.

As an example, it is worthwhile to consider the event occurring on 2016-06-24 in the realized GBP/USD variance. In Table 8 we report the corresponding maxima in terms of sigma-events (σ-events) for all realized FX variances. From Table 8 we see that this event corresponds to a 47σ-event. How high is the likelihood that the log-normal distribution would generate such an event? From Table 8 we observe that the probability of this event's occurrence, provided the underlying distribution is log-normal (e.g., $P(X > 47\sigma|LGN)$) corresponds to 5.90E-05. How high is the likelihood that a power-law distribution with parameter vector $(\widehat{\alpha}, \widehat{x}_{MIN}) = (2.6085, 0.0113)$ would generate such an event? From Table 8 we observe that the probability of this event's occurrence, provided the underlying distribution is governed by a power law (e.g., $P(X > 47\sigma|PL)$) corresponds 8.77E-04. That means, that the arrival of such an event is 15 times more likely for the power law. According to Benoit Mandelbrot, events of this magnitude never happen.[6] However, this result is in line with Renò and Rizza (2003), who argue that the log-normal distribution is unable to capture the tail behavior of variance distributions. In this regard, Taleb (2020) proposed an application of

[6] See Benoit Mandelbrot's lecture at MIT is available at https://www.youtube.com/watch?v=ock9Gk_aqw4.

Table 4
Estimated power-law functions for the first subsample.

Distribution	$\widehat{\alpha}$	95% CI	$\left\lvert\frac{(3-\widehat{\alpha})}{\widehat{\sigma}}\right\rvert$	$\widehat{x}_{MIN}$	GoF test (p-value)	N	N_{PL}
AUD/USD	2.1262	[2.0361; 2.2163]	19.01	0.0136	0.2270	2007	29.90%
CAD/USD	2.8334	[2.5970; 3.0699]	1.38	0.0206	0.1220	2007	11.51%
CHF/USD	2.7934	[2.6068; 2.9800]	2.17	0.0159	0.7870	2007	17.69%
EUR/USD	2.7359	[2.5463; 2.9255]	2.73	0.0072	0.8610	2007	16.04%
GBP/USD	2.3691	[2.2389; 2.4993]	9.50	0.0102	0.0130	2007	21.18%
JPY/USD	2.5833	[2.4295; 2.7371]	5.31	0.0134	0.6780	2007	20.28%
NOK/USD	3.2075	[2.7949; 3.6201]	0.99	0.0553	0.7870	2007	5.48%
NZD/USD	2.5767	[2.3770; 2.7737]	4.21	0.0353	0.5970	2007	12.26%
SEK/USD	2.2568	[2.1710; 2.3430]	16.90	0.0122	0.0000	2007	40.71%

This table reports the estimates for power-law models $p(x) = (\alpha - 1)x_{MIN}^{\alpha-1}x^{-\alpha}$ using MLE. Tail exponent α is estimated as

$$\widehat{\alpha} = 1 + N\left(\sum_{i=1}^{N} ln\left(\frac{x_i}{x_{MIN}}\right)\right)^{-1},$$

where $\widehat{\alpha}$ denotes the MLE estimator, and N denotes the number of observations, provided $x_i \geq x_{MIN}$. In this model, estimate $\widehat{x}_{MIN}$ is assessed via the KS statistic D, which is the maximum distance between the CDFs of the data and the fitted model:

$$D = MAX_{x \geq x_{MIN}} |S(x) - P(x)|,$$

where $S(x)$ is the CDF of the data for the observation with a value at least x_{MIN}, and $P(x)$ is the CDF for the power-law model that best fits the data in the region $x \geq x_{MIN}$. The estimate of the $\widehat{x}_{MIN}$ is the value of x_{MIN} that minimizes D. Clauset et al.'s (2009) GoF test generates a p-value that quantifies the plausibility of the hypothesis. This test compares D with distance measurements for comparable synthetic data sets drawn from the hypothesized model, and the p-value is defined as the fraction of synthetic distances that are larger than the empirical distance. Given a significance level of 5%, the power-law null hypothesis is not rejected because the difference between the empirical data and the model can be attributed to statistical fluctuations alone. The implementation of this test was detailed by Clauset et al. (2009, p. 675–678). N_{PL} denotes the percentage of sample observations governed by a power-law process, whereas CI denotes the confidence interval for $\widehat{\alpha}$. The sample period used in this table is from 2006 - 05-16 to 2014-03-07.

Table 5
Estimated power-law functions for the second subsample.

Distribution	$\widehat{\alpha}$	95% CI	$\left\lvert\frac{(3-\widehat{\alpha})}{\widehat{\sigma}}\right\rvert$	$\widehat{x}_{MIN}$	GoF test (p-value)	N	N_{PL}
AUD/USD	2.8085	[2.6438; 2.9732]	2.28	0.0095	0.4020	2007	23.07%
CAD/USD	2.9950	[2.7884; 3.2016]	0.05	0.0064	0.2950	2007	17.84%
CHF/USD	2.6711	[2.4498; 2.8924]	2.91	0.0083	0.8110	2007	10.91%
EUR/USD	2.7499	[2.5749; 2.9249]	2.80	0.0031	0.1660	2007	19.13%
GBP/USD	3.1706	[2.8748; 3.4664]	1.13	0.0116	0.5470	2007	10.31%
JPY/USD	2.5946	[2.3958; 2.7934]	4.00	0.0080	0.5640	2007	12.31%
NOK/USD	2.5002	[2.3683; 2.6321]	7.43	0.0121	0.1210	2007	24.76%
NZD/USD	3.0522	[2.9291; 3.2934]	4.21	0.0132	0.1040	2007	13.85%
SEK/USD	2.7354	[2.5731; 2.8977]	3.19	0.0097	0.9990	2007	21.87%

This table reports the estimates for power-law models $p(x) = (\alpha - 1)x_{MIN}^{\alpha-1}x^{-\alpha}$ using MLE. Tail exponent α is estimated as

$$\widehat{\alpha} = 1 + N\left(\sum_{i=1}^{N} ln\left(\frac{x_i}{x_{MIN}}\right)\right)^{-1},$$

where $\widehat{\alpha}$ denotes the MLE estimator, and N denotes the number of observations, provided $x_i \geq x_{MIN}$. In this model, estimate $\widehat{x}_{MIN}$ is assessed via the KS statistic D, which is the maximum distance between the CDFs of the data and the fitted model:

$$D = MAX_{x \geq x_{MIN}} |S(x) - P(x)|,$$

where $S(x)$ is the CDF of the data for the observation with value at least x_{MIN}, and $P(x)$ is the CDF for the power-law model that best fits the data in the region $x \geq x_{MIN}$. The estimate of the $\widehat{x}_{MIN}$ is the value of x_{MIN} that minimizes D. Clauset et al.'s (2009) GoF test generates a p-value that quantifies the plausibility of the hypothesis. This test compares D with distance measurements for comparable synthetic data sets drawn from the hypothesized model, and the p-value is defined as the fraction of synthetic distances that are larger than the empirical distance. Given a significance level of 5%, the power-law null hypothesis is not rejected because the difference between the empirical data and the model can be attributed to statistical fluctuations alone. The implementation of this test was detailed by Clauset et al. (2009, p. 675–678). N_{PL} denotes the percentage of sample observations governed by a power-law process, whereas CI denotes the confidence interval for $\widehat{\alpha}$. The sample period used in this table is from 2014 - 03-10 to 2021-11-19.

Table 6
Co-fractality matrix for the first subsample.

Panel A. Total observations governed by a power law.									
Currency	AUD/USD	CAD/USD	CHF/USD	EUR/USD	GBP/USD	JPY/USD	NOK/USD	NZD/USD	SEK/USD
#	599	230	354	321	422	407	110	246	806
Panel B. Matrix of coinciding observations.									
AUD/USD	599	194	228	240	305	260	107	233	455
CAD/USD		230	116	131	169	137	81	138	197
CHF/USD			354	208	184	172	76	126	292
EUR/USD				321	209	160	94	155	314
GBP/USD					422	204	99	183	357
JPY/USD						407	76	161	281
NOK/USD							110	92	110
NZD/USD								246	218
SEK/USD									806
Panel C. Co-fractality matrix in terms of its strong form.									
AUD/USD	1	0.32	0.38	0.40	0.51	0.43	0.18	0.39	0.56**
CAD/USD		1	0.33	0.41	0.40	0.34	0.35	0.56	0.24
CHF/USD			1	0.59	0.44	0.42	0.21	0.36	0.36
EUR/USD				1	0.50	0.39	0.29	0.48	0.39
GBP/USD					1	0.48	0.23	0.43	0.44
JPY/USD						1	0.19	0.40	0.35
NOK/USD							1	0.37	0.14
NZD/USD								1	0.27
SEK/USD									1
Panel D. Co-fractality matrix in terms of its weak form.									
AUD/USD	1	0.84***	0.64***	0.75***	0.72***	0.64***	0.97***	0.95***	0.76***
CAD/USD		1	0.50	0.57	0.40	0.60**	0.74***	0.60**	0.86***
CHF/USD			1	0.65***	0.73***	0.49	0.69***	0.51	0.82***
EUR/USD				1	0.65***	0.50	0.85***	0.63***	0.98***
GBP/USD					1	0.50	0.90***	0.74***	0.69***
JPY/USD						1	0.69***	0.65***	0.69***
NOK/USD							1	0.84***	1.00***
NZD/USD								1	0.89***
SEK/USD									1

This table reports the co-fractality matrix for realized FX rate variances for the first subsample.
** Statistically significant at a 5% level.
*** Statistically significant at a 1% level.

Bayes' rule to examine to plausibility of some distribution conditional on the occurrence of observed extreme events. Using this application of Bayes' rule, the conditional probability that the underlying distribution is log-normally distributed, given that a 47σ-event occurred on 2016-06-24 in the realized GBP/USD variance, is defined as

$$P(LGN|E) = \frac{P(LGN)P(E|LGN)}{(1-P(LGN))P(E|PL) + P(LGN)P(E|LGN)},$$

where $P(LGN|E)$ is the probability that the distribution is log-normal given that the corresponding event occurred, $P(E|LGN)$ is the probability of the event given that the distribution is log-normal, and $P(E|PL)$ is the probability of the event given that the distribution is governed by a power-law process with $(\hat{\alpha}, \hat{x}_{MIN})$. Assuming various probabilities for $P(LGN)$, Table 9 reports the computed likelihoods $P(LGN|E)$.

From Table 9 we observe that when assuming that the log-normal distribution and the power-law process are equally likely, the likelihood that the distribution is log-normal given a 47σ-event is 6.30%. Even if we assume that the probability that the underlying data-generating process is log-normally distributed is as high as 90%, the likelihood that the distribution is lognormal given a 47σ-event is only 37.71%. We see that even if we make very unreasonable assumptions concerning the likelihood that the underlying distribution is log-normal, the probability that the underlying distribution is log-normal given the arrival of such an extreme event is still less than tossing a fair coin. Thus, according to Taleb (2020) reasoning we can rule out the log-normal distribution as corresponding distribution for the underlying data-generating process governing the realized variance of the GBP/USD FX rate.

In the same manner, we can explore the plausibility of the log-normal distribution versus power laws for the remaining realized FX variances. From Table 9 we observe that when we assume that the log-normal distribution and the power-law process are equally likely, for six out of nine realized FX rate variances, the probability that the underlying distribution is log-normal given the arrival of the corresponding maximums are <25%. Even when we assume that the log-normal distribution is 70% likely, still for six out of nine realized FX rate variances, the probability that the underlying distribution is log-normal given the arrival of the corresponding maximums are <40%. Overall, testing the log-normal distribution against power laws, we see that for the vast majority of realized FX rate variances we can clearly rule out the log-normal distribution as the corresponding underlying data-generating process for realized FX variances.

5.2. Are the chosen values for the cut-offs accurately describing the data-generating power law?

It is well-known that maximum-likelihood estimators for the power law exponent are very sensitive to the chosen value for the cut-off, that is, $x_{\min}$. Selecting a proper cut-off is, of course, a tricky matter. For instance, Lux (2000) attempted to replicate the study of Krämer and Runde (1996), who estimated the tail index (e.g., power law exponent) for the German stock index DAX as well as 26 individual constitutes over the 1960 to 1992 sample period. He argues that the enormous discrepancies between his replication attempt and the results reported in Krämer and Runde (1996) are due to differences in the chosen cut-offs.

Table 7
Co-fractality matrix for the second subsample.

Currency	AUD/USD	CAD/USD	CHF/USD	EUR/USD	GBP/USD	JPY/USD	NOK/USD	NZD/USD	SEK/USD
Panel A. Total observations governed by a power law.									
#	463	355	216	374	206	245	493	277	437
Panel B. Matrix of mutual observations.									
AUD/USD	463	210	128	214	103	125	255	210	215
CAD/USD		355	100	157	90	92	205	157	172
CHF/USD			216	172	53	89	135	87	141
EUR/USD				374	97	129	227	145	235
GBP/USD					206	72	98	78	95
JPY/USD						245	118	101	103
NOK/USD							493	176	272
NZD/USD								277	151
SEK/USD									437
Panel C. Co-fractality matrix in terms of its strong form.									
AUD/USD	1	0.45	0.28	0.46	0.22	0.27	0.52	0.45	0.46
CAD/USD		1	0.28	0.42	0.25	0.26	0.42	0.44	0.39
CHF/USD			1	0.46	0.25	0.36	0.27	0.31	0.39
EUR/USD				1	0.26	0.34	0.46	0.39	0.54
GBP/USD					1	0.29	0.20	0.28	0.22
JPY/USD						1	0.24	0.36	0.24
NOK/USD							1	0.36	0.55
NZD/USD								1	0.35
SEK/USD									1
Panel D. Co-fractality matrix in terms of its weak form.									
AUD/USD	1	0.59**	0.59*	0.57*	0.50	0.51	0.55	0.76***	0.49
CAD/USD		1	0.46	0.42	0.44	0.38	0.58**	0.57	0.39
CHF/USD			1	0.80***	0.26	0.41	0.63***	0.40	0.65***
EUR/USD				1	0.47	0.53	0.61***	0.52	0.63***
GBP/USD					1	0.35	0.48	0.38	0.46
JPY/USD						1	0.48	0.41	0.42
NOK/USD							1	0.64***	0.62***
NZD/USD								1	0.55
SEK/USD									1

This table reports the co-fractality matrix for realized FX rate variances for the second subsample.
* Statistically significant at a 10% level.
** Statistically significant at a 5% level.
*** Statistically significant at a 1% level.

Relatedly, Lux (2001) provides a statistical analysis of high-frequency recordings of the German stock index DAX over the 1988 to 1995 sample period. His study focuses on the limiting behavior characterizing the tail regions of the empirical distribution. In both studies, Lux employs a *data-driven* method to estimate the optimal tail indices by selecting the value of N (see Eq. (8)) by minimizing the mean squared error (MSE) of the following function:

$$N = argmin_N E\left[(\widehat{\gamma}(N) - \gamma)^2\right],$$

where $\widehat{\gamma} = 1/\widehat{\alpha}$. In this regard, Lux (2000, p. 646) highlights: "In view of these problems of implementations, the recent development of methods for *data-driven* selection of the tail sample constitutes an important advance." Following Lux (2000) argument, the current research employs a more recently developed *data-driven* approach, as proposed in Clauset et al. (2009), to select the optimal power law exponents for our realized FX rate variances. As outlined in Section 3.1., this approach is based on the so-called Kolmogorov–Smirnov (KS) distance which is the maximum distance between the cumulative density functions (CDFs) of the data and the fitted power-law model as defined by Eq. (9):

$$D = MAX_{x \geq x_{MIN}} |S(x) - P(x)|,$$

where $S(x)$ is the CDF of the data for the observation with a value of at least x_{MIN}, and $P(x)$ is the CDF for the power-law model that best fits the data in the region $x \geq x_{MIN}$. Using this approach, the optimal cut-off, $\widehat{x}_{MIN}$, is then the value of x_{MIN} that minimizes D. While both approaches are *data-driven*, a refinement of Clauset et al.'s (2009) approach is that based on the optimal D, a GoF test can be derived which generates a p-value that quantifies the plausibility of the power-law null hypothesis. This p-value is obtained from comparing the empirical $\widehat{D}$, estimated via Eq. (9), with simulation-based values for D derived from the theoretical null model.[7] From Table 2 we see that for at least seven

[7] As mentioned earlier, the exact implementation of this test was detailed by Clauset et al. (2009, p. 675–678). Since this test assumes the power law model as the null model, one could refer to this test as a weak test. There is, however, an important issue that needs to be taken into account. Note that under certain circumstances, it might be reasonable to use the hypothesized distribution as null model as opposed to the alternative. As typical example one could, for instance, refer to the time-honored augmented Dickey–Fuller test (ADF) test which assumes the process of higher order of integration under the null hypothesis. This type of test is often used to test a non-stationary process against a stationary process. Specifically, the integrated process is assumed under the null hypothesis, whereas stationarity is presumed under the alternative. The logic is here that when using mistakenly a non-stationary process in an econometric model requiring stationary data, the impact of the error is enormous as statistical inferences derived from such as model are inevitably wrong. The same argument seems to be relevant in the current research context: Using mistakenly a distribution that exhibits strong Pareto-type behavior (e.g., $\alpha < 3$) in an econometric model requiring thin-tailed distributions, the impact of the error is enormous as statistical inferences derived from such as model are inevitably invalid because obtained results will be sample-specific. It is for this reason that even if the GoF test based on the optimal KS distance $\widehat{D}$ is weak in a statistical sense, it seems to be a reasonable choice.

Table 8
Low probability events.

Distribution	Event day	Sigma	$P(E\|LGN)$	$P(E\|PL)$	$\frac{P(E\|PL)}{P(E\|LGN)}$
AUD/USD	2006-12-25	45σ	7.04E-05	2.59E-06	0.0368
CAD/USD	2008-10-29	25σ	6.44E-04	0.00470	7.2981
CHF/USD	2015-01-15	68σ	1.22E-05	1.06E-05	0.8689
EUR/USD	2008-03-17	68σ	1.22E-05	3.15E-09	2.58E-04
GBP/USD	2016-06-24	47σ	5.90E-05	8.77E-04	14.8644
JPY/USD	2008-10-24	26σ	5.61E-04	0.00360	6.4171
NOK/USD	2020-03-20	48σ	5.42E-05	7.73E-04	14.2620
NZD/USD	2012-02-01	63σ	1.71E-05	6.37E-05	3.7251
SEK/USD	2020-03-19	45σ	7.04E-05	0.0020	28.4091

For each foreign exchange market $i = \left\{\frac{AUD}{USD}, \frac{CAD}{USD}, \dots, \frac{SEK}{USD}\right\}$, the time series vectors are standardized and extreme events from Table 1 are collected. This table reports the corresponding deviation in terms of sigma (σ) and the probabilities for its occurrence given that the distribution is log-normal $P(E|LGN)$. We also report the probability for power law distributions $P(E|PL)$ using the estimated optimal parameter vector $(\widehat{\alpha}_i, \widehat{x}_{MIN,i})$ from Table 2.

realized FX rate variances, the GoF test cannot reject the power law null hypothesis. Moreover, from the comparison between power law and log-normal distribution, as elaborated on in Section 4.1., it became evident that power laws are more accurately describing the extreme events for at least six of out nine realized FX rate variance processes. Finally, let us turn our attention to the Hill plots which illustrate the set of possible MLE candidates for the data-generating power laws. Figs. 2–10 show the Hill plots for the estimated realized FX variances. Inspection of Figs. 2–10 shows that for at least five out of nine realized FX rate variances (e.g., AUD/USD, CHF/USD, EUR/USD, NOK/USD, NZD/USD) it is true that

$$\widehat{\alpha} < 3 \forall x_{MIN}.$$

This is strong evidence supporting the results from Table 2 documenting that the power law exponents are below 3. Recall that power law exponents below 3 imply in the current study's research setting that the variances of realized FX rate variances are mathematically undefined. Finally, Clauset et al. (2009) pointed out that using the Hill plot, it is common practice to choose the cut-off $\widehat{x}_{MIN}$ beyond which $\widehat{\alpha}$ is stable. From Fig. 3 in Clauset et al. (2009, p. 670), it is evident that this value corresponds to the saddle point in a $\widehat{\alpha}/\widehat{x}_{MIN}$-graph. Comparing Fig. 2–10 with Fig. 3 in Clauset et al. (2009, p. 670), we see that the optimal cut-offs derived from the optimal KS distances D coincide with the first local maximums in the Hill plots for all realized FX variances, except for CAD/USD and SEK/USD.[8] In this regard, it is noteworthy that the GoF test rejects the power-law null hypothesis only for those two exceptions (e.g., CAD/USD, SEK/USD). Overall, we can interpret these findings as evidence supporting the view that the approach used here to select the most accurate cut-offs provides a proper description of the underlying data-generating power-law processes.

5.3. *What is the difference between co-fractality and the tail dependence coefficient (TDC)?*

This study evaluates the concept of co-fractality in relation to the concept of correlation. The concept of correlation is the most wide-

[8] In Figures 1–9, the optimal cut-offs derived from the optimal KS distances D are highlighted by grey arrows.

Table 9
Bayes' rule.

$P(LGN)$	0.50	0.70	0.90	1
AUD/USD	0.9645	0.9845	0.9959	1
CAD/USD	0.1205	0.2423	0.5522	1
CHF/USD	0.5351	0.7287	0.9120	1
EUR/USD	0.9997	0.9999	1.0000	1
GBP/USD	0.0630	0.1357	0.3771	1
JPY/USD	0.1348	0.2667	0.5838	1
NOK/USD	0.0655	0.1406	0.3869	1
NZD/USD	0.2116	0.3851	0.7073	1
SEK/USD	0.0340	0.0759	0.2406	1

We apply Bayes' rule as outlined in detail in Taleb (2020, p.52) to explore how likely it is that the data generating processes of the realized foreign exchange rate variances are governed by a log-normal process as opposed to a power-law process with $(\widehat{\alpha}_i, \widehat{x}_{MIN,i})$ where $i = \left\{\frac{AUD}{USD}, \frac{CAD}{USD}, \dots, \frac{SEK}{USD}\right\}$, given the low probability events which we observed in Table 8. According to Bayes' rule, the conditional probability that the underlying distribution governing those events follows a log-normal distribution (LGN), conditional on extreme event occurrences, is defined as:

$$P(LGN|E) = \frac{P(LGN)P(E|LGN)}{(1 - P(LGN))P(E|PL) + P(LGN)P(E|LGN)},$$

where $P(LGN|E)$ is the probability that the distribution is log-normal given that the corresponding event occurred, $P(E|LGN)$ is the probability of the event given that the distribution is log-normal, and $P(E|PL)$ is the probability of the event given that the distribution is governed by a power-law process with $(\widehat{\alpha}_i, \widehat{x}_{MIN,i})$. Assuming various probabilities for $P(LGN)$, the table reports the computed likelihoods $P(LGN|E)$.

spread methodological concept used for conducting empirical finance research. It is worth noting that with respect to economic research, the Nobel prize has been awarded at least half a dozen times for research derived from correlation-based methodologies.[9] One could argue that in the corresponding literature, other methodologies have been proposed that share common features with the concept of co-fractality. Perhaps, the methodological concept that is closet to co-fractality is the so-called *tail dependence coefficient* (TDC).[10] Frahm, Junker, and Schmidt (2005) document that three types of TDC estimations have been proposed in the literature, that is, TDC estimations which are based on either (*i*) a specific distribution or a family of distributions, or (*ii*) a specific copula or a family of copulas, or (*iii*) a nonparametric model. Given the research context of the current study, the *upper tail-dependence coefficient* (upper TDC) is perhaps the most relevant counterpart for our realized FX rate variances which is defined as,

[9] Nobel prizes for economic-related research derived from correlation-based methodologies have been awarded to Harry Markowitz (1990), William Sharpe (1990), Robert Merton (1997), Myron Scholes (1997), Robert Engle (2003), Eugene Fama (2013), and Lars Hansen (2013).

[10] Other metrics presenting some alternatives to the concept of correlation are the covariation or the co-difference (see Garel, d'Estampes, & Tjøstheim, 2005; Janicki & Weron, 2021; Samorodnitsky & Taqqu, 2017; Weron, Burnecki, Mercik, & Weron, 2005). Another methodology related to the concept of co-fractality is the concept of exceedance correlation or asymmetric correlation as proposed by Ang and Chen (2002). Exceedance correlation measures correlation asymmetry by evaluating the behavior in the tails of the distribution, and as a consequence, the statistic is not model-specific. There is, however, an important difference between the concept of exceedance correlation or asymmetric correlation, proposed by Ang and Chen (2002), and the concept of co-fractality: exceedance correlation presumes that some correlation exists. On the other hand, the current research shows that correlation is not defined for some financial data governed by power laws with $\alpha < 3$.

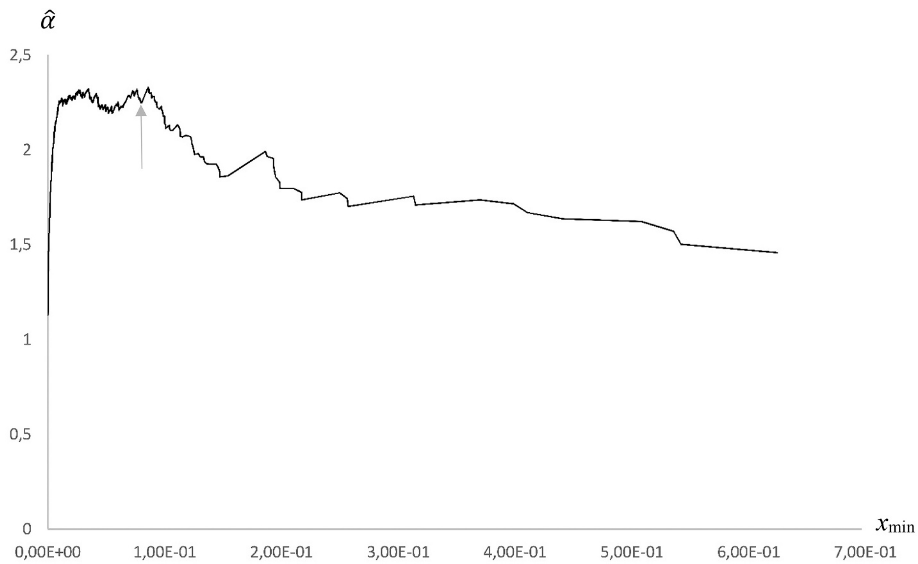

Fig. 2. Hill plot for the AUD/USD realized variance. This figure shows the Hill plot for the AUD/USD realized variance. On the *y*-axis, the graph shows the evolution of $\hat{\alpha}$ depending on the cut-off x_{MIN} which is shown on the *x*-axis. The $\hat{\alpha}$ is obtained from using the MLE estimator of the model $p(x) = (\alpha - 1)x_{MIN}^{\alpha-1}x^{-\alpha}$ where $\hat{\alpha} = 1 + N\left(\sum_{i=1}^{N} ln\left(\frac{x_i}{x_{MIN}}\right)\right)^{-1}$, where N denotes the number of observations, provided $x_i \geq x_{MIN}$. The grey arrow illustrates the location of the MLE with respect to the optimal KS distance D. The sample period used in this table is from 2006 - 05-16 to 2021-11-19.

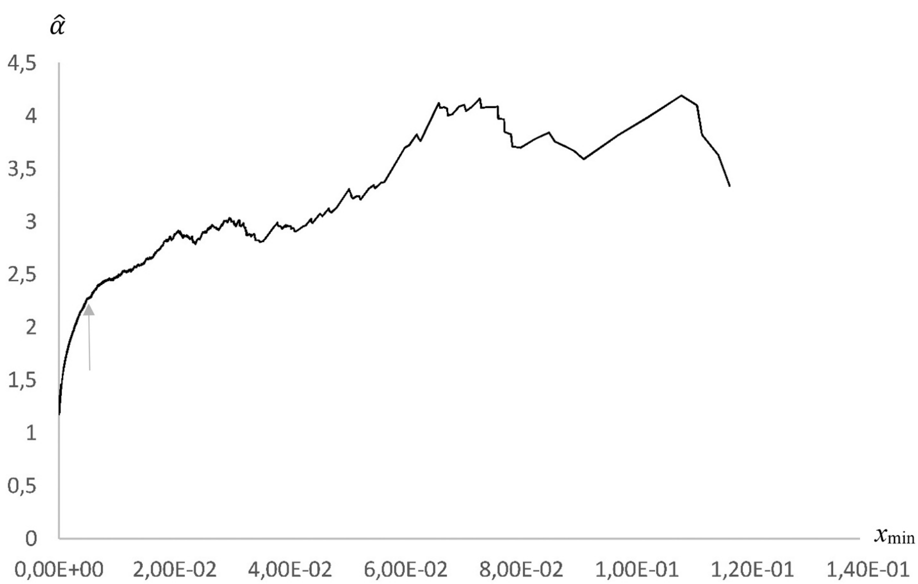

Fig. 3. Hill plot for the CAD/USD realized variance. This figure shows the Hill plot for the CAD/USD realized variance. On the *y*-axis, the graph shows the evolution of $\hat{\alpha}$ depending on the cut-off x_{MIN} which is shown on the *x*-axis. The $\hat{\alpha}$ is obtained from using the MLE estimator of the model $p(x) = (\alpha - 1)x_{MIN}^{\alpha-1}x^{-\alpha}$ where $\hat{\alpha} = 1 + N\left(\sum_{i=1}^{N} ln\left(\frac{x_i}{x_{MIN}}\right)\right)^{-1}$, where N denotes the number of observations, provided $x_i \geq x_{MIN}$. The grey arrow illustrates the location of the MLE with respect to the optimal KS distance D. The sample period used in this table is from 2006-05-16 to 2021-11-19.

$$\lambda_U = \lim_{t \to 1^-} P\{G(X) > t | H(Y) > t\},$$

where (X, Y) is a random pair with joint cumulative distribution function F and marginals G (for X) and H (for Y). Hence, the upper TDC, λ_U, corresponds in essence to the probability that one margin exceeds a high threshold under the condition that the other margin exceeds a high threshold. Note that Cirillo and Taleb (2016, p. 1489) point out that in risk management the top 5% of the observations are typically treated as the extremes of a distribution, which appears to be a common practice which is also in line with earlier relevant literature (Gnedenko, 1943; Gumbel, 1958). To make the results from an application of the upper TDC estimation comparable with our proposed concept of co-fractality, we choose a similar set-up; that is, for each subsample, we use the top 5% of each realized FX rate variance and construct binary vectors that have values of 1 if the realized FX rate variances are in the top 5% of its corresponding distribution and a value of 0 otherwise. As described in Section 3.3, we count the coincides for each pair of binary vectors and divide this figure by the total of extremal values. The relative frequency of coincides is treated as drawings from a Bernoulli distribution, where co-dependency requires relative frequencies to be statistically significantly higher than 0.50.

The results are reported in Tables A.2 and A.3 in the appendix. From Table A.2 we observe that in the first subsample, six realized FX rate variance pairs exhibit statistically significant co-dependencies in their tails. Interestingly, from Table A.3 we observe that only one of those pairs (NZD/USD-AUD/USD) exhibits also statistically significant co-dependency in their tails in the second subsample. Surprisingly, from 30 realized FX rate variance pairs that did not exhibit statistically significant co-dependencies in their tails in the first sample, 29 pairs did not either exhibit statistically significant co-dependencies in their tails in the second sample. This result is similar to what we documented earlier for co-fractality: Whereas lack of co-fractality between realized FX rate variance pairs in the first subsample had some predictive power for expecting absence of co-fractality in the second subsample, the presence of co-fractality has been shown to be not as persistent as its absence.

The question arises what are the benefits from using co-fractality as opposed the TDC? Frahm et al. (2005) point out that it is difficult to conclude whether (X, Y) is tail dependent or not, given one has only

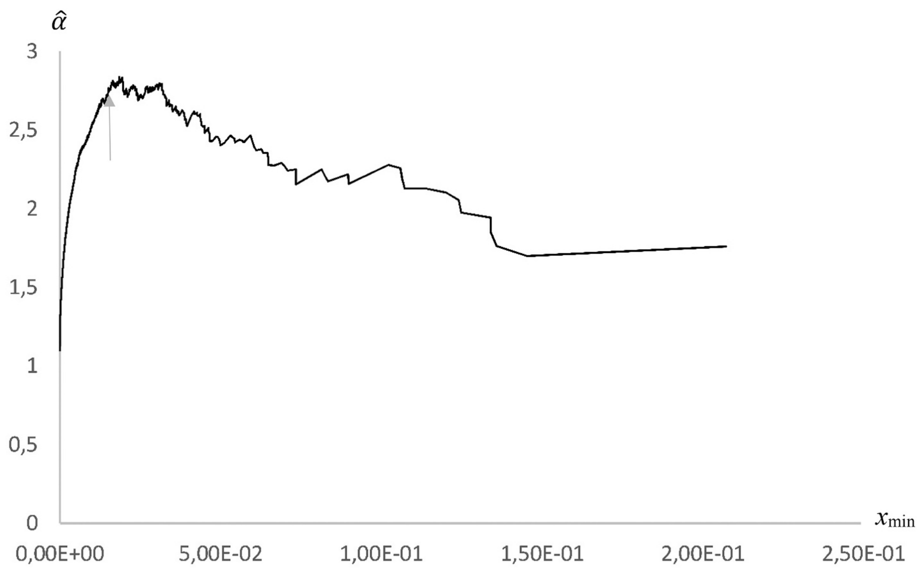

Fig. 4. Hill plot for the CHF/USD realized variance. This figure shows the Hill plot for the CHF/USD realized variance. On the *y*-axis, the graph shows the evolution of $\hat{\alpha}$ depending on the cut-off x_{MIN} which is shown on the *x*-axis. The $\hat{\alpha}$ is obtained from using the MLE estimator of the model $p(x) = (\alpha - 1)x_{MIN}^{\alpha-1}x^{-\alpha}$ where $\hat{\alpha} = 1 + N\left(\sum_{i=1}^{N} ln\left(\frac{x_i}{x_{MIN}}\right)\right)^{-1}$, where N denotes the number of observations, provided $x_i \geq x_{MIN}$. The grey arrow illustrates the location of the MLE with respect to the optimal KS distance D. The sample period used in this table is from 2006-05-16 to 2021-11-19.

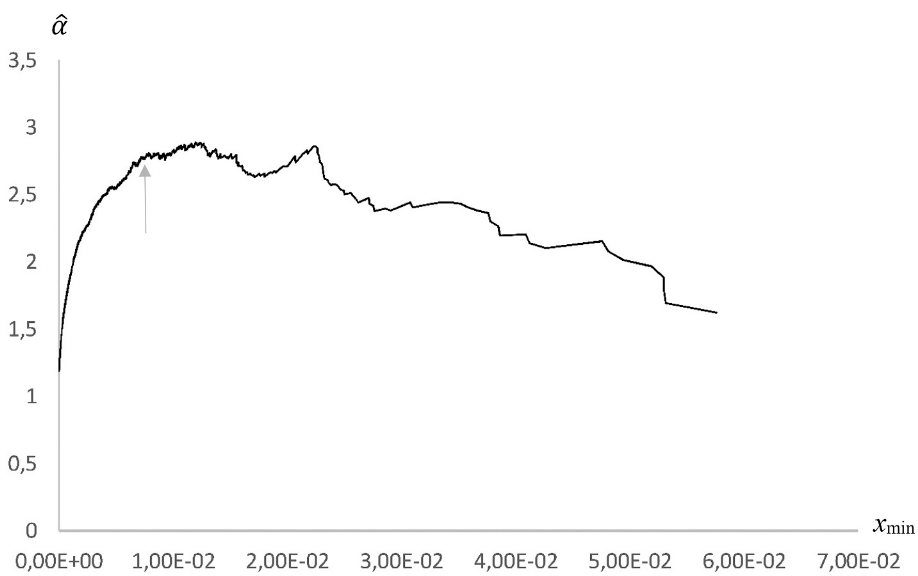

Fig. 5. Hill plot for the EUR/USD realized variance. This figure shows the Hill plot for the EUR/USD realized variance. On the *y*-axis, the graph shows the evolution of $\hat{\alpha}$ depending on the cut-off x_{MIN} which is shown on the *x*-axis. The $\hat{\alpha}$ is obtained from using the MLE estimator of the model $p(x) = (\alpha - 1)x_{MIN}^{\alpha-1}x^{-\alpha}$ where $\hat{\alpha} = 1 + N\left(\sum_{i=1}^{N} ln\left(\frac{x_i}{x_{MIN}}\right)\right)^{-1}$, where N denotes the number of observations, provided $x_i \geq x_{MIN}$. The grey arrow illustrates the location of the MLE with respect to the optimal KS distance D. The sample period used in this table is from 2006-05-16 to 2021-11-19.

finitely many observations (x_1, y_1), …, (x_n, y_n) of (X, Y) available. The authors highlight that as for tail-index estimation, one can always specify thin-tailed distributions which produce sample observations suggesting heavy tails even for large sample sizes.[11] Co-fractality can be considered a data-driven approach, where it is first analyzed whether some specific data exhibit Paretian tails, respectively, are partly governed by fractal processes. Thereby, the power-law exponents are estimated via MLE. In this regard, Clauset et al. (2009, p. 668-669) argue: "…we note that the MLEs are only guaranteed to be unbiased in the asymptotic limit of large sample size, $n \to \infty$. For finite data sets, biases are present but decay as $O(n^{-1})$ for any choice of $x_{\min}$. For very small data sets, such biases can be significant but in most practical situations they can be ignored because they are much smaller than the statistical error of the estimator, which decays as $O(n^{-1/2})$. Our experience suggests that $n \geq 50$ is a reasonable rule of thumb for extracting reliable parameter estimates." Since we have $n \geq 50$ in the tail regions for all MLEs, our estimation results can be deemed reliable.[12] As outlined in Sections 3.1 and 4.2, the power-law models are tested using GoF tests as proposed in Clauset et al. (2009). Hence, the crucial difference between the TDC and co-fractality is that co-fractality assesses in the first step the specific data-generating (power-law) process, and in the second step, potential co-dependencies between those parts of the distributions governed by power-law processes are evaluated.

Whereas the TDC is designed for investigating co-dependencies between tail regions of some distributions, co-fractality is focused on exploring co-dependencies between relatively substantial parts of distributions: For instance, considering the realized variance for AUD/USD, we see from Table 2 that >30% of the variance distribution is governed by a power law. Hence, we cannot argue here with a "tail region" because 30% means a considerable part of the overall distribution. Finally, we argue that extreme events can occur in the part of the

[11] Frahm et al. (2005) provide a detailed simulation study on this issue.

[12] From Table 2 we see that we have at least 4.75% of 4014 observation in the tail region.

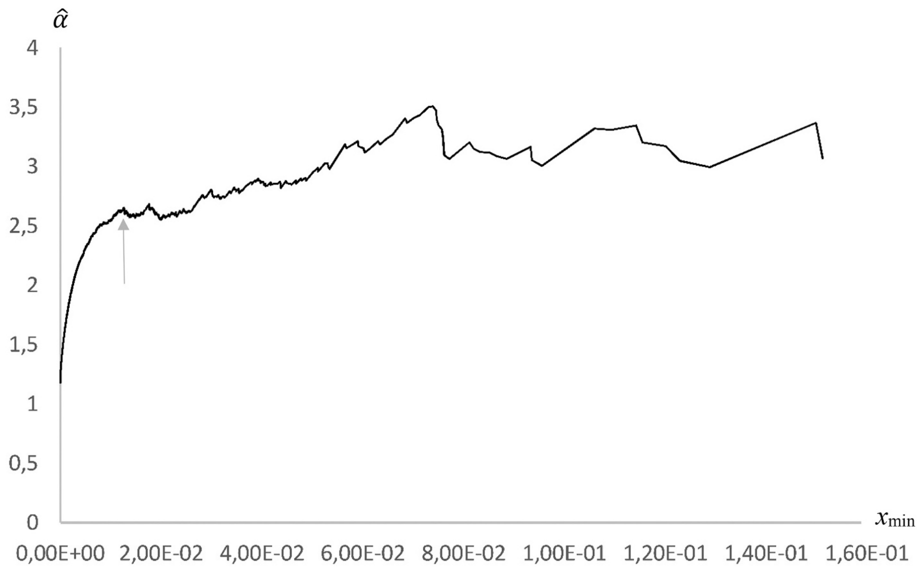

Fig. 6. Hill plot for the GBP/USD realized variance. This figure shows the Hill plot for the GBP/USD realized variance. On the y-axis, the graph shows the evolution of $\hat{\alpha}$ depending on the cut-off x_{MIN} which is shown on the x-axis. The $\hat{\alpha}$ is obtained from using the MLE estimator of the model $p(x) = (\alpha - 1)x_{MIN}^{\alpha-1}x^{-\alpha}$ where $\hat{\alpha} = 1 + N\left(\sum_{i=1}^{N} ln\left(\frac{x_i}{x_{MIN}}\right)\right)^{-1}$, where N denotes the number of observations, provided $x_i \geq x_{MIN}$. The grey arrow illustrates the location of the MLE with respect to the optimal KS distance D. The sample period used in this table is from 2006-05-16 to 2021-11-19.

distribution governed by a power law process–especially power laws with $\alpha < 3$ which allow for "wild behavior"–which appears to be a feature that simple TDC estimation does not account for. Therefore, we would like to argue that co-fractality serves as a more general approach to identify co-dependencies is some specific parts of distributions that may be subject to "wild behavior".

Finally, the concept of co-fractality takes explicitly into account that not all processes that are analyzed would exhibit power-law behaviors to the same degree. For instance, from Table 2 we observed that only 4.75% of the realized variance for the NOK/USD are governed by a power law process, whereas the corresponding figure is 31.61% for the AUD/USD exchange rate. Hence, the maximal fraction that could be subject to potential co-dependency in terms of co-fractality is limited to only 4.75% of the overall sample. TDC estimation – based on the empirical distribution – would in this case treat 0.25% of the distribution as "extremes" even though the corresponding observations are generated from the thin-tailed part of the distribution.

5.4. *Are our results a statistical artefact due to employing the Parkinson estimator to measure FX rate variances?*

The current research follows Grobys' (2021) recent study by using the Parkinson estimator to measure realized FX rate variances. On the one hand, Shu and Zhang (2006), who analyze the relative performance of four range-based volatility estimators including Parkinson, Garman-Klass, Rogers-Satchell, and Yang-Zhang estimators for S&P 500 index data, find that all those price range estimators perform very well. On the other hand, Molnár (2012) argues that when taking into account the noise of range-based volatility estimators, the best estimator appears to be the Garman and Klass (1980) estimator. Hence, to explore whether the assessed power-law property of realized FX rate variances is an artefact manifested in the possibly larger variance produced from the Parkinson estimator, we compare our annualized realized FX rate variances based on daily data with those obtained from using the Garman-Klass estimator. Specifically, in line with Garman and Klass (1980), the realized annualized realized FX rate variances based on daily data are computed as

$$\sigma_{i,t}^2 = T\left(0.5\left[ln\left(\frac{H_{i,t}}{L_{i,t}}\right)\right]^2 - [2ln(2) - 1]\left[ln\left(\frac{C_{i,t}}{O_{i,t}}\right)\right]^2\right),$$

where $H_{i,\ t}$, $L_{i,\ t}$, $C_{i,\ t}$ and $O_{i,\ t}$ denote the highest, lowest, closing and opening price for FX market i on day t, $\sigma_{i,\ t}^2$ denotes FX market i's corresponding realized annualized daily variance, and $T = 250$. The descriptive statistics are reported in Table A.4 in the appendix.[13] Notably, from Table A.4 we observe that for seven out of nine realized FX rate variances, the maximums are even higher for the Garman-Klass estimator as opposed to the Parkinson estimator. Indeed, only the realized FX rate variances for AUD/USD and NOK/USD are lower when computing the realized variances using the Garman-Klass estimator. Next, in Table A.5 the results are reported for estimating the power-law exponents, the corresponding cut-offs assessed via the optimal KS distance D, and GoF tests in line with Clauset et al. (2009).

Strikingly, the we find strong evidence for that our power-law hypothesis is confirmed as the optimal power-law exponents are estimated to be below three for all realized FX rate variances. These results are in line with Table 2 and confirm that the variances of variances are undefined. Furthermore, our findings indicate that for all realized FX rate variances – except for the realized variance for SEK/USD – the estimated power law exponents are more than two standard deviations below $\alpha = 3$. An astonishing result is that the point estimates for six out of nine realized FX variances, the estimated power-law exponents using the Parkinson estimator are within the 95% confidence interval for the power-law exponents estimated using the Garman-Klass estimator. That means, for the vast majority of realized FX rate variances, both estimators are statistically indistinguishable from each other. Overall, we find that the power-law property of realized FX rate variances is not an artefact manifested in the possibly larger variance produced by the Parkinson estimator.

5.5. *Financial implications of co-fractality*

To illustrate some financial implications of our proposed concept of

[13] Note that using the Garman-Klass estimator, we lose five sample observation (e.g., $T = 4009$ as opposed to $T = 4014$ when using the Parkinson estimator). One drawback of the Garman-Klass estimator may be that if $[2ln(2) - 1]\left[ln\left(\frac{C_{i,t}}{O_{i,t}}\right)\right]^2 > 0.5\left[ln\left(\frac{H_{i,t}}{L_{i,t}}\right)\right]^2$, the variance is undefined.

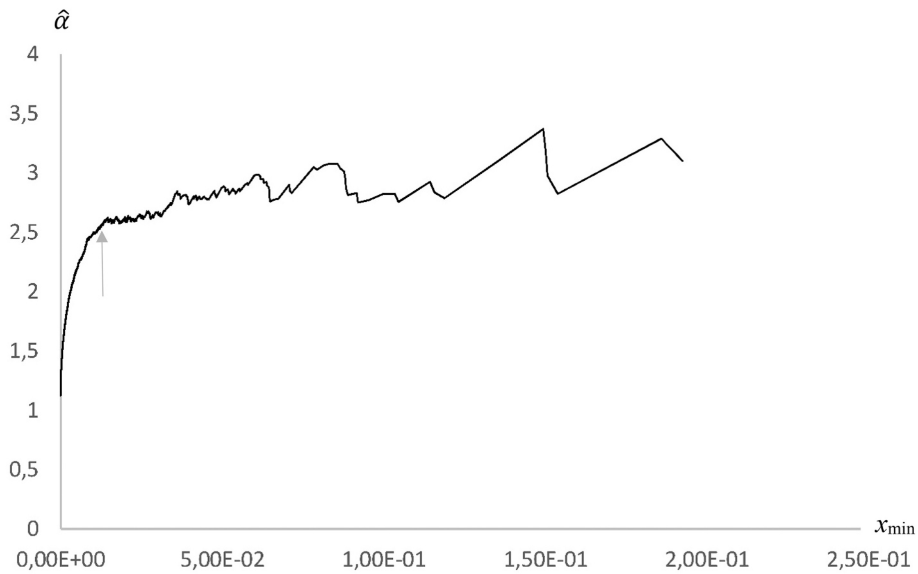

Fig. 7. Hill plot for the JPY/USD realized variance. This figure shows the Hill plot for the JPY/USD realized variance. On the *y*-axis, the graph shows the evolution of $\hat{\alpha}$ depending on the cut-off x_{MIN} which is shown on the *x*-axis. The $\hat{\alpha}$ is obtained from using the MLE estimator of the model $p(x) = (\alpha - 1)x_{MIN}^{\alpha-1}x^{-\alpha}$ where $\hat{\alpha} = 1 + N\left(\sum_{i=1}^{N} ln\left(\frac{x_i}{x_{MIN}}\right)\right)^{-1}$, where N denotes the number of observations, provided $x_i \geq x_{MIN}$. The grey arrow illustrates the location of the MLE with respect to the optimal KS distance D. The sample period used in this table is from 2006-05-16 to 2021-11-19.

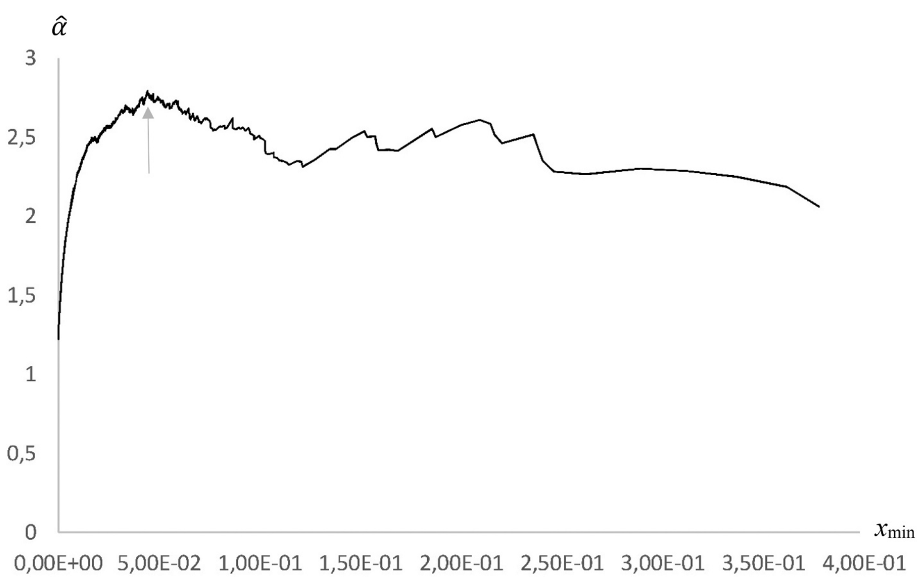

Fig. 8. Hill plot for the NOK/USD realized variance. This figure shows the Hill plot for the NOK/USD realized variance. On the *y*-axis, the graph shows the evolution of $\hat{\alpha}$ depending on the cut-off x_{MIN} which is shown on the *x*-axis. The $\hat{\alpha}$ is obtained from using the MLE estimator of the model $p(x) = (\alpha - 1)x_{MIN}^{\alpha-1}x^{-\alpha}$ where $\hat{\alpha} = 1 + N\left(\sum_{i=1}^{N} ln\left(\frac{x_i}{x_{MIN}}\right)\right)^{-1}$, where N denotes the number of observations, provided $x_i \geq x_{MIN}$. The grey arrow illustrates the location of the MLE with respect to the optimal KS distance D. The sample period used in this table is from 2006-05-16 to 2021-11-19.

co-fractality, let us consider the following two scenarios: First, let us consider N stochastic processes which do not exhibit any co-fractality. This set-up is illustrated in Fig. 11a, where analogously to Figs. 1a – d, the thin-tailed realizations are highlighted in white, whereas the regimes governed by power laws are highlighted in various grey levels. We see from Fig. 11a that power-law regimes are non-overlapping across the N vectors which corresponds to the extreme case where $\lambda^s = \lambda^w = 0$. Recall from Table 8 that the uncertainty in the FX market exhibited numerous extreme events ranging from 25σ events (e.g., CAD/USD realized variance on 2008-10-29) to 68σ events (e.g., EUR/USD realized variance on 2008-03-17). Let us for simplicity assume that all N stochastic processes which are visualized in Fig. 11a are governed by the same power-law process like the CAD/USD realized variance, that is,

$$p(x_1) = p(x_2) = \ldots = p(x_N) = Cx^{-\alpha},$$

with $C = (\alpha - 1)x_{MIN}^{\alpha-1}$, and according to Table 2, $(\alpha, x_{MIN}) = (2.3893, 0.0066)$. If all N stochastic processes are held in a portfolio, the probability of the arrival of extreme events that are equal or larger than 25σ is, according to Table 8, $p = 0.00470$ regardless of time t. Since the thin-tailed part of the distributions do virtually not allow for the occurrences of such extreme events, we can ignore it in the calculation.[14]

Second, let us consider N stochastic processes which exhibit strong co-fractality. This set-up is illustrated in Fig. 11b, where again the thin-tailed realizations are highlighted in white, whereas the regimes governed by power laws are highlighted in various grey levels. We see from Fig. 11b that power-law regimes are overlapping across the N vectors which corresponds to the extreme case where $\lambda^s = 1$. From Fig. 11b we see that the power-law regimes for all N stochastic processes overlap, which occurs between $t = 0$ and $t = 1$. Again, let us for simplicity assume that all N stochastic processes visualized in Fig. 11b are governed by the

[14] For instance, the probability that a $\chi^2(1)$ would generate such an event is $p = 5.73\text{E-}07 \approx 0$.

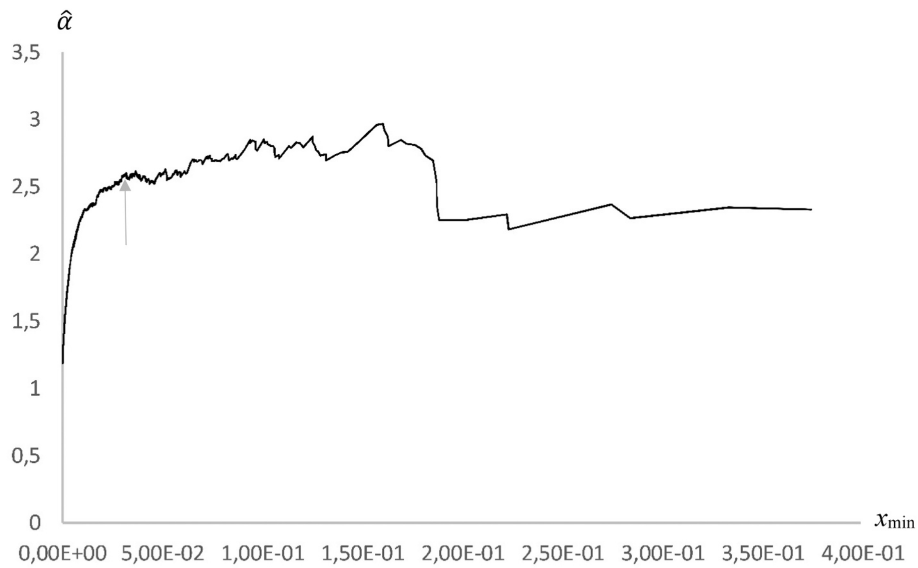

Fig. 9. Hill plot for the NZD/USD realized variance. This figure shows the Hill plot for the NZD/USD realized variance. On the y-axis, the graph shows the evolution of $\hat{\alpha}$ depending on the cut-off x_{MIN} which is shown on the x-axis. The $\hat{\alpha}$ is obtained from using the MLE estimator of the model $p(x) = (\alpha - 1)x_{MIN}^{\alpha-1}x^{-\alpha}$ where $\hat{\alpha} = 1 + N\left(\sum_{i=1}^{N} ln\left(\frac{x_i}{x_{MIN}}\right)\right)^{-1}$, where N denotes the number of observations, provided $x_i \geq x_{MIN}$. The grey arrow illustrates the location of the MLE with respect to the optimal KS distance D. The sample period used in this table is from 2006-05-16 to 2021-11-19.

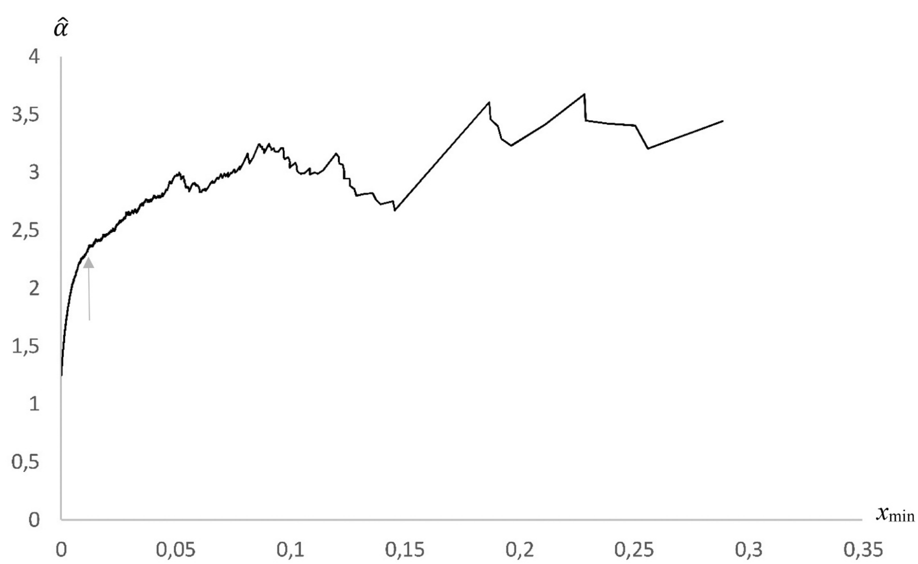

Fig. 10. Hill plot for the SEK/USD realized variance. This figure shows the Hill plot for the SEK/USD realized variance. On the y-axis, the graph shows the evolution of $\hat{\alpha}$ depending on the cut-off x_{MIN} which is shown on the x-axis. The $\hat{\alpha}$ is obtained from using the MLE estimator of the model $p(x) = (\alpha - 1)x_{MIN}^{\alpha-1}x^{-\alpha}$ where $\hat{\alpha} = 1 + N\left(\sum_{i=1}^{N} ln\left(\frac{x_i}{x_{MIN}}\right)\right)^{-1}$, where N denotes the number of observations, provided $x_i \geq x_{MIN}$. The grey arrow illustrates the location of the MLE with respect to the optimal KS distance D. The sample period used in this table is from 2006-05-16 to 2021-11-19.

same power-law process like in the previous example. The question arises, what is the probability of the arrival of extreme events that are equal or larger than 25σ? Obviously, from $t = 1$ to T, the probability of the occurrences of such extreme events is negligible because all distributions are contemporaneously governed by some thin-tailed distribution. However, between $t = 0$ and $t = 1$, the risk is accelerated by a substantial margin because all distributions are contemporaneously in a power-law regime allowing extreme events to occur. Using the same assumptions like in the previous example, the probability of the occurrences of extreme events that are equal or larger than 25σ is then 0.00470 N. More precisely, holding a portfolio of $N = 100$ stochastic processes, the probability of the occurrences of extreme events that are equal or larger than 25σ is, hence, 47%. It is perhaps for this reason that Mandelbrot (2008, p. 266) critically argues: "When you pick a stock by the conventional [correlation-based] method, you may actually be adding risk rather than reducing it."

6. Conclusion

Periods of high variance are manifestations of extreme events. Traditional finance research has typically relied on the concept of correlation for risk diversification. For instance, Markowitz (1952) modern portfolio theory, Sharpe (1964) capital asset pricing model, and factor models such as the Fama and French (1993, 2015, 2018) models, are built on the fundamental concept of correlation. This study argues that correlation is not defined for processes governed by power laws for

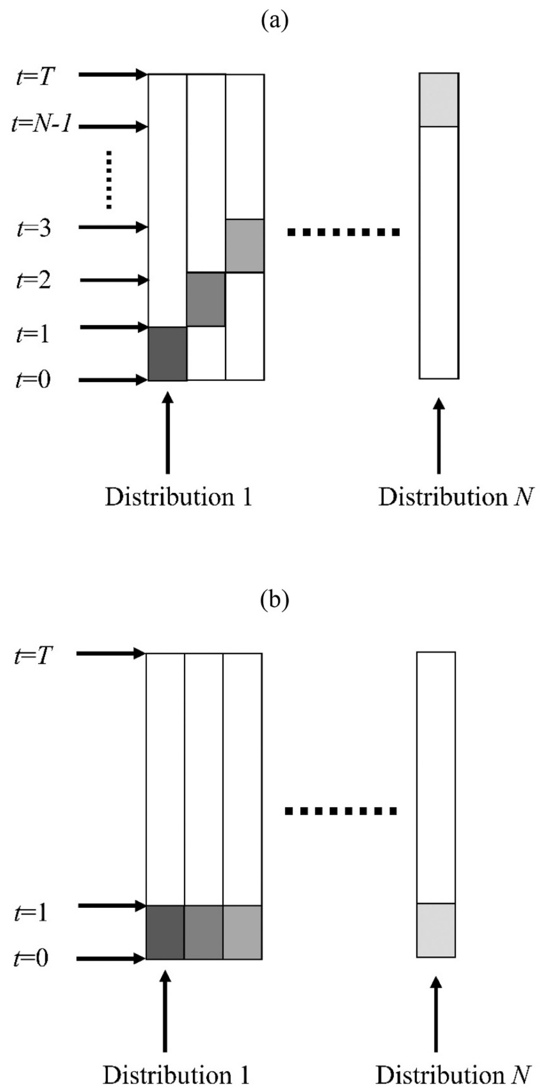

Fig. 11. Diversification of extreme risks.
Fig. 11a illustrates N stochastic processes which do not exhibit any co-fractality. Analogously to Figs. 1a – d, the thin-tailed realizations are highlighted in white, whereas the regimes governed by power laws are highlighted in various grey levels. Power-law regimes are non-overlapping across the N vectors which corresponds to the extreme case where $\lambda^s = \lambda^w = 0$. Fig. 11b illustrates N stochastic processes which exhibit strong co-fractality. Power-law regimes are overlapping between $t = 0$ and $t = 1$ across the N vectors which corresponds to the extreme case where $\lambda^s = 1$. Risk is accelerated by a substantial margin because all distributions are contemporaneously in a power-law regime allowing extreme events to occur.

which $\alpha < 3$. As a remedy, we define the concept of co-fractality as a metric measuring the co-dependence of power law regimes across distributions. Specifically, we distinguish between co-fractality in its strong and weak form because the number of realizations governed by power laws differs between distributions; in our sample the percentage of the part of the distributions governed by power-law processes varies between 4.75% and 31.61%.

Confirming the results of earlier studies, the findings of this study indicate that realized variance processes for G10 currencies exhibit Paretian tails. Because our findings indicate that the power laws for all currencies exhibit exponents that are statistically below 3 ($\alpha < 3$), we infer that theoretical second moments for the variances—or the variances of variances—do not exist. This implies, in turn, that correlations between variances are not defined either. Extending earlier studies, we show that the tail exponents are stable in the sense that the null hypothesis that the power-law exponent is $\alpha < 3$ cannot be rejected across independent samples.

We further explore whether the co-dependencies of power-law regimes between realized FX rate variances change over time. While the evidence suggests that the presence of co-fractality does not appear to be persistent across samples, the absence of co-fractality seems to be more persistent than its presence. Indeed, all seven realized FX rate variance pairs not exhibiting weak co-fractality in the first subsample do not show evidence of any co-fractality in the later subsample either. This finding suggests that tail risks can be diversified. Power laws with tail exponent $\alpha < 3$ generate unforeseen extreme events; that is, past data do not capture the potential impact of extreme events in the future. Hence, risk diversification should focus on the tails as opposed to the thin-tailed part of the data (e.g., $x < x_{\min}$). Since the concept of correlation is not designed as a tool that could be used for portfolio or risk management in the presence of power-law processes with tail exponent $\alpha < 3$, the concept of co-fractality appears to be a remedy to solve such concrete financial problems that traditional approaches cannot solve.

The current study opens various avenues for future research. For instance, referring to the usage of power laws to model financial market data, Lux and Alfarano (2016) commented, "The power laws in returns and in volatility seem to be intimately related: none of them was ever observed without the other and it, therefore, seems warranted to interpret them as the joint essential characteristics of financial data" (p. 4). While the current research explicitly explores power-law interdependencies in the second moments of some financial assets (i.e., G10 currencies), future research is encouraged to investigate this issue for other moments. Given that power laws are a stylized fact for financial market data, future research could also explore power-law interdependencies for the first or second moment for equities or for cryptocurrencies. Moreover, future research is needed to validate the result of this study for expanded samples. Finally, the simple and straightforward methodological approach proposed in this study, where co-dependencies are derived from transformations of random processes into some Bernoulli-distributed auxiliary variables, can be arbitrarily extended to modeling co-dependencies of many other fat-tailed distributions – not necessarily ones that exactly follow power-law decays. These issues are, however, beyond the scope of the current research and are therefore left for future research.

Data availability

Foreign exchange rate data used in this study is publicly available.

Appendix A. Appendix

Table A.1
Correlation matrices.

Currency	AUD/USD	CAD/USD	CHF/USD	EUR/USD	GBP/USD	JPY/USD	NOK/USD	NZD/USD	SEK/USD
Panel A. Correlation matrix for the overall sample from 2006-05-16 to 2021-11-19.									
AUD/USD	1	0.00	0.00	0.00	0.00	−0.01	0.00	0.00	0.00
CAD/USD		1	0.06***	0.02	0.40***	0.43***	0.30***	0.07***	0.44***
CHF/USD			1	0.00	0.02	0.04**	0.04**	0.01	0.05***
EUR/USD				1	0.01	0.04**	0.01	0.01	0.00
GBP/USD					1	0.52***	0.34***	0.05***	0.46***
JPY/USD						1	0.22***	0.06***	0.36***
NOK/USD							1	0.05***	0.69***
NZD/USD								1	0.08***
SEK/USD									1
Panel B. Correlation matrix for the sample from 2006-05-16 to 2014-03-07.									
AUD/USD	1	−0.01	−0.01	0.00	0.00	−0.01	−0.01	0.00	−0.01
CAD/USD		1	0.23***	0.01	0.57***	0.42***	0.60***	0.06***	0.57***
CHF/USD			1	0.03	0.24***	0.15***	0.28**	0.02	0.28***
EUR/USD				1	0.01	0.04*	0.02	0.00	0.00
GBP/USD					1	0.58***	0.67***	0.06***	0.63***
JPY/USD						1	0.48***	0.06**	0.45***
NOK/USD							1	0.07***	0.79***
NZD/USD								1	0.07***
SEK/USD									1
Panel C. Correlation matrix for the sample from 2014 to 2014-03-10 to 2021-11-19.									
AUD/USD	1	0.54***	0.04*	0.37***	0.26***	0.32***	0.63***	0.90***	0.82***
CAD/USD		1	0.09***	0.37***	0.27***	0.31***	0.42***	0.46***	0.39***
CHF/USD			1	0.09***	0.00	0.04*	0.03	0.06***	0.03
EUR/USD				1	0.46***	0.57***	0.30***	0.36***	0.41***
GBP/USD					1	0.59***	0.28***	0.26***	0.36***
JPY/USD						1	0.20***	0.35***	0.28***
NOK/USD							1	0.60***	0.72***
NZD/USD								1	0.78***
SEK/USD									1

This table reports the correlation matrix for realized FX rate variances for the overall sample and subsamples.
* Statistically significant at a 10% level.
** Statistically significant at a 5% level.
*** Statistically significant at a 1% level.

Table A.2
Tail dependence coefficient for the first subsample.

	AUD/USD	CAD/USD	CHF/USD	EUR/USD	GBP/USD	JPY/USD	NOK/USD	NZD/USD	SEK/USD
AUD/USD	1	0.5960***	0.3535	0.4848	0.5556**	0.4242	0.5152	0.6768***	0.4848
CAD/USD		1	0.3232	0.4343	0.5354	0.3737	0.5253	0.5253	0.4545
CHF/USD			1	0.5152	0.3939	0.2929	0.3737	0.2929	0.3939
EUR/USD				1	0.5253	0.3333	0.5859***	0.4040	0.5152
GBP/USD					1	0.3939	0.5556**	0.4747	0.5253
JPY/USD						1	0.3737	0.4141	0.3535
NOK/USD							1	0.5253	0.6566***
NZD/USD								1	0.4747
SEK/USD									1

Using the empirical distributions, we define the empirical tail dependence coefficient as the relative fraction of coincides in the 5% upper tail of the underlying variance.
** Statistically significant at a 5% level.
*** Statistically significant at a 1% level.

Table A.3
Tail dependence coefficient for the second subsample.

	AUD/USD	CAD/USD	CHF/USD	EUR/USD	GBP/USD	JPY/USD	NOK/USD	NZD/USD	SEK/USD
AUD/USD	1	0.4040	0.2929	0.3131	0.2424	0.3131	0.4646	0.5556**	0.4343
CAD/USD		1	0.2828	0.2626	0.2020	0.2828	0.3333	0.4444	0.3131
CHF/USD			1	0.6768***	0.1616	0.2828	0.3434	0.3030	0.3535
EUR/USD				1	0.1919	0.2727	0.3636	0.3232	0.4545
GBP/USD					1	0.2222	0.2626	0.2727	0.2626
JPY/USD						1	0.2323	0.3535	0.2323
NOK/USD							1	0.3838	0.4949

(continued on next page)

Table A.3 (*continued*)

	AUD/USD	CAD/USD	CHF/USD	EUR/USD	GBP/USD	JPY/USD	NOK/USD	NZD/USD	SEK/USD
NZD/USD								1	0.3737
SEK/USD									1

Using the empirical distributions, we define the empirical tail dependence coefficient as the relative fraction of coincides in the 5% upper tail of the underlying variance.

** Statistically significant at a 5% level.

*** Statistically significant at a 1% level.

Table A.4

Descriptive statistics.

Exchange rate	AUD/USD	CAD/USD	CHF/USD	EUR/USD	GBP/USD	JPY/USD	NOK/USD	NZD/USD	SEK/USD
Mean	0.1368	0.0088	0.0130	0.3136	0.0344	0.0091	0.0185	0.0206	0.0158
Median	0.0079	0.0048	0.0052	0.0046	0.0050	0.0046	0.0097	0.0085	0.0084
Maximum	241.2387	0.3642	13.8377	1187.7764	1.2456	0.5358	2.2629	19.8368	1.6670
Date of maximum	2006-12-25	2008-10-29	2015-01-15	2008-03-17	2016-06-24	2008-10-24	2020-03-20	2012-02-01	2020-03-19
Minimum	3.3621E-06	1.7479E-05	2.0055E-07	1.3363E-05	8.5145E-06	3.6122E-07	1.1825E-04	2.1406E-05	1.4454E-04
Std.Dev.	5.3699	0.0151	0.2201	18.7684	1.6106	0.0201	0.0557	0.3143	0.0363
Skewness	44.7512	8.9298	61.9117	63.2237	63.2956	12.6614	26.3595	62.5630	26.6588
Kurtosis	2001.8233	145.0671	3886.4707	4000.9855	4007.2082	236.3705	900.2472	3944.8998	1108.4176
Start of the sample	2006-05-16	2006-05-16	2006-05-16	2006-05-16	2006-05-16	2006-05-16	2006-05-16	2006-05-16	2006-05-16
End of the sample	2021-11-19	2021-11-19	2021-11-19	2021-11-19	2021-11-19	2021-11-19	2021-11-19	2021-11-19	2021-11-19
Total observations	4009	4009	4009	4009	4009	4009	4009	4009	4009

This table reports the descriptive statistics for the annualized daily realized variance for AUD/USD, CAD/USD, CHF/USD, EUR/USD, GBP/USD, JPY/USD, NOK/USD, NZD/USD, and SEK/USD. The annualized daily realized variances for each FX market i are in line with those of Garman and Klass (1980), computed as

$$\sigma^2_{i,t} = T\left(0.5\left[ln\left(\frac{H_{i,t}}{L_{i,t}}\right)\right]^2 - [2ln(2)-1]\left[ln\left(\frac{C_{i,t}}{O_{i,t}}\right)\right]^2\right),$$

where $H_{i,\,t}$, $L_{i,\,t}$, $C_{i,\,t}$ and $O_{i,\,t}$ denote the highest, lowest, closing and opening price for FX market i on day t, $\sigma^2_{i,\,t}$ denotes FX market i's corresponding realized annualized variance, and $T = 250$. Publicly available daily data on FX rates against the USD were retrieved from finance.yahoo.com.

Table A.5

Estimated power-law functions for variances using the Garman and Klass (1980) estimator.

Distribution	$\widehat{\alpha}$	95% CI	$\left\lvert\frac{(3-\widehat{\alpha})}{\widehat{\sigma}}\right\rvert$	$\widehat{x}_{MIN}$	GoF test (p-value)	N	N_{PL}
AUD/USD	2.3063	[2.2334; 2.3797]	18.66	0.0133	0.1410	4009	30.81%
CAD/USD	2.5329	[2.4351; 2.6307]	9.36	0.0103	0.0000	4009	23.55%
CHF/USD	2.7817	[2.6246; 2.9388]	2.72	0.0172	0.6880	4009	12.32%
EUR/USD	2.8198	[2.6466; 2.9930]	2.04	0.0170	0.4280	4009	10.58%
GBP/USD	2.7325	[2.5799; 2.8851]	3.44	0.0159	0.4360	4009	12.35%
JPY/USD	2.6781	[2.5199; 2.8363]	3.99	0.0176	0.4780	4009	10.78%
NOK/USD	2.6818	[2.5390; 2.8246]	4.37	0.0301	0.0300	4009	13.33%
NZD/USD	2.4579	[2.3771; 2.5387]	13.16	0.0140	0.0000	4009	31.23%
SEK/USD	2.9792	[2.7617; 3.1967]	0.19	0.0395	0.2460	4009	7.93%

This table reports the estimates for power-law models $p(x) = (\alpha - 1)x_{MIN}^{\alpha-1}x^{-\alpha}$ using MLE. Tail exponent α is estimated as

$$\widehat{\alpha} = 1 + N\left(\sum_{i=1}^{N} ln\left(\frac{x_i}{x_{MIN}}\right)\right)^{-1},$$

where $\widehat{\alpha}$ denotes the MLE estimator, and N denotes the number of observations, provided $x_i \geq x_{MIN}$. In this model, estimate $\widehat{x}_{MIN}$ is assessed via the KS statistic D, which is the maximum distance between the CDFs of the data and the fitted model:

$$D = MAX_{x \geq x_{MIN}} \lvert S(x) - P(x) \rvert,$$

where $S(x)$ is the CDF of the data for the observation with a value at least x_{MIN}, and $P(x)$ is the CDF for the power-law model that best fits the data in the region $x \geq x_{MIN}$. The estimate of the $\widehat{x}_{MIN}$ is the value of x_{MIN} that minimizes D. Clauset et al. (2009) GoF test generates a p-value that quantifies the plausibility of the hypothesis. This test compares D with distance measurements for comparable synthetic data sets drawn from the hypothesized model, and the p-value is defined as the fraction of synthetic distances that are larger than the empirical distance. Given a significance level of 5%, the power-law null hypothesis is not rejected because the difference between the empirical data and the model can be attributed to statistical fluctuations alone. The implementation of this test was detailed by Clauset et al. (2009, p. 675–678). N_{PL} denotes the percentage of sample observations governed by a power-law process, whereas CI denotes the confidence interval for $\widehat{\alpha}$. The sample period used in this table is from 2006-05-16 to 2021-11-19.

References

Alexander, C. (2009). *Market Risk Analysis, Value at Risk Models*. Hoboken, USA: John Wiley and Sons.

Andersen, T. G., Bollerslev, T., Diebold, F. X., & Ebens, H. (2001). The distribution of realized stock return volatility. *Journal of Financial Economics, 61*, 43–76.

Andersen, T. G., Bollerslev, T., Diebold, F. X., & Labys, P. (2001a). Modeling and forecasting realized volatility. *Econometrica, 71*, 579–625.

Andersen, T. G., Bollerslev, T., Diebold, F. X., & Labys, P. (2001b). The distribution of realized exchange rate volatility. *Journal of the American Statistical Association, 96*, 42–55.

Ang, A., & Chen, J. (2002). Asymmetric correlations of equity portfolios. *Journal of Financial Economics, 63*(3), 443–494.

Aven, T. (2016). Risk assessment and risk management: Review of recent advances on their foundation. *European Journal of Operational Research, 253*, 1–13.

Boldanov, R., Degiannakis, S., & Filis, G. (2016). Time-varying correlation between oil and stock market volatilities: Evidence from oil-importing and oil-exporting countries. *International Review of Financial Analysis, 48*, 209–220.

Cirillo, P., & Taleb, N. N. (2016). Expected shortfall estimation for apparently infinite-mean models of operational risk. *Quantitative Finance, 16*, 1485–1494.

Clauset, A., Shalizi, C. R., & Newman, M. E. J. (2009). Power law distributions in empirical data. *SIAM Review, 51*, 661–703.

Cont, R., Potters, M., & Bouchaud, J. (1997). Scaling in stock market data: Stable laws and beyond. In B. Dubrulle, F. Graner, & D. Sornette (Eds.)*, 7. Scale invariance and beyond. Centre de physique des Houches* (pp. 75–85). Springer. ISBN 978-3-540-640 0 0-4.

Da Fonseca, J., & Zhang, W. (2019). Volatility of volatility is (also) rough. *Journal of Futures Markets, 39*(5), 600–611.

Dalio, R. (2017). *Principles*. New York, USA: Simon & Schuster.

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

Fama, E., & French, K. R. (2015). A five-factor asset pricing model. *Journal of Financial Economics, 116*, 1–22.

Fama, E., & French, K. R. (2018). Choosing factors. *Journal of Financial Economics, 128*, 234–252.

Fama, E. F. (1963). Mandelbrot and the stable Paretian hypothesis. *Journal of Business, 36*, 420–429.

Frahm, G., Junker, N., & Schmidt, R. (2005). Estimating the tail-dependence coefficient: Properties and pitfalls. *Insurance: Mathematics & Economics, 37*, 80–100.

Garel, B., d'Estampes, L., & Tjøstheim, D. (2005). Revealing some unexpected dependence properties of linear combinations of stable random variables using symmetric covariation. *Communications in Statistics-Theory and Methods, 33*(4), 769–786.

Garman, M., & Klass, M. (1980). On the estimation of security price volatilities from historical data. *Journal of Business, 53*, 67–78.

Gnedenko, D. V. (1943). Sur la distribution limité du terme d'une série aléatoire. *Annals of Mathematics, 44*, 423–453.

Gopikrishnan, P., Meyer, M., Amaral, L., & Stanley, H. (1998). Inverse cubic law for the distribution of stock price variations. *The European Physical Journal B - Condensed Matter and Complex Systems, 3*, 139–140.

Grobys, K. (2021). What do we know about the second moment of financial markets? *International Review of Financial Analysis, 78*, Article 101891.

Grobys, K., Junttila, J., Kolari, J. W., & Sapkota, N. (2021). On the stability of stablecoins. *Journal of Empirical Finance, 64*, 207–223.

Grobys, K., & Kolari, J. W. (2022). Why Benoit Mandelbrot was right: Fractality and realized foreign exchange rate variances, working paper. available at: https://www.researchgate.net/publication/358901267_Why_Benoit_Mandelbrot_was_right_Fractality_and_realized_foreign_exchange_rate_variances.

Gumbel, E. J. (1958). *Statistics of extremes*. Cambridge: Cambridge University Press.

Hall, J., Brorsen, B., & Irwin, S. (1989). The distribution of futures prices: A test of the stable paretian and mixture of normals hypotheses. *The Journal of Financial and Quantitative Analysis, 24*, 105–116.

Harvey, C. R., Liu, Y., & Zhu, H. (2016). And the cross-section of expected returns. *Review of Financial Studies, 29*, 5–68.

Hill, B. (1975). A simple general approach to inference about the tail of a distribution. *Annals of Statistics, 3*, 1163–1174.

Hou, K., Xue, C., & Zhang, L. (2020). Replicating anomalies. *Review of Financial Studies, 33*, 2019–2133.

Janicki, A., & Weron, A. (2021). *Simulation and chaotic behavior of alpha-stable stochastic processes*. CRC Press.

Jansen, D. W., & de Vries, C. G. (1991). On the frequency of large stock returns: Putting booms and busts into perspective. *The Review of Economics and Statistics, 73*, 18–24.

Kaplanski, G., & Levy, H. (2015). Value-at-risk capital requirement regulation, risk taking and asset allocation: a mean–variance analysis. *European Journal of Finance, 21*(3), 215–241.

Krämer, W., & Runde, R. (1996). Stochastic properties of German stock returns. *Empirical Economics, 21*, 281–306.

Lau, A.-L., Lau, H.-S., & Wingender, J. (1990). The distribution of stock returns: New evidence against the stable model. *Journal of Business and Economic Statistics, 8*, 217–223.

Lux, T. (1996). The stable paretian hypothesis and the frequency of large returns: An examination of major German stocks. *Applied Financial Economics, 6*, 463–475.

Lux, T. (2000). On moment condition failure in German stock returns: An application of recent advances in extreme value statistics. *Empirical Economics, 25*, 641–652.

Lux, T. (2001). The limiting extremal behavior of speculative returns: An analysis of intra-daily data from the Frankfurt stock exchange. *Applied Financial Economics, 11*, 299–315.

Lux, T., & Alfarano, S. (2016). Financial power laws: Empirical evidence, models, and mechanisms. *Chaos, Solitons and Fractals, 88*, 3–18.

Mandelbrot, B. (1963). The variation of certain speculative prices. *Journal of Business, 36*, 394–419.

Mandelbrot, B. (2008). *The (Mis)Behaviour of Markets*. London, UK: Profile Books.

Markowitz, H. M. (1952). Portfolio selection. *Journal of Finance, 7*, 77–91.

Molnár, P. (2012). Properties of range-based volatility estimators. *International Review of Financial Analysis, 23*, 20–29.

Palomba, G., & Riccetti, L. (2012). Portfolio frontiers with restrictions to tracking error volatility and value at risk. *Journal of Banking & Finance, 36*(9), 2604–2615.

Parkinson, M. (1980). The extreme value method for estimating the variance of the rate of return. *Journal of Business, 53*(1), 61–65.

Renò, R., & Rizza, R. (2003). Is volatility lognormal? Evidence from Italian futures. *Physica A: Statistical Mechanics and its Applications, 322*, 620–628.

Samorodnitsky, G., & Taqqu, M. S. (2017). *Stable non-Gaussian random processes: Stochastic models with infinite variance: Stochastic modeling.* Routledge.

Serra-Garcia, M., & Gneezy, U. (2021). Nonreplicable publications are cited more than replicable ones. *Science Advances, 7*(21), 1705.

Sharpe, W. F. (1964). Capital asset prices a theory of market equilibrium under conditions of risk. *Journal of Finance, 19*, 425–442.

Shu, J. H., & Zhang, J. E. (2006). Resting range estimators of historical volatility. *Journal of Futures Markets, 26*, 297–313.

Taleb, N. N. (2010). *The black swan*. New York, NY: Random House.

Taleb, N. N. (2020). *Statistical consequences of fat tails: Real world Preasymptotics, epistemology, and applications, papers and commentary*. STEM Academic Press.

Werner, T., & Upper, C. (2002). *Time variation in the tail behavior of bund futures returns. Discussion paper 25/02.* Deutsche Bundesbank 2002.

Weron, A., Burnecki, K., Mercik, S., & Weron, K. (2005). Complete description of all self-similar models driven by Lévy stable noise. *Physical Review E, 71*(1), Article 016113.

White, E., Enquist, B., & Green, J. L. (2008). On estimating the exponent of power-law frequency distributions. *Ecology, 89*, 905–912.