

Equity's Randomness to Expect

Working Paper, this version: January 14, 2022

Wolfgang Schadner

University of St. Gallen, Switzerland

ARTICLE INFO

Keywords:

complex behavior
investor expectations
implied fractals
rough volatility
volatility term-structure
risk-neutral moments
non-linear phenomena

ABSTRACT

It is fact that investor behavior is complex, translating into non-linear asset price dynamics. This fact implies that auto-correlation in returns varies over time and occasionally differs from zero. Being aware of that, the questions about measurement of expected return randomness, its predictability and the economic implications arise. While much attention has been paid to option-implied densities, less on their scaling behavior. The econometric explanation for implied variance scaling non-linear in horizon is via the underlying's correlation structure, which is well explained by fractal geometry. Persistence, as such, is the crucial determinant herein. This work introduces the well established fractal toolbox to financial options, allowing to estimate expected return persistence in a model-free manner. The derived measure is an ex-ante gauge of randomness and market efficiency, with direct effects on financial stability, asset prices and risk. While persistence describes the randomness of returns, the multifractal strength explains how much it can vary over time. Market-crashes, for example, are a structural break in persistence. The concept of implied fractals is empirically applied upon daily data of the S&P 500 index. A surprisingly strong predictability from ex-ante implied on future realized return persistence is detected. It is demonstrated that this pay-off can be replicated from a portfolio of variance swaps, reducing the term-structure of variance risk premia down to a single quantity termed 'fractal risk premium'. This portfolio enables a direct bet on the degree of random walk. It is empirically found to negatively correlate with market fear; a feature which makes it particularly interesting for hedging against crashes.


1. Introduction

"Three states of matter - solid, liquid, and gas - have long been known. An analogous distinction between three states of randomness - mild, slow, and wild - arises from the mathematics of fractal geometry." - B.B. Mandelbrot

The Hurst exponent H determines the self-similarity of returns, closely tied to the degree of randomness,

H $\begin{cases} > 0.5 \dots \text{trending (slow; positive auto-corr.)} \\ = 0.5 \dots \text{random walk (mild; no auto-corr.)} \\ < 0.5 \dots \text{anti-persistent (wild; negative auto-corr.)} \end{cases}$

. The closer returns are to $H = 0.5$, the higher the degree of randomness. Without randomness, there is no risk, hence H plays a crucial role in financial markets. Since asset prices are determined by investor expectations, a natural economic interest lies upon the expectation for H . The concept of self-similar returns is easily explained visually (Fig.1).

 wolfgang.schadner@unisg.ch (W. Schadner)

 www.sites.google.com/view/wolfgang-schadner/ (W. Schadner)

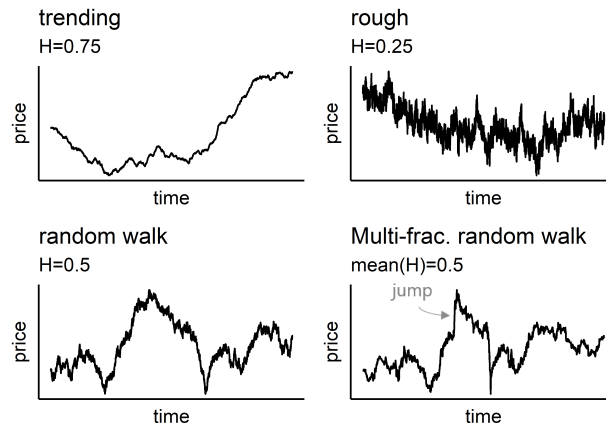


Figure 1: Simulated paths of fractal motions, H determines the self-similarity of returns (linked to randomness). When return auto-correlation is zero (fully random), then $H = 0.5$. In case H varies over time, then prices are multifractal. Typical characteristics of multifractality are volatility clustering and jump-like behavior. Broad empirical evidence shows that equity returns are significantly multifractal.

Loosely speaking, financial markets are characterized by times of mild and wild randomness. Take the time-series of the S&P 500 and break it into fractions of crises and non-crises. What you observe is high volatility during short periods of crises, and low volatility during long periods of recovery. Plot the fraction's variance against the period length, and what you observe is that variance does not scale 1:1 proportional in time. The described pattern is better known as volatility clustering, first recognized in the 1960s by Benoît Mandelbrot - the man who made fractals popular. As of today, the concept of generalized scaling laws is extensively used in overall science.¹ Fractals, in general, are used to explain a time-series self-affine structure; the synonym 'auto-correlation' might be more appealing for economists. If this structure changes over time, one is speaking of multi-fractality. If not, the series is understood as monofractal. Rough volatilities, as became recently popular, are an example for the latter. While a series' distribution is one main feature, its correlation structure is the second one. The implications are severe. In the context of finance, fractals directly affect market-efficiency, cause volatility clustering (Fig.2), determine price stability (Fig.1), open new investment strategies, are sourced in investor behavior and befall structural changes on market-crashes. The more it is important to understand a market's fractal nature, especially from an ex-ante point of view.

This work introduces the statistical mechanics required to connect forward-looking information (option prices, implied volatility) to the expected self-similarity in returns. The link is derived in a model-free manner, with a simple intuition behind it: The term-structure of variance swaps. The concept is empirically applied upon daily data of the S&P 500 index (1996-2021). Expected self-similarity (H) is found to have three remarkable characteristics. First, it strongly deviates from the critical 0.5 level for most of the time. This implies that the classic random walk assumption is misleading for explaining investor expectations. Second, there are certain patterns associated to bull- and bear-market phases. Returns are expected to be trending during growing market phases, and befall harsh drops into anti-persistence with the outbreak of fear. This brings up the third interesting pattern, that persistence and anti-persistence occasionally cancel out over an entire business cycle. Hence, investor expectations might follow a

¹For example, consider Elsevier's Scopus online library. Searching the database on 10/19/2021 in paper titles, abstract and keywords for 'fractals' returns more than 94'000 studies related to that topic, excluding the fields of mathematics and physics, still over 40'000 remain. In comparison, 'interest rate' is only used in about 32'000 papers, 'CAPM' in around 3'000.

multifractal random walk for the long run (mean of H close to 0.5). The time-varying behavior of expected randomness rises a broad variety of economic questions, where only a few of them can be addressed within this work. Among them: **(i)** Does implied persistence predict future realizations? **(ii)** What does it mean for market efficiency? **(iii)** How can persistence beliefs be traded? And, **(iv)** how does it relate to financial stability? The paper searches for answers to these questions based on novel insights.

Empirical evidence for realized fractals is very broad and reaches back to [1–4]. More recent works study fractal patterns of realized returns across different asset classes, such as stocks ([5–12]; high-frequency: [13, 14]), currencies ([7, 15]), crypto ([16–18]), fixed income ([7, 19]), commodities ([7, 18, 20]), or the fractal cross-correlation among asset classes ([21]). Fractal analyses of stock market indices are given by [22–25]. That the multifractality in returns exceeds spurious sources (finite-size, heavy-tail distribution) was documented in [6, 26–28] and many publications thereafter. Other researcher focused on the roughness of realized variance ([29–31]), spot implied volatility ([32, 33]) or the persistence of liquidity ([34]) and trading volume ([31, 35]). Fractal analysis has been also used to test for market efficiency ([19, 36–41]), to explain the behavior of sentiment ([42]), or to predict market crashes ([43, 44]). [45–47] study the breakdown in the correlation structure due to crashes in greater detail. On the theoretical side, important contributions are given by Mandelbrot’s multifractal model of asset returns ([48]), fractional Black-Scholes ([49, 50]) and the recent trend for rough volatility models ([32, 51, 52]). It is also Mandelbrot (e.g., [48]) to demonstrate that GARCH models are unfavorable in describing a return’s correlation nature. This is because either a very large number of coefficients is needed to match the true dynamics, making interpretation difficult to impossible. Or, when using only a handful of coefficients, then the long-range memory is falsely described. Either way, GARCH-like approaches require a larger number of observations for stable serial correlation estimates, hence are improper in describing current behavior (cp. [47]). Different to that, fractal methods can reduce the observation number down to a fifth of that (see [53, 54]), hence enable a more accurate measurement of current persistence. This fact is important for understanding the true market behavior.

The paper presented here is significantly distinct from the publications mentioned above, yet surprisingly simple. Methods and analyses exist for realized series; the novel contribution is to set the multifractal ideas into an ex-ante perspective. In this sense, it is demonstrated how the term-structure of implied variance can be combined with the statistical mechanics of multifractals. The output are ex-ante estimates of local Hurst exponents (current persistence, randomness) and the implied singularity spectrum (variation in Hurst exponents). Implementation of the non-parametric framework is easy and requires no specific knowledge nor skills. Comparable model-based approaches seem noncompetitive for the estimation of implied persistence. On the one hand, there is [48] (and derivations of it) which can be only solved from simulations. On the other, [49, 50] require the unrealistic assumptions of returns being mono-fractal and normally distributed. Besides the econometric perspective, the second contribution of the presented paper comes from an economic point of view, extending existing knowledge by answering questions **(i)-(iv)** based on the gathered empirical evidence. For this purpose daily data of the S&P 500 index (1996-2020) is investigated. A surprisingly strong predictability of ex-ante implied on future-realized return persistence is detected. This observation is economically puzzling because the degree of market efficiency becomes forecastable. It is shown that investors can trade corresponding beliefs from a long-short portfolio of variance swaps. This portfolio pays the difference between expected and future-realized persistence, hence is termed the ‘fractal risk premium’. The premium is negatively correlated with market fear, hence can be used to hedge against crash risk.

The reading is set up as follows. Section 2 provides an introduction to fractals and presents the methodology. Section 3 carries out the empirical analysis. Section 4 utilizes the collected insights to discuss the

economic questions. Section 5 concludes.

2. Methodology

The power of multifractal analysis is that it studies the intermittent return structure by combining both fundamental classes of analyses ([54]): (1) model based analyses of the time-series' distribution and (2) the dependency of the ordering. Class (1) analyses show that returns are heavy-tailed, assuming independent returns. Class (2) analyses show that this iid assumption does not generally hold, but assume normal distribution. Since financial markets are exposed to both effects simultaneously, they are generally better described from a multifractal point of view.

To familiarize with the concept, consider a power-law in the fashion of

$$F_q(\tau) = c(q)\tau^{h(q)}$$

. This builds the very core of multifractal formalism, which will be used to describe the geometry of investor expectations. In our context, F_q denotes the fluctuation function, directly tied to standard deviations observed throughout the sample. τ represents the time-scale, or window-size, over which the fluctuations are computed. The coefficient $c(q)$ is usually of less interest, so it is convenient to reduce the power-law to a proportional form,

$$F_q(\tau) \sim \tau^{h(q)} \tag{1}$$

² The variable q corresponds to the moment order, like $q = 2$ for variance. A special emphasis is put on the generalized Hurst exponents $h(q)$. If all $h(q)$ exponents equal the same value, then one observes a monofractal series. In case of $h(q)$ being clearly non-constant, then returns are multifractal. The second-order Hurst exponent $H \equiv h(2)$ takes on a special role, as it describes the series' persistence. Classic Brownian motion is thus a special case of constant $h(q) = 0.5, \forall q$. When $H = 0.5$ holds only on average, then one is speaking of multifractal random walk. A typical characteristic of such are price jumps and volatility clustering. This generalization can be derived from fractal formalism, which is discussed in greater detail within the next section. The estimation of F_q is presented in 2.2.

2.1. Fractal Formalism

The mathematical foundation of multifractality was significantly emphasized by Benoît Mandelbrot. Multiplicative cascades (e.g., [48, 55]) play a crucial role herein. Loosely speaking, the procedure is to take a signal's scale dependent measure (like variance or entropy) and break it into fractions over the time-line. Repeat this procedure many times and what you observe is a geometry in the fashion of Fig.2. The figure nicely visualizes the volatility clustering effect. Mathematical details about the derivation of multiplicative cascades are discussed in greater detail by [48, 55, 56]. More important for economic applications is the key result they derive. Generally, a multifractal system is described by

$$\langle |R(\tau)|^q \rangle \sim \tau^{\zeta(q)+1}$$

for a process R having stationary increments, with $\zeta(q)$ denoting the scaling function and using $\langle \cdot \rangle$ for expectations.³ This power-law relationship can be statistically assessed under different measures, like fluctuation

²Since $c(q)$ is scale independent it is frequently left out during empirical studies. Here, $c(q)$ might be economically interpreted as the base level of volatility (at $\tau = 1$). It is related to the average mass that preserves under the canonical multifractal measure. That is, the average of Fig.2. Going into technical details is clearly off-topic here, hence the reader interested might address [55, 56].

³Mandelbrot and many other denote the scaling function by ' τ '. This can be confusing in the context of financial options where ' τ ' is widely used for the time to maturity. Therefore, the scaling function is denoted by ' ζ ' instead.

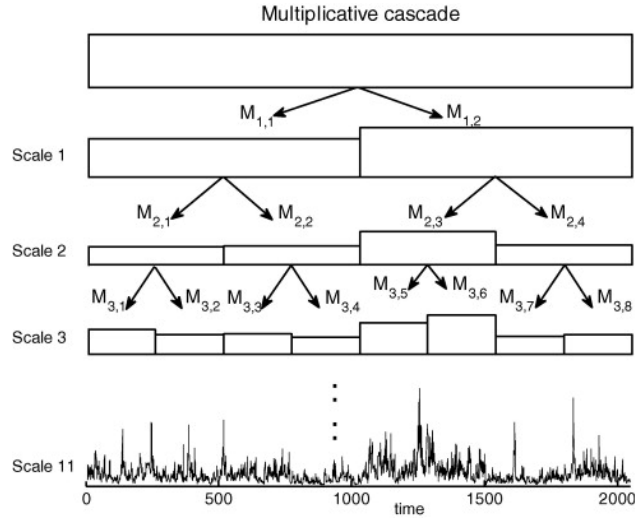


Figure 2: Multiplicative cascades: Take the scale-dependent measure (e.g., variance) and break it into fractions. Multifractality characterizes that the mass is different in these fractions. Repeat this process for smaller and smaller time scales, and the volatility clustering effect becomes visible. Picture taken from [54].

functions or entropy (see [57, 58] for examples). In case returns have a time-varying mean, the stationary assumption is no longer fulfilled. This issue can be resolved by working with a detrended measure instead. For this sake, consider the multifractal detrended fluctuation analysis (MF-DFA; [59]) from now on. As of today it is one of the most widely used techniques for multifractal analyses (cp. [11]). The superior performance of MF-DFA in capturing local correlation structures was empirically compared to other methods in [47]. In this sense,

$$F_q(\tau)^q \sim \tau^{\zeta(q)+1} \quad (2)$$

describes the multifractal behavior. The scaling-function ζ will be linear for a monofractal series (neglecting spurious sources for now). In case of multifractality, ζ is strictly convex as proven by [55, 56]. Mandelbrot also proposes the alternative wording of uni- and multi-scaling to describe mono- and multifractality.⁴ The two different scaling coefficients $\zeta(q)$ and $h(q)$ are interrelated by

$$\zeta(q) = qh(q) - 1$$

. Therefore, a linear $\zeta(q)$ implies a flat $h(q)$ for every level of q , which describes monofractality. Differently, a convex $\zeta(q)$ implies a non-constant $h(q)$ and the series will be called multifractal. As introduced earlier, the second order Hurst exponent $H \equiv h(2)$ determines the self-similarity. Substituting for $h(q)$ and taking logarithms of Eq.2,

$$\ln F_q(\tau) \sim h(q) \ln(\tau) \quad (3)$$

yields a simple linear relationship that can be easily fit by regressing $\ln F_q(\tau)$ against $\ln(\tau)$ for different levels of q . The slope coefficient then corresponds to $h(q)$.

⁴Not to be confused with 'multi-scaling' as in [60], where the word is used to describe existence of multiple fractal regimes, i.e. non-linear log-fluctuation plots.

From a mathematical perspective, the generalized Hurst exponents are only indirect measures of multifractality. Hölder singularities α are the direct ones. The corresponding density $f(\alpha)$ is called the multifractal- or singularity spectrum. The singularity measures are connected to the above quantities by

$$\alpha = \frac{\partial \zeta}{\partial q} \quad \text{and} \quad f(\alpha) = \inf_{\alpha} [\alpha q - \zeta(q)]$$

. The multifractal spectrum generally takes on a form of a smooth inverted parabola, potential shapes are discussed in [48]. Generally, the richer the multifractal strength, the more are $h(q)$ exponents dispersed and the wider the singularity spectrum becomes. Hence, the following two became common measures for multifractal strength,

$$\Delta h = h(q_{min}) - h(q_{max}) \quad \text{and} \quad \Delta \alpha = \alpha(q_{min}) - \alpha(q_{max})$$

. While the moment order q can theoretically take on any real value ($q \in \mathbb{R}$), it is empirically more reasonable to introduce some boundaries of q_{min} and q_{max} . The particular choice of the highest/lowest moment is discussed in the empirical section.

Local Behavior. Multifractality ($\Delta h > 0$ and $\Delta \alpha > 0$) implies persistence to vary over time. This is in contradiction with the classic random walk assumption (and the large body of literature building upon) assuming zero persistence throughout time. A natural economic interest now lies upon the current degree of persistence, as it has high implications on stability, investments and risk-management. H as derived from Eq.3 relates to the persistence for the overall horizon, the local one will be denoted by H_t . The term 'local' means as observed at a point in time t and the small time-window around it, hence H_t can be interpreted as the current degree of persistence.

First, it will be explained how H_t is evaluated for realized returns using existing literature. Then, the necessary step to estimate expected H_t as implied by options prices is derived. As pointed out in [54, 57], there are generally two formalisms to estimate H_t ; one via Legendre transform and one from large deviation formalism. The latter fits the power-law ' $F_t(\tau) \sim H_t \ln \tau$ ' directly for every observation date t . This is numerically less stable for realized time-series (see [57, 61]), hence the indirect Legendre transform should be generally preferred. For this sake I follow [62] for the realized return series. It comes with the great advantage that once the model is calibrated on the overall horizon, it allows to reliably estimate correlation structures on small time-scales. Hence, the routine outperforms other methods in capturing true local correlations (cp. [47]). Let c_0 denote the intercept of the fit from Eq.3 at $q = 2$. Next, choose a new set of time-scales τ' . The new time-scales should be small enough to properly capture local fluctuations, the parameter choice is discussed in greater detail at the empirical section. For all τ' , estimate the new local fluctuations $F_t(\tau')$ which can be interpreted as root-mean-squared-errors. The local-in-time $H_t(\tau')$ are then geometrically interpolated from

$$H_t(\tau') = \frac{c_0 + H \ln \tau' - \ln F_t(\tau')}{\ln n - \ln \tau'} + H \quad (4)$$

, which spans a time-conditional distribution of $H_t(\tau')$ exponents. Note that the large deviation formalism links the density of all $H_t(\tau)$ exponents (across t and τ) to the multifractal spectrum. Hence, it is less surprising that [62] argue that the width of the time-conditional distribution of $H_t(\tau)$ represents local multifractality. The mean of this density relates to the numerically stable estimate of current persistence:

$$H_t = \text{mean}\{H_t(\tau')\} \quad (5)$$

In the case of option implied information, the indirect estimation via Legendre formalism is less appealing as the forward-looking nature gets lost. However, and this is good news, fluctuation functions are directly traded in the form of variance swaps (or indirectly via options). Hence, implied fluctuation functions will be numerically more stable than realized once. This is because otherwise arbitrageurs will monetize noise from expectations. Empirical support for the stability argument is given at section 3. From this follows, that the direct large deviation formalism is occasionally more reliable here. Therefore, data as observed on t can be used to evaluate current persistence directly from a linear fit of

$$\ln F_{t,\mathbb{Q}}(\tau) \sim \langle H_t \rangle \ln \tau \quad (6)$$

with $F_{t,\mathbb{Q}}(\tau)$ referring to the square-root of implied variance, discussed further below. For example, when using simple OLS to assess Eq.6, then the regression coefficient $\langle H_t \rangle$ is the mean of the time-conditional $H_t(\tau')$ distribution.

Multiple Regimes. If returns follow dynamics of even greater complexity, then one fractal regime might not be sufficient in describing the geometry of returns. Such a phenomenon can be detected by investigating the $\ln F_q$ versus $\ln \tau$ plot. If, after correcting for spurious sources, it evolves non-linear across the scale range τ , then the underlying process is likely to be driven by multiple fractal regimes. For an unbiased description of the overall dynamics it is then necessary to detect the cross-overs and separate the different regimes (see e.g., [63]). A finance example is given by [42], who documents that market sentiment is driven by an anti-persistent regime which is overlapped by noise at very large scales. As for the returns of the S&P 500, however, this work finds that the fractal geometry is best described by a single regime. Hence, the scaling law of Eq.1 is appropriate for all reasonable horizons τ . A piece wise definition of $h(q)$ exponents for different ranges of τ would relax this assumption, but is not necessary for the application at hand.

2.2. Estimating Fluctuation Functions

Let S_t denote the price of an asset as observed on time $t \in \{1, \dots, n\}$. The log-return for one time-step (e.g., daily) ahead is thus given as

$$r_t = \ln S_{t+1} - \ln S_t$$

. Obviously, r_t is a noise like process. Using log-returns instead of discrete ones is not necessary, but it will evolve handy at our application because the aggregate sequence $R_t(\tau)$ from simple summation,

$$R_t(\tau) = \sum_{k=t}^{t+\tau} r_k \quad (7)$$

for a time-window of τ will equal the period return from t to $t + \tau$, $R_t(\tau) \equiv \ln S_{t+\tau} - \ln S_t$. Drawing parallels to the notation of the MF-DFA founders ([59]), one shall observe that t refers to the segment and τ to the segment-size. Consider expectations $\langle \cdot \rangle$ can be taken under the physical \mathbb{P} or the risk-neutral \mathbb{Q} probability measure, dependent on the context (realized or implied). Log-returns almost surely have a drift component $\langle r_t \rangle \neq 0$ that also varies over time. Hence, r_t is likely to be non-stationary to some degree. At this point, the local de-trending of DFA comes into play,

$$\text{detrending:} \quad R_t(\tau) - \langle R_t(\tau) \rangle$$

to make returns locally stationary. For realized returns, it is usually sufficient to estimate $\langle R_t(\tau) \rangle$ from a linear fit. This is equivalent to substituting r_k in Eq.7 by $r_k - \langle r_k \rangle$, using the arithmetic mean over $[t, t + \tau]$ for the latter term. Higher-order or other de-trending routines can be also applied, see [59] for a discussion.

Older methods that do not remove local trends were proven to be biased for non-stationary series (see [2]), an example is the still used (e.g., [33]) R/S analysis based on [64]. After de-trending, the time-local fluctuation function $F_t^2(\tau)$ is now given by

$$F_t^2(\tau) = \langle (R_t(\tau) - \langle R_t(\tau) \rangle)^2 \rangle \quad \text{and} \quad F_t(\tau) = \sqrt{F_t^2(\tau)} \quad (8)$$

. For realized series, this is calculated from

$$\text{realized: } F_t^2(\tau) = \frac{1}{\tau} \sum_{k=t}^{t+\tau} (R_k(\tau) - \langle R_t(\tau) \rangle)^2$$

At this point, the crucial step shall be observed that establishes the link between the statistics to the finance literature. The future price development of S is unknown ex-ante, making it infeasible to approximate individual returns a priori. Since [65], however, finance literature came up with the theoretical foundations for risk-neutral densities as implied by currently traded option prices. This allows to compute ex-ante densities for $R_t(\tau)$. From the reasoning of such densities, risk-neutral moments can be derived. The (probably) most prominent example of such is the CBOE implied volatility index, short 'VIX' representing one-month-ahead expectation. The key idea behind this and similar methodologies is to derive risk-neutral moments from the economic value of contracts replicating the pay-off to the n th power. Following [66], for example, the value of the volatility contract can be generally defined as $R_t(\tau)^2$. And consequently, the definition of risk-neutral variance evolves as

$$\text{R.N. Variance}_t(\tau) = \langle (R_t(\tau) - \langle R_t(\tau) \rangle_{\mathbb{Q}})^2 \rangle_{\mathbb{Q}}$$

. It is now easy to see that risk-neutral variance is the ex-ante equivalent to the fluctuation function of Eq.9,

$$\text{R.N. Variance}_t(\tau) \equiv F_{t,\mathbb{Q}}^2(\tau)$$

. This builds the bridge between non-linear dynamics and measurable investor expectations.

So to speak, realized fluctuations are computed from log-returns r_t , and forward-looking ones from option implied variance. The subsequent steps are common again. Once time-local fluctuations for different levels of τ are estimated, they need to be transformed into q -order fluctuations $F_q(\tau)$ with q denoting the particular moment. For this purpose, the sequences of $F_t^2(\tau)$ will be broken down into segments of size τ . The original routine of [59] chooses the segments to not overlap for the sake of computational effort (step-size = τ). Note that non-overlapping segments make the estimates of $F_q(\tau)$ sensitive to the particular choice of τ levels. Hence, using maximally overlapping windows instead (i.e., step-size = 1) was found to have clear advantages with respect to robust estimation of $F_q(\tau)$ as well as the detection of non-linearities within (see [42, 58, 60, 67]). Thought researchers exposed to very long sequences might prefer non-overlapping segments due to computation time; the procedure below presents the more robust fully overlapping ones. In this sense, moment wise fluctuations follow the statistical mechanics of

$$F_q(\tau) = \sqrt[q]{\frac{1}{n} \sum_{t=1}^n (F_t(\tau))^q} \quad (9)$$

. Having estimates of the q -wise fluctuations enables to investigate the correlation dynamics in greater detail.

2.3. Trading Fractal Risk

Finance literature put much effort into studying the variance risk premium, which is basically defined as the difference between realized and risk-neutral log-return variation. The standard instrument to trade such exposure is via variance swaps. Innovations in that field point to the works of [68] and [69], who, independently of each other, demonstrate that a portfolio of OTM options and futures can be constructed to replicate the value of a variance swap. Further, the term structure of the variance risk premium was recently studied by [70]. Their conclusion, the premium term-structure varies over time and is dependent on market phases. As discussed above, risk-neutral variance of log-returns is the direct representative of the expected fluctuation function. Hence, the term-structure of variance risk premia can be reduced down to fractal risk exposure. The trading of such exposure can be achieved as follows.

A standard variance swap (i.e., 'the log-contract' [71]) tracks the entropy of log-returns, and not the variance directly. When assuming stationary returns, one could thus replace the fluctuation functions directly with the (de-annualized) prices of such variance swaps. In this case, MF-DFA is basically replaced by the entropy scaling analysis, see [57]. In the presence of jumps and other effects, however, it is known that the standard variance swap will be a biased proxy of risk-neutral variance. For this sake, consider an alternative variance swap (like the one of [72]) which replicates the centered second moment (Eq.8). Generally, a variance swap exchanges realized variation for some pre-specified strike \tilde{V} . Under the risk-neutral measure, the value of the variance swap is zero at inception such that \tilde{V} equals the \mathbb{Q} -expectation for future realized variation. In practice the swap contracts are quoted in annualized terms, hence risk-neutral fluctuation functions are equivalent to buying τ units of \tilde{V} , $V_t(\tau) := \tau \tilde{V}_t(\tau) \equiv F_{t,\mathbb{Q}}^2(\tau)$. Since $V_t(\tau)$ represents a traded price it is not put into expectation brackets. From the multifractal formalism we know that

$$V_t(\tau) = V_t(1) \tau^{2\langle H_t(\tau) \rangle},$$

$$\ln V_t(\tau) = 2\langle H_t(\tau) \rangle \ln \tau + \ln V_t(1)$$

and

$$\langle H_t(\tau) \rangle = \frac{1}{\ln 2\tau} [\ln V_t(\tau) - \ln V_t(1)] \quad (10)$$

. $V_t(1)$ can be interpreted as the price of a short-term variance swap. From the large deviation formalism we know that the same relation can be used for an approximate of realized Hurst exponent. This proxy will be denoted by $\vec{H}'_t(\tau)$. Now, let us define the horizon-specific fractal risk premium $FRP_t(\tau)$ as the difference between realized and expected Hurst exponents,

$$FRP_t(\tau) := \vec{H}'_t(\tau) - \langle H_t(\tau) \rangle$$

. Substituting for Eq.10 and rearranging terms, one observes

$$FRP_t(\tau) = \frac{1}{2 \ln \tau} \left[\ln \frac{\vec{V}_t(\tau)}{V_t(\tau)} - \ln \frac{\vec{V}_t(1)}{V_t(1)} \right]$$

. So, the term in the brackets represents the log-return of the $V_t(\tau)$ contract minus the log-return of $V_t(1)$. Therefore, $FRP_t(\tau)$ is replicated by a portfolio with a long position in a far-end variance swap and a short position in a near-end swap, overall weighted by one over $2 \ln \tau$. In real world applications, it might not be feasible to trade the very short end variance swap because it is highly affected by market frictions, requires frequent re-balancing for long-term strategies, or simply might not be available. Therefore, it is more convenient to replicate the return from a swap contract that reaches a bit further in maturity. Say, τ_{min} is the

shortest feasible horizon for a traded variance swap falling not into the mentioned pitfalls (e.g., $\tau_{min} = 1$ month). The return for $V_t(1)$ is then approximately replicated by

$$\ln \frac{\vec{V}_t(1)}{V_t(1)} \approx \tau^{-2\langle H_t \rangle} \ln \frac{\vec{V}_t(\tau_{min})}{V_t(\tau_{min})}$$

where $\langle H_t \rangle$ can be estimated from Eq.6. Hence, instead of shorting $V_t(1)$ to replicate $FRP_t(\tau)$, one can use $\tau^{-2\langle H_t \rangle}$ units of $V_t(\tau)$. From a replicating portfolio using maturities of one month and longer it is unlikely that the realized measure $\vec{H}'_t(\tau)$ reflects local correlation behavior, simply because time-windows are too large. Thus, $\vec{H}'_t(\tau)$ should be interpreted as the persistence realized over $[t, t + \tau]$ rather than current persistence (i.e., H_t).

In case of mono-fractality, $FRP_t(\tau)$ should be constant across τ , and trading the long-short portfolio is already a direct bet on the random walk of the underlying. This changes in presence of multi-scaling. As demonstrated by [62], the mean of the time-conditional $H_t(\tau)$ distribution is a more stable descriptor of local persistence, hence

$$\langle H'_t \rangle = \frac{1}{n} \sum_{i=1}^n \langle H_t(\tau_i) \rangle$$

is an sophisticated proxy for a discrete set of variance swaps. Therefore, investing into a portfolio of equally-weighted $FRP_t(\tau)$ positions,

$$FRP_t = \frac{1}{n} \sum_{i=1}^n FRP_t(\tau_i)$$

represents a more stable bet on the difference between expected and realized persistence. In an economic sense, FRP_t can thus be interpreted as a persistence risk-premium, hence the constructed portfolio can be called the 'persistence swap'. Technically speaking, the persistence swap allows traders to directly bet on the degree of random-walk. This allows investors to directly bet on market-efficiency.

3. Empirical Illustration: S&P 500

This section sketches an empirical illustration of implied fractals upon the example of the S&P 500 index. The data frequency is daily, spanning the horizon of January 1996 to December 2020. To demonstrate that the derived fractal coefficients are not just some statistical artifacts but come with clear implications, connections are drawn to realized return structures. The option data used is based on OptionMetrics and was provided by Prof. Grigory Vilkov.⁵ It contains model-free estimates of simple variance swaps ([72]). There are two reasons why I decided to use this data set. First, the data is publicly available, hence replication of this work should be easy. Second, simple variance swaps track the *centered* second moment and is thus de-trended. In case of time-varying mean returns this should deliver more reliable results than using more common entropy tracker (VIX-like contracts).⁶ Data for the realized price series of the S&P 500 is derived from Thomson Reuter's Datastream, building on adjusted closing prices.

⁵I am thankful to Prof. Grigory Vilkov, Frankfurt School of Finance. The data is publicly available at www.vilkov.net/codedata.

⁶Scaling laws of fluctuation functions using log- or simple returns are almost identical up to a minor approximation error. The difference in scaling coefficients (H) between the two is 0.0002 using realized returns over the total horizon.



Figure 3: Realized local persistence and log-price of the S&P 500 index. Shaded areas indicate major and a few minor rapid downturns. A common feature is the harsh break in the correlation structure, recognized as drops in persistence (H_t). Hence, crashes cause strong variation in H_t , with multifractality being the consequence. Note: $H_t = 0.5$ means random walk, above is trending and below is anti-persistent.

3.1. Parameter Choice

- *Time-Scale, τ .* The option dataset covers implied variance estimates with target maturities of 1, 3, 6, 9 and 12 months. This choice seems reasonable from two aspects. First, it is known that the very short-end of the volatility term structure is highly affected by noise. Hence, the lowest time-scale of one month avoids to fall inside that range. Second, the very far end of the term-structure is typically illiquid such that expectations are not properly represented. With the largest time-scale of one year this should not be an issue. The number of interim time-scales should be linked to the goodness of fit from the regressions of Eq.3; if the fit is weak, then a larger number of interim time-scales is necessary to stabilize the regression. For the application at hand, the choice of 34 interim steps turns out to be sufficient. Realized fluctuation functions are estimated from returns, hence are more exposed to heavy-tail events and other heterogeneous structures. Thus, τ for the realized series is chosen differently, following suggestions of [53]. Generally, the objective is to choose τ for the realized signal as broad as possible, but non-surprisingly there is a trade-off. Going to small runs the risk of observing an artificial dip on the left-hand-side of the F_q plot, which arises out of limited precision in the measurement of the series (e.g., prices with only two digits). The smallest time-scale chosen for this empirical study is thus 10 days. The largest time-scale is a trade-off to the estimation robustness of the fluctuation functions, the larger τ the smaller the number of observations used for Eq.3. Also, it is documented (e.g., [42]) that if the series is exposed to noise, than this causes a flat-out in the F_q plot on the right-hand-side. Hence, the particular choice of the largest time-scale was set to $\tau_{max} = \frac{n}{10}$. The proper choice of both, τ_{min} and τ_{max} , is verified upon visual inspection of the F_q plot, which is put to the appendix. The interim time-scales are constructed from log-equally spacing between τ_{min} and τ_{max} , using 34 steps. The log-equally spacing ensures that the explanatory variable in Eq.3 is equally spaced, which supports robustness of the fit.

- *Fluctuation Order, q .* Multifractal DFA principally allows to use any real number for the moment order, $q \in \mathbb{R}$. It is recommended to use both positive and negative values for q to derive both sides of the singularity spectrum. It is particularly handy to set $q_{max} = -q_{min}$ as it allows to directly investigate the symmetry of the singularity spectrum from visual inspection. The symmetry is economically interesting as it uncovers information about the leveling of small versus large fluctuations. One recognizes that the sensitivity of F_q is directly proportional to the level of q . Hence, a too large q_{max} potentially causes corrupted tails in the singularity spectrum, which can be examined from visual inspection. For this work, the proper choice of

q_{max} was empirically assessed on the realized return series. The procedure was to first estimate the singularity spectrum for a broader range of $q_{max} = 15$, and then reducing down to the range of valid tails. From this, the choice fell onto $q_{max} = -q_{min} = 7.5$. The moment orders are then chosen as $q \in [-7.5, 7.5]$ with interim step size of 0.5. This allows a smooth visualization of $f(\alpha)$.

- *Secondary Time-Scale, τ'* . Option data is per definition time-conditional. Further, by no-arbitrage principles one can assume that investor expectations are properly incorporated into current option prices at time t , such that the entire term-structure concurrently changes with updates in the market outlook. From this reasoning, I set $\tau' = \tau$ estimating time-local implied fractals. However, things are again different at realized series. The trade-off is between choosing τ' large enough to give proper estimates of $H_t(\tau')$ but also small enough to be considered local at t (see e.g., [53]). Using daily financial data, I presume one- to four trading weeks with daily interim steps to be a reasonable choice, $\tau' \in [5, 21]$. By Eq.5, the estimate of local persistence H_t is less sensitive to the choice of τ' as for example originally proposed by [53], using only one particular level of τ' instead of a range.

3.2. Anatomy of Meltdowns

A human can take a pen and draw a square, the sides and angles describe its geometry. Humans can also cause financial markets to crash; it is then fractal formalism to describe the geometry. Crashes are recognized in the distribution domain as heavy tails, the scaling of distributions tells an even richer story. A common feature of market crashes are negative price jumps, caused by non-linear human behavior ([73]). Volatility is generally considered to be mean-reverting. A jump in price causes a shock in volatility for the near future, while the expected long-end towards which it converges will be less affected. The realized fluctuation term-structure thus befalls a sudden non-parallel change, taking the form of a flattening. The coefficient H_t is strictly tied to the slope of the log term-structure, hence strongly drops when markets crash. If H_t is not constant, it means that multiple scaling coefficients are necessary to describe the distribution's scaling behavior. Hence, returns are multifractal in the presence of jumps. As for volatility, it will befall high levels during short periods of crashes, until it oscillates around a longer lasting natural level. Volatility clustering thus means that it does not scale always the same throughout time, hence multiple fractal structures exist. This simple formulae-free statistical mechanics sum up [48]'s fractal description of volatility clustering and price jumps.

Before studying implied fractals, attention is put upon the realized behavior of correlation structures throughout time. To distinguish between implied and future-realized metrics, no superscript will be used for the former, and '→' for the latter in the subsequent reading. The focus is thus on realized local persistence \vec{H}_t for now, before diving deeper into ex-ante observations. The Hurst exponent is computed according to Section 2. Routines are left aligned such that estimates represent future-realized behavior, unknown ex-ante. The output is visualized in Fig.3. This plot uncovers important insights that are in line with the above explained dynamics. When studying the series of realized persistence \vec{H}_t one recognizes that it oscillates around the critical 0.5 level. This meaning that random walk might hold on average. A pattern that can be further observed is that \vec{H}_t is mainly within the trending region (> 0.5) during bull markets, and in the 'wild-randomness' domain at downturns. Thought, this pattern is more pronounced in the last 20 year, but less clear for the times previous to the Dot-com bubble. The shaded areas at Fig.3 indicate examples of sudden draw-downs of the S&P 500 index. It is recognized that the (small and larger) crashes are associated to structural breaks in self-affinity, noticed as drops in \vec{H}_t . The plot thus allows to gather a deeper understanding of the market's realized multifractal behavior. The times of mild ($\vec{H}_t \approx 0.5$) and wild ($\vec{H}_t < 0.5$) randomness are clearly visible. The stronger \vec{H}_t deviates from random walk, the easier it is for econometric models to make

Table 1

Multifractal statistics of implied expectations and realized returns of the S&P 500 index. The Hurst exponent $H = h(2)$ determines the degree of persistence; for random-walk series, $H = 0.5$. The spreads Δh and $\Delta\alpha$ indicate how much H can vary over time, this is also known as multifractal strength. After correcting for spurious sources (FSE, heavy-tails), the multifractality remains significant.

	implied	realized
H	0.524	0.448
Δh	0.233	0.218
$h(q)_{min}$	0.397	0.367
$h(q)_{max}$	0.630	0.585
$\partial^2 \zeta(q) / \partial q^2$	-0.008	-0.007
$\Delta\alpha$	0.378	0.390
$\Delta\alpha_{spur}$	0.101	0.247
$\Delta\alpha_{nlc}$	0.277	0.143
R^2	0.999	0.998
R^2_{iid}	0.158	-0.091

accurate predictions. Thus, knowing \vec{H}_t ex-ante would come at great advantages. But how to predict \vec{H}_t ? An educated guess is to rely on the wisdom of the crowd, using H_t as implied by traded option prices.

3.3. Large-Scale Fractals: Implied vs. Realized

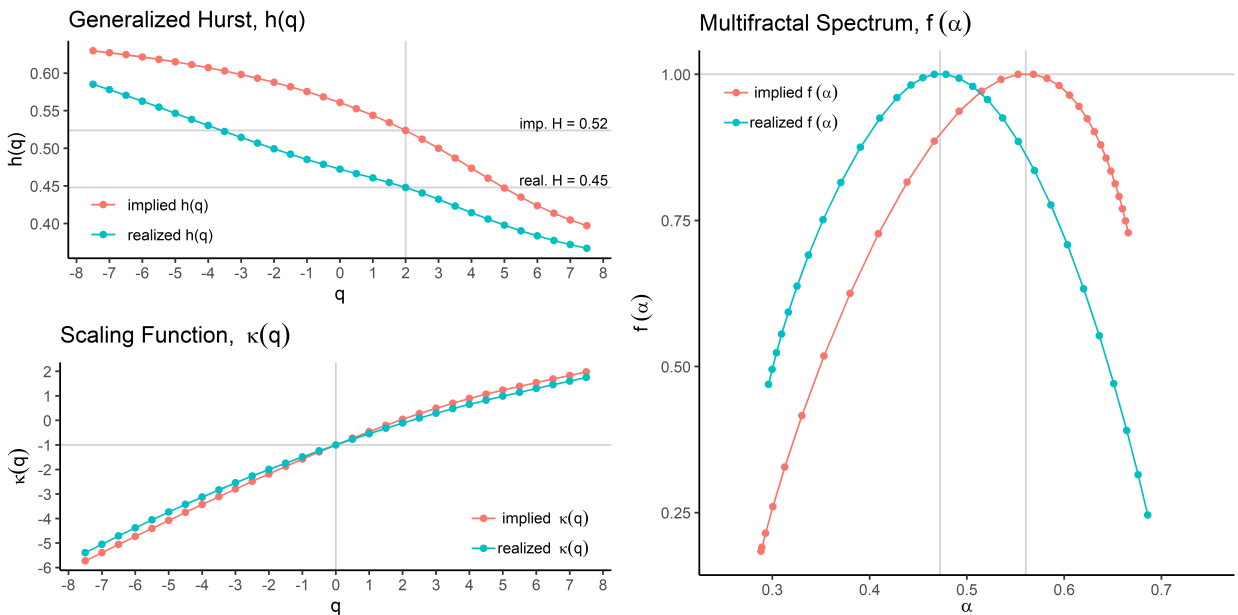


Figure 4: Visualization of the S&P 500's fractal characteristics. The red lines relate to ex-ante expectations as implied by option prices, the green ones to the statistics as derived from the index's realized returns.

The non-flatness of realized local Hurst exponent (Fig.3) demonstrates that the S&P 500 is likely to be multifractal. To examine whether the non-linear behavior is significant it is useful to build on the Legendre formalism and study fractals on the large scale. While empirical evidence is already given for realized




Figure 5: The randomness to expect: time-series of implied versus future-realized return persistence of the S&P 500 index. The overlap of the two series is very high. Hence, ex-ante persistence is a qualified predictor of future realized randomness.

returns, the more interesting question is how they compare to implied fractals. To tackle this question implied fractals are computed according to Section 2, the discussion of economic implications is put to Section 4. A good point to start with is the output of Fig.3.3. The plot of generalized Hurst exponents $h(q)$ is non-flat for the implied expectations as well as for the realized returns. Given that both $h(q)$ functions are not flat, one observes at first sight that multifractality is not only given at realized returns, but also ex-ante. The realized $h(q)$ function seems approximately linear, while the implied one features a modest curvature. Despite the shape, the slopes of the $h(q)$ functions seem very similar. What clearly differs between ex-ante and realized $h(q)$ is the intercept. It is observed that $H = 0.52$ for the implied and $H = 0.45$ for the realized signal. This means that risk-neutral expectations, more or less, believe in random-walk, while realized returns have a tendency for anti-persistence. The gap between implied and realized H potentially diminishes after correcting for the variance risk premium. The second plot of interest is the scaling function $\zeta(q)$. Both functions intersect at $\zeta(0) = -1$, which is a general feature of fluctuations. As the $h(q)$'s are not constant, it is less surprising that the scaling functions are not linear, hence also reject unscaling behavior. Special attention should be put on the curvature. Tab.1 reports the convexity of the scaling function, estimated by

$$\partial^2 \zeta(q) / \partial q^2 = \frac{[\zeta(q_{max}) - \zeta(0)][\zeta(0) - \zeta(q_{min})]}{[q_{max} - q_{min}]^2}$$

. Here we observe that both functions are convex, with the implied one showing a slightly stronger pattern than the realized ones. This is consistent with the Δh estimates and indicates that both series are multifractal. Visual inspection of $\zeta(q)$ confirms the proper choice of q_{max} . The latter is also confirmed from studying the tails of the multifractal spectra: The smooth inverted parabola shape as discussed in [55] holds up to q_{max} and down to q_{min} . The peak of $f(\alpha)$ being equal to one is an universal characteristic. That this peak is shifted to the left for the realized and more to the right for the implied expectations also indicates that the former has a lower degree of persistence. Besides the locations, the spectra width's seem equivalent, so the raw multifractal strengths are similar. A more loosely interpretation of the spectrum width is as the variation of randomness. A difference between the two is given at the skew of the distributions. The realized spectrum demonstrates a right skew while the implied one a left skew. This uncovers the leveling effect.

From a statistical point of view it should be mentioned that spurious effects exist that artificially inflate the width of the multifractal spectrum. The first of such sources is applying Eq.3 on non-linear log-fluctuation functions. There are a few reasons why $\log-F_q(\tau)$ plots can become non-linear. The most commonly known ones are insufficient de-trending ([59]), an improper choice of τ_{min} and τ_{max} ([53]) or the existence of multiple fractal regimes ([42]). Hence, inspecting the fluctuation functions visually is important. For this study,

linearity at the described setting holds; plots are attached at the appendix. Once $F_q(\tau)$ functions per se are separated from inappropriate areas, still two spurious sources remain. Those two sources are the finite size ([74, 75]) and the heavy-tail distribution effect ([76]). For realized signals, it is convenient to capture the two spurious effects from running an identical MF-DFA upon a surrogate sequence of the original returns. The surrogate sequence should be transformed such that it preserves the same degree of linear correlation (i.e., identical H) and the same distribution of x , while destroying all non-linear correlation (i.e., true multifractality). To achieve that, I follow [42] and use the iterated amplitude adjusted Fourier transform (iAAFT; [77]). The spread in singularity exponent of the surrogate sequence $\Delta\alpha_{sur}$ thus represents spurious multifractality. Deducting it from the raw spread $\Delta\alpha$ hence leaves the degree of non-linear correlation $\Delta\alpha_{nlc}$, which is the measure of pure multifractality. The pure multifractal width can thus be easily estimated for realized returns. For implied expectations without data for the incremental process, however, assessment is difficult. To come up with a sophisticated approximate I impose two assumptions. First, the horizon for implied volatilities is theoretically unbound, hence there will be no finite size effect. Second, the heavy-tail effect is the same for the implied as for realized sequence. Both assumptions might be relaxed for follow up research, but are an educated guess here. After deducting spurious spreads from raw ones, it is observed that significant multifractality as measured upon $\Delta\alpha_{nlc}$ remains for both. This implies non-linear correlation risk. It should be further emphasized that regression fits of the fractal formalism on fluctuation functions are measured by the coefficient of determination R^2 and reported in the last row of Tab.1. As can be seen, the multifractal framework explains the market behavior very well with goodness of fits larger or equal 0.998. For comparison, the iid assumption ($h(q) = 0.5, \forall q$) yields R_{iid}^2 's of 0.158 and -0.0913, hence clearly fails to properly explain market dynamics.⁷

3.4. Local Behavior and future-realized Returns

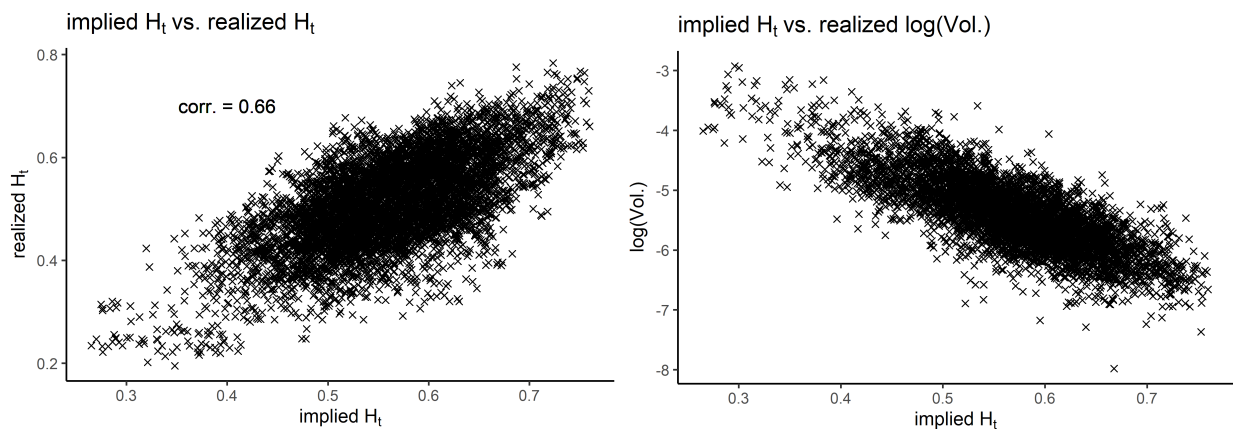


Figure 6: Time-conditional implied persistence versus future realized persistence (left) and next day variance (right).

The large-scale analysis confirms that the S&P 500 is significantly multi-fractal, with similar singularity spectra for realized returns and implied beliefs. In the economic context, however, the local measure of randomness, H_t , might be of greater interest. The analysis sketches comparisons of implied H_t to future-realized \bar{H}_t and the next-day realized volatility $\bar{\sigma}$. For the latter, data from the Oxford-Man Institute is used. It contains daily variance estimates as computed on high-frequency returns (5 minutes interval).⁸

⁷ R_{iid}^2 are computed as coefficients of determination from fitting Eq.3 with fixed $h(q) = 0.5$.

⁸ The data is publicly available at <https://realized.oxford-man.ox.ac.uk/>.

Let us first investigate the time-series of H_t and \bar{H}_t ; Fig.6 draws the plot. Just from visual inspection it is already apparent that the assumption of constant \bar{H}_t does not hold. Both series, implied and future-realized, seem to oscillate around the magical 0.5 level. Hence, random walk might hold on average. However, local persistence occasionally deviates from 0.5 quite strongly, especially during strong boom markets and crashes. It is remarkably that implied H_t closely tracks the future realized one. At this point I would like to emphasize that H_t is completely derived from ex-ante variances, while \bar{H}_t uses realized returns; two different datasets from two different sources. The more it is surprising that the both series are overlapping in such a high degree. The intuitive explanation for why this is the case may arise from arbitrage-freeness of the volatility term-structure. When liquid option prices of different maturities are free of arbitrage, then they embed an unbiased (risk-neutral) expectation of future variance. As a consequence, the implied power-law scaling will be also unbiased, hence implied H_t should closely track the future realized one. This principle explains the high Pearson's correlation of 0.663 between H_t and \bar{H}_t . A second observation is that implied H_t is on average larger than future-realized \bar{H}_t by 0.050. Such a pattern occurs when the variance risk premium is large for near-term options but decreases the longer the maturity. Hence, a 'persistence risk premium' in the sense of $H_t > \bar{H}_t$ agrees with the empirical observations (and economic arguments) of [70] that the expected variance risk premium term-structure is on average downward sloping. As discussed before, another pattern that also immediately catches the eye is the behavior around the three crises that occurred during the observation horizon (Dot-com 2002, Financial 2008, COVID 2020). H_t seems to be above average during bullish market phases, but is harshly dropping with the out-break of a crisis event. The very same statement holds for realized persistence \bar{H}_t . From this follows that implied Hurst exponents might be interpreted in terms of market mood: neutral if 0.5, bullish if above and pessimistic if below. This idea was already sketched in [78], arguing via the fractional Black-Scholes model ([49, 50]) though. Following the sentiment interpretation, growing market phases are characterized by trending returns, while pessimistic times show strongly overreacting patterns. Technically speaking, the more H_t differs from 0.5, the less random the market is and the easier it becomes for econometric models to predict short term returns. This is statistical mechanic is quite intuitive since local persistence is a measure of auto-correlation. Therefore, as ex-ante implied H_t closely tracks future realized \bar{H}_t , it is a highly qualified measure for the randomness to expect, with economic implications to market sentiment, market efficiency, econometric forecasting and investment strategies.

It is thus of particular interest to study realized return structures in greater detail. Fig.6 visualizes the plots of implied H_t versus the realized persistence and realized next-day's log-volatility. The plots allow to gather a deeper understanding of the linkages between implied fractals to future return structures. From visual inspection it is observed that, as presumed, the relation between H_t and \bar{H}_t is linear. The Pearson's correlation is high at 0.66. Given Eq.2, a linear relation is also expected between H_t and log-volatility. This is confirmed by Fig.6. Here the Pearson's correlation is even stronger at -0.82. This significant relation embeds a very high potential for implied fractals to predict the high-frequency return variance for the next day. Since there is a strong link between H_t and $\log-\bar{\sigma}$, it will mean that returns are heteroskedastic in H_t ; a fact that can be useful for econometric forecasts of short-term returns. The negative relations states that when implied H_t is low today, then there will be great intra-day return variance tomorrow. Hence, it will be more difficult to predict short-term returns for times when $H_t < 0.5$. Whenever $H_t > 0.5$, it seems that investors are safe to not expect strong negative price drops for the next day. Market crashes, characterized by strong negative price jumps, seem to happen only for times of $H_t < 0.5$. Implied multifractality thus seems to explain future risks well. The picture looks like market turmoils of high volatility are only possible

Table 2

OLS regression of implied H_t and implied volatility on future realized persistence and next-day variance. Option implied information is a significant predictor for the future degree of randomness.

	<i>dependent variable</i>			
	real. \vec{H}_{t+1}		real. $\ln \vec{\sigma}_{t+1}$	
Intercept	0.071*** (0.006)	0.071*** (0.006)	-3.536*** (0.072)	-3.745*** (0.062)
imp. H_t	0.785*** (0.011)	0.785*** (0.01)	-12.821*** (0.125)	-12.449*** (0.109)
imp. $\ln \sigma'_t$		-0.146*** (0.003)		1.650*** (0.040)
Obs.	6288	6288	5265	5265
Adj. R^2	0.440	0.575	0.666	0.748

Note:

* $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

when ex-ante H_t indicates anti-persistence for the returns to follow.

To quantify the predictive potential of implied fractals, simple OLS regressions for the linear relations as detected in Fig.6 are run. That is, ex-ante data versus future realized \vec{H}_t and $\log\vec{\sigma}_{t+1}$. Two models are used each; one using only H_t as the predictor and one additionally including the log implied volatility $\ln \sigma_t$ with a maturity of 30 days. To not fall into the multicollinearity trap, the latter is orthogonalized into $\ln \sigma'_t$ by first regressing H_t on $\ln \sigma_t$ and then deducting the fit. The results from the four regressions are presented in Tab.2. All regression coefficients turn out to be highly significant. This is less surprising given the large number of observations. The first regression, H_t on \vec{H}_t , indicates a positive intercept and the coefficient smaller one. These two patterns favor the existence of a fractal risk premium. The goodness-of-fit $R^2 = 0.44$ was already indirectly reported before, since at this setting it equals the square root of the Pearson correlation. When additionally including the implied volatility, the predictive potential as measured upon R^2 increases to 0.58. This observations allows the conclusion that future realized persistence and consequently market efficiency can indeed be predicted. The predictive regression of the ex-ante measures on the next intra-day return variance sketches an even stronger picture.

3.5. Fractal Risk Premium

This section studies the time-series of the realized fractal risk premium. For this purpose it builds on synthetic prices of simple variance swaps ([72]), hence uses the same data as above. FRP is calculated according to section 3.3. Since $\tau \in \{1, 3, 6, 9, 12\}$ months, it follows that \vec{H}'_t is an average Hurst exponent over a longer horizon, as opposed to \vec{H}_t which describes current correlation behavior. The time-series of estimated fractal risk premium is displayed in Fig.7. On average, the fractal risk-premium is negative, hence the investor taking the opposite positions makes capital gains most of the times. However, the special pattern arises that FRP_t spikes significantly when markets show sudden downturns. The pattern is less surprising considering the multifractal dynamics caused by a fast varying short-end and slow varying long-end of fluctuations. This characteristic makes the persistence swap attractive for investors seeking for insurance against unexpected market crashes. (... TBD ...)




Figure 7: Log-price of the S&P 500 versus the price of the persistence swap. The contract is designed as a long-short position of variance swaps with different maturities, monetizing the difference between implied and realized persistence (1 year). The premium is thoroughly negative during bull markets, but befalls remarkable spikes on market downturns. This feature makes the persistence swap interesting for hedging against market crashes.

4. Discussion

Coincidence?

With an eye on Fig.5, when implied persistence has been largely unnoticed so far, how does it come that it highly overlaps with future realized \bar{H}_t ? A plausible explanation is that much effort has been paid on arbitrage-free calibration of volatility dynamics, long before fractal methods received attention in the discipline of finance. In that sense, investment research works hard to make precise predictions for future variance over a wide domain of horizons. The S&P 500 index itself and the options written upon are highly liquid, which makes systematic mismatches costly. Hence, significant mis-beliefs are unlikely to survive for the long run. Therefore, the term-structure of implied volatilities will be highly correlated with the term-structure of future-realized volatilities. Recap that the term-structure fluctuation determines the scaling coefficient, it is thus natural that implied H_t is strongly correlated with future realized \bar{H}_t .

So to say, when ex-ante expectations are not systematically mis-estimating future realizations, implied H_t will be a good predictor of \bar{H}_t . A generally noticed fact about volatilities is their mean reverting behavior. This means that the short-end volatility has a tendency to converge towards its long-end. Hence, near-term volatility fluctuates more strongly and the far-end rather slowly. With the outbreak of an adverse events, volatilities - implied and realized - spikes up drastically for the short-end but modestly for the long one. Hence, one observes a structural break in the fluctuation scaling behavior, and a drastic drop in local Hurst exponent. Therefore, it is natural that market crashes are not just spikes in variance, but also structural breaks in the correlation function. From the mean-reverting variance it follows that local Hurst exponents will also revert back to a natural level.

Comparable Approaches

Implied Hurst exponents were previously computed for example in [78], who addressed the topic building on the fractional Black-Scholes market model ([49, 50]). This model utilizes the fractional Brownian motion, which implies two crucial assumptions. First, it assumes that returns are normally distributed, hence in contradiction with implied volatility smirks. Second, it assumes returns are monofractal, thus rejects that H_t will vary in time. Those two flaws might be disfavored for empirical

researchers. The non-parametric implied MF-DFA as proposed within this work does not require such or any further specific assumptions. Other methodological issues arise at the estimation of realized fractals. It should be mentioned that the very first empirical work that introduced the statistics of persistence is the rescaled-range (R/S) analysis of Hurst ([64]), thus the wording 'Hurst exponent'. Since then, it was recognized ([2]) that if the series is non-stationary, then estimates will be biased and a proper de-trending is obligatory. I guess most researcher agree upon the fact that mean returns are time-varying, thus building on simple R/S does not make much of a sense for financial return data. A recent example is given by the contradicting results of [33] (R/S) and [42] (MF-DFA), both running persistence analyses for the VIX index. Another flaw arises from the proper usage of the empirical toolbox. In that sense, using R/S or DFA on simple rolling windows does not properly represent the local correlation structure. This is due to two reason. Either windows are chosen to long such that local effects are blurred and smoothed out to long-term average. Or, windows are too small such that estimating correlations from a handful of observations is highly unreliable. This is why [53, 57] came up with the discussed geometrical work-around (Eqs.4). The bias of rolling window estimates is demonstrated by [47], visualizing the problematic from simulated series. They bring clear evidence that H_t under [53]'s MF-DFA approach truly captures the local correlation structures, while this is not the case for rolling window estimates (e.g., used in [43]).

From the above stated reasons it is less surprising that the MF-DFA method experienced a tremendous growth of usage within its 20 years of existence. Thought, while handy to match with ex-ante variance, the method is also not perfect. A document issue arises at a very low or even negative domain of $h(q)$ exponent, where the method might not be reliable (cp. [59]). While this is not an issue for random-walk-like series (e.g., equity returns), oscillating around $H = 0.5$, this occasionally becomes reasonable for variances that are documented to be rough.

Market-Efficiency

The early version of [79]'s EMH (efficient market hypothesis) considers market returns to follow a classic random walk. This clearly only holds for $H \approx 0.5$. The approximate sign enters here because of trading costs and other frictions, such that small deviations are not profitable to monetize. Thus, when H deviates strongly from the magical 0.5 level, then weak-form EMH will not hold. This is why several papers empirically studied EMH based on the Hurst exponent, examples are [38–40]. Much evidence points into the direction that the Hurst exponent of a financial market oscillates around 0.5 or is slightly above. It should be further highlighted that many studies up to today use outdated empirical methods that were proven to be biased for a market's non-linear dynamics. [64]'s R/S analysis is one example for this. Another one, also quite frequently used, are rolling-window routines of second-order methods (e.g, DFA, R/S). Note that [43] provide evidence that second-order methods require at least 200 observations in order to yield reliable estimates of H . This aspect is crucial, and [47] demonstrate very clearly how the rolling window routines fail to capture local correlation behavior. It is especially due to two reasons, either the the windows are chosen to small (e.g., below 200 observations), then the estimation via second-order methods runs into unreliable estimates. Or, windows are chosen to large, such that the estimates gets averaged out and do not properly represent local behavior. Similar arguments are brought in [53, 57], from which the MF-DFA approach is emphasized to be superior in that aspect. Nonetheless, from the empirical insights of the correlation patterns the fractal market hypothesis (FMH) evolved, which puts the focus mainly on the overall Hurst exponent H . The broad consensus supports the idea that large markets are information efficient, yet show significant fractal patterns. The frictions are argued to arise due to trading costs and high bid-ask spreads. Overall, however, the frictions are argued to be non-profitable and $H \approx 0.5$ holds approximately. While much of the FMH arguments focus on the H exponent itself, Mandelbrot [73] brings in a different perspective from the evidence of multifractality. It is widely accepted

that financial markets are characterized by $H \approx 0.5$ and Δh , in simple words, this is referred to as the multifractal random walk. Mandelbrot's point of view on this behavior is that even though correlation structures and changes within exist, markets are still efficient as long as the fractal patterns are unpredictable.

The empirical evidence as found above contradicts this statement by large. Ex-ante beliefs on H_t seems to predict future realized ones quite well. This evidence is thus finance-theoretically challenging. On the one hand, market efficiency causes that public information is properly represented in option prices, such that they adequately represent the expected future development. On the other hand, the term structure of implied variance uncovers that beliefs strongly deviate from the $H_t = 0.5$ level quite often, see Fig.5. This deviation together with the predictability of local correlation structure rises the potential for trading strategies, which simply use option and historical return data. As mentioned before, a feature of market-crashes is the structural break in correlation ([80, 81]), recognized as a harsh drop in local persistence. With an eye on Fig.5, one recognizes that at the harsh drops that initialize a crisis, H_t is slightly lagged compared to \bar{H}_t . Therefore, sudden jumps are still unpredictable.

To close the discussion on market efficiency, I would like to emphasize that technical analysis frequently misunderstands the concept of Hurst exponents. Fractal structures are used to argue in favor of their trading signals. The key issue mistaken herein is the difference between 'trend' and 'trending'. The former corresponds to the arithmetic mean, the latter one to fractals, both being literally decoupled from each other. Consequently, a time-series can be trending but still show a negative trend, and vice versa. Hence, this is one potential pitfall in the interpretation of Hurst exponents. The second one is also simple; there is simply no statistical nor economic reason why the Hurst exponent as realized on past returns should repeat itself in the future. Quite the opposite seems the case, considering [42]'s empirical evidence that H_t is highly anti-self-similar. Hence, using H estimates from historical returns is very likely to be useless for permanently profitable trading strategies.

Investment Opportunities

Fractal market hypothesis relaxes classic EMH under two conditions. First, Hurst exponents deviate only slightly from 0.5 such that monetizing auto-correlated returns is not profitable after trading costs and market frictions. Second, when returns are significantly multifractal, then markets are still efficient as long as Hurst exponents are unpredictable. With the empirical insights from above, the entire situation becomes puzzling. If option prices were significantly misestimating local persistence, then this would enable arbitrage opportunities in the option market. On the other hand, this implies that Hurst exponents and consequently short-term returns become predictable, especially for the times where H_t is quite different from 0.5. The empirical evidence from above points into the direction that especially the latter is the case. This raises potential for investment strategies that monetize expected auto-correlation in returns. At this point it is important to not mistake the wording 'trend' by 'trending', as they can be mathematically independent. Stock returns might have a negative trend (mean < 0) while simultaneously being trending ($H > 0.5$); the other way around is also possible. It seems that technical analysts occasionally mistake the usage of Hurst exponents as 'trend' to motivate their approach. Also, the wording of stock momentum - referring to a positive return in the previous month - forms a potential fallacy. The one is trend continuation, but this positive mean return can also happen at $H < 0.5$. This pattern can be seen for example at the early recovery period from the financial crisis in 2009. Overall, however, the picture looks like bull markets have trending returns and bear market anti-persistent ones (cp. Fig.3).

Consider now an investor who wants to monetize the predictable persistence of daily S&P 500 returns, his/her strategy could look as follows. Investments in the index are not directly possible, hence one has to

use ETFs or futures. For this sake price data of the SPDR S&P 500 ETF (SPY) is used. Based on [82] one can assume transaction costs in the size of 1.25 basis points per trade. Since implied H_t overestimates realized H_t on average, the physically expected $\langle H_t \rangle_{\mathbb{P}}$ will be computed as the implied one plus the average difference (time-expanding) between implied and realized. The investor now allocates between the risk-free rate (1 month treasury bill) and the ETF. Each day, the allocation decision is based on the information of $\langle H_t \rangle_{\mathbb{P}}$ and the one-day historical return r_t .⁹ The investor expects a positive return for the next day if (i) $\langle H_t \rangle_{\mathbb{P}} > 0.5$ and $r_t > 0$ or if (ii) $\langle H_t \rangle_{\mathbb{P}} < 0.5$ and $r_t < 0$. In such cases the investment is into the ETF. Otherwise the index is expected to have a negative return and the investment is into the risk-free rate. The performance of this strategy compared to the S&P 500 is now as follows. Run over the entire horizon (1996-2020), the index had a mean return of 8.7% p.a. while the strategy outperformed this at 10.8% and 9.2% (before/after trading costs). At the same time realized standard deviations were higher for the index at 18.9% compared to 14.3% of the strategy. This results into Sharpe ratios of 0.346 (index) versus 0.601 and 0.490 (before/after costs). Hence, the randomness timing strategy outperforms the index in absolute and relative terms. Especially when it comes to crash risk, the strategy seems to be very favorable as the max-drawdown is reduced from -0.596 (index) to -0.243 (strategy). This little example demonstrates the usefulness of option implied persistence. (another simple example: invest only if $H_t > 0.55$.)

5. Conclusion

Volatility clustering is a well established stylized fact, which was first described in the 1960s inherently via fractal geometry. Yet it is surprising that finance literature heavily sticks to the strict assumption of classic random walk. It is fact that asset returns follow highly complex dynamics, caused by investor's non-linear behavior and economic interaction effects. Hence, it is unreasonable to assume that the degree of randomness is always the same. The basic machinery of financial markets is to aggregate investor expectations into asset prices. The more it is important to understand the dynamics of such forward-looking believes. What fractal geometry basically does is to describe how fluctuations scale in horizon. This scaling behavior, in turn, maps into the correlation structure of returns, hence determines the degree of randomness. Aim of the presented paper is to show how these fundamental statistical mechanics can be applied upon the term-structure of option implied volatilities, reflecting investor's forward-looking expectations. When assuming that (liquid) option prices are free of arbitrage, then the natural consequence is that the slope of the log-term structure H_t relates to the \mathbb{Q} implied expectation of future return persistence. Though it is known that this slope varies substantially over time (especially on crashes), it contradicts the assumption that investors believe in classic random walk. Technically speaking, the greater the variability of H_t , the stronger the multi-fractality. A stylized feature of such is the anatomy of market meltdowns, characterized by a significant break in the correlation structure. This can be early detected from option implied H_t . Exposure to persistence can be achieved from a long-short portfolio of variance swaps with different maturities. This makes betting on random walk straight forward. The associated risk mark-up is coined 'fractal risk premium' FRP herein. Empirical investigation upon the S&P 500 uncovers that option implied persistence closely tracks temporal persistence of future realized returns. Hence, the randomness to expect is well captured within option data. The consequence is that the degree of future randomness is well predictable from the scaling of implied volatilities, which leaves a puzzle in terms of market efficiency. As for FRP , it is found to be thoroughly negative, except on market crashes. Hence, the persistence bet can be used as an insurance tool to hedge against surprising market downturns. The economic/behavioral reasons why and how long and short-ends of the volatility term-structure vary have been already extensively discussed in finance literature. Distinct to that, this work simply uses fractal formalism to describe the geometry of expectations, enabling a holistic description.

⁹End-of-day data is used, hence it is assumed that trades are made with the closing of markets.

References

- [1] E. E. Peters, "Fractal Structure in the Capital Markets," *Financial Analysts Journal*, vol. 45, no. 4, pp. 32–37, 1989.
- [2] A. W. Lo, "Long-Term Memory in Stock Market Prices," *Econometrica*, vol. 59, no. 5, pp. 1279–1313, 1991.
- [3] E. E. Peters, *Chaos and Order in the Capital Markets*. New York: John Wiley and Sons, 1991.
- [4] E. E. Peters, *Fractal Market Analysis: Applying Chaos Theory to Investment and Economics*. New York: John Wiley and Sons, 1994.
- [5] A. Turiel and C. J. Pérez-Vicente, "Multifractal geometry in stock market time series," *Physica A: Statistical Mechanics and its Applications*, vol. 322, pp. 629–649, 2003.
- [6] A. Turiel and C. J. Pérez-Vicente, "Role of multifractal sources in the analysis of stock market time series," *Physica A: Statistical Mechanics and its Applications*, vol. 355, no. 2, pp. 475–496, 2005.
- [7] M. I. Bogachev, J. F. Eichner, and A. Bunde, "Effect of Nonlinear Correlations on the Statistics of Return Intervals in Multifractal Data Sets," *Physical Review Letters*, vol. 99, no. 24, p. 240601, 2007.
- [8] J. Alvarez-Ramirez, E. Rodriguez, and G. Fernandez-Anaya, "Time-varying Hurst exponent for US stock markets," *Physica A: Statistical Mechanics and its Applications*, vol. 387, no. 24, pp. 6159–6169, 2008.
- [9] G. Du and X. Ning, "Multifractal properties of Chinese stock market in Shanghai," *Physica A: Statistical Mechanics and its Applications*, vol. 387, no. 1, pp. 261–269, 2008.
- [10] P. Caraianni, "Evidence of multifractality from emerging European stock markets," *PLoS ONE*, vol. 7, no. 7, 2012.
- [11] D. Grech, "Alternative measure of multifractal content and its application in finance," *Chaos, Solitons and Fractals*, vol. 88, pp. 183–195, 2016.
- [12] H.-Y. Wang and T.-T. Wang, "Multifractal analysis of the Chinese stock, bond and fund markets," *Physica A: Statistical Mechanics and its Applications*, vol. 512, pp. 280–292, 2018.
- [13] J. Kwapien, P. Oswiecimka, and S. Drozd, "Components of multifractality in high-frequency stock returns," *Physica A: Statistical Mechanics and its Applications*, vol. 350, no. 2-4, pp. 466–474, 2005.
- [14] D. Gu and J. Huang, "Multifractal detrended fluctuation analysis on high-frequency SZSE in Chinese stock market," *Physica A: Statistical Mechanics and its Applications*, vol. 521, pp. 225–235, 2019.
- [15] S. Drozd, J. Kwapien, P. Oswiecimka, and R. Rak, "The foreign exchange market: Return distributions, multifractality, anomalous multifractality and the Epps effect," *New Journal of Physics*, vol. 12, 2010.
- [16] K. H. Al-Yahyaee, W. Mensi, and S.-M. Yoon, "Efficiency, multifractality, and the long-memory property of the Bitcoin market: A comparative analysis with stock, currency, and gold markets," *Finance Research Letters*, vol. 27, pp. 228–234, 2018.
- [17] D. Stosic, D. Stosic, T. B. Ludermir, and T. Stosic, "Multifractal behavior of price and volume changes in the cryptocurrency market," *Physica A: Statistical Mechanics and its Applications*, vol. 520, pp. 54–61, 2019.
- [18] M. M. Ghazani and R. Khosravi, "Multifractal detrended cross-correlation analysis on benchmark cryptocurrencies and crude oil prices," *Physica A: Statistical Mechanics and its Applications*, vol. 560, pp. 125–172, 2020.
- [19] C. Aloui, S. J. H. Shahzad, and R. Jammazi, "Dynamic efficiency of European credit sectors: A rolling-window multifractal detrended fluctuation analysis," *Physica A: Statistical Mechanics and its Applications*, vol. 506, pp. 337–349, 2018.
- [20] Y. Shao, "Does Crude Oil Market Efficiency Improve After the Lift of the U.S. Export Ban? Evidence From Time-Varying Hurst Exponent," *Frontiers in Physics*, vol. 8, 2020.
- [21] S. J. H. Shahzad, S. M. Nor, W. Mensi, and R. R. Kumar, "Examining the efficiency and interdependence of US credit and stock markets through MF-DFA and MF-DXA approaches," *Physica A: Statistical Mechanics and its Applications*, vol. 471, pp. 351–363, 2017.
- [22] A. Z. Gorski, S. Drozd, and J. Speth, "Financial multifractality and its subtleties: an example of DAX," *Physica A*, vol. 316, pp. p.496–510, 2002.
- [23] L. Zunino, A. Figliola, B. M. Tabak, D. G. Pérez, M. Garavaglia, and O. A. Rosso, "Multifractal structure in Latin-American market indices," *Chaos, Solitons and Fractals*, vol. 41, no. 5, pp. 2331–2340, 2009.
- [24] C.-R. Dominique and L. E. Rivera-Solis, "Mixed fractional Brownian motion, short and long-term Dependence and economic conditions: the case of the S&P-500 Index," *International Business Management*, vol. 3, pp. 1–6, 2011.
- [25] A. Sensoy, "Generalized Hurst exponent approach to efficiency in MENA markets," *Physica A: Statistical Mechanics and its Applications*, vol. 392, no. 20, pp. 5019–5026, 2013.
- [26] W.-X. Zhou, "The components of empirical multifractality in financial returns," *Europhysics Letters*, vol. 88, no. 2, p. 28004, 2009.
- [27] J. Barunik, T. Aste, T. Di Matteo, and R. Liu, "Understanding the source of multifractality in financial markets," *Physica A: Statistical Mechanics and its Applications*, vol. 391, no. 17, pp. 4234–4251, 2012.
- [28] R. J. Buonocone, T. Aste, and T. Di Matteo, "Measuring multiscaling in financial time-series," *Chaos, Solitons and Fractals*, vol. 88, pp. 38–47, 2016.
- [29] Z.-Q. Jiang and W.-X. Zhou, "Multifractal analysis of Chinese stock volatilities based on the partition function approach," *Physica A: Statistical Mechanics and its Applications*, vol. 387, no. 19, pp. 4881–4888, 2008.
- [30] J. Gatheral, T. Jaisson, and M. Rosenbaum, "Volatility is rough," *Quantitative Finance*, vol. 18, pp. 933–949, jun 2018.
- [31] S. J. H. Shahzad, J. A. Hernandez, W. Hanif, and G. M. Kayani, "Intraday return inefficiency and long memory in the volatilities of forex markets and the role of trading volume," *Physica A: Statistical Mechanics and its Applications*, vol. 506, pp. 433–450, 2018.
- [32] G. Livieri, S. Mouti, A. Pallavicini, and M. Rosenbaum, "Rough volatility: Evidence from option prices," *IJSE Transactions*, vol. 50, pp. 767–776, sep 2018.
- [33] G. M. Caporale, L. Gil-Alana, and A. Plastun, "Is market fear persistent? A long-memory analysis," *Finance Research Letters*, vol. 27, pp. 140–147, 2018.
- [34] R. Yan, D. Yue, X. Chen, and X. Wu, "Non-linear characterization and trend identification of liquidity in China's new OTC stock market based on multifractal detrended fluctuation analysis," *Chaos, Solitons and Fractals*, vol. 139, p. 110063, 2020.
- [35] L. G. Moyano, J. de Souza, and S. M. Duarte Queirós, "Multi-fractal structure of traded volume in financial markets," *Physica A: Statistical*

- Mechanics and its Applications*, vol. 371, no. 1, pp. 118–121, 2006.
- [36] D. O. Cajueiro and B. M. Tabak, “The Hurst exponent over time: testing the assertion that emerging markets are becoming more efficient,” *Physica A: Statistical Mechanics and its Applications*, vol. 336, no. 3–4, pp. 521–537, 2004.
- [37] L. Zunino, B. M. Tabak, A. Figliola, D. G. Pérez, M. Garavaglia, and O. A. Rosso, “A multifractal approach for stock market inefficiency,” *Physica A: Statistical Mechanics and its Applications*, vol. 387, no. 26, pp. 6558–6566, 2008.
- [38] E. Onali and J. Goddard, “Are European equity markets efficient? New evidence from fractal analysis,” *International Review of Financial Analysis*, vol. 20, no. 2, pp. p.59–67, 2011.
- [39] L. Kristoufek and M. Vosvrda, “Measuring capital market efficiency: Global and local correlations structure,” *Physica A: Statistical Mechanics and its Applications*, vol. 392, pp. 184–193, 2013.
- [40] L. Kristoufek and M. Vosvrda, “Measuring capital market efficiency: long-term memory, fractal dimension and approximate entropy,” *The European Physical Journal B*, vol. 87, no. 162, 2014.
- [41] S. Ali, S. J. H. Shahzad, N. Raza, and K. H. Al-Yahyaee, “Stock market efficiency: A comparative analysis of Islamic and conventional stock markets,” *Physica A: Statistical Mechanics and its Applications*, vol. 503, pp. 139–153, 2018.
- [42] W. Schadner, “On the persistence of market sentiment: A multifractal fluctuation analysis,” *Physica A: Statistical Mechanics and its Applications*, vol. 581, p. 126242, 2021.
- [43] D. Grech and Z. Mazur, “Can one make any crash prediction in finance using the local Hurst exponent idea?,” *Physica A: Statistical Mechanics and its Applications*, vol. 336, no. 1, pp. 133–145, 2004.
- [44] Z. Tan, J. Liu, and J. Chen, “Detecting stock market turning points using wavelet leaders method,” *Physica A: Statistical Mechanics and its Applications*, vol. 565, 2021.
- [45] E. Green, W. Hanan, and D. Heffernan, “The origins of multifractality in financial time series and the effect of extreme events,” *The European Physical Journal B*, vol. 87, no. 6, pp. 1–9, 2014.
- [46] R. Hasan and S. M. Mohammad, “Multifractal analysis of Asian markets during 2007–2008 financial crisis,” *Physica A: Statistical Mechanics and its Applications*, vol. 419, pp. 746–761, 2015.
- [47] D. A. Bloch, “A practical guide to quantitative portfolio trading,” *Available at SSRN 2543802*, 2014.
- [48] B. B. Mandelbrot, A. J. Fisher, and L. E. Calvet, “A Multifractal Model of Asset Returns,” (Yale University), Cowles Foundation Discussion Paper No. 1164, 1997.
- [49] Y. Hu and B. Øksendal, “Fractional White Noise Calculus and Applications to Finance,” *Infinite Dimensional Analysis, Quantum Probability and Related Topics*, vol. 6, no. 1, pp. 1–32, 2003.
- [50] R. J. Elliott and J. Van Der Hoek, “A General Fractional White Noise Theory And Applications To Finance,” *Mathematical Finance*, vol. 13, pp. p.301–330, mar 2003.
- [51] C. Bayer, P. Friz, and J. Gatheral, “Pricing under rough volatility,” *Quantitative Finance*, vol. 16, no. 6, pp. 887–904, 2016.
- [52] G. Morelli and P. Santucci de Magistris, “Volatility tail risk under fractionality,” *Journal of Banking & Finance*, vol. 108, p. 105654, 2019.
- [53] E. Ihlen, “Introduction to Multifractal Detrended Fluctuation Analysis in Matlab,” *Frontiers in Physiology*, vol. 3, p. 141, 2012.
- [54] E. A. F. Ihlen, “Multifractal analyses of human response time: potential pitfalls in the interpretation of results,” 2014.
- [55] L. E. Calvet, A. J. Fisher, and B. B. Mandelbrot, “Large deviations and the distribution of price changes,” (Yale University), Cowles Foundation Discussion Paper No. 1165, 1997.
- [56] F. Maglione, “Multifractality in Finance: A deep understanding and review of Mandelbrot’s MMAR,” *University Ca’Foscari of Venice, Dept. of Economics Research Paper Series No.*, vol. 5, 2015.
- [57] E. A. F. Ihlen, “Multifractal analyses of response time series: A comparative study,” *Behavior Research Methods*, vol. 45, no. 4, pp. 928–945, 2013.
- [58] P. Mukli, Z. Nagy, and A. Eke, “Multifractal formalism by enforcing the universal behavior of scaling functions,” *Physica A: Statistical Mechanics and its Applications*, vol. 417, pp. 150–167, 2015.
- [59] J. W. Kantelhardt, S. A. Zschiegner, E. Koscielny-Bunde, S. Havlin, A. Bunde, and H. Stanley, “Multifractal detrended fluctuation analysis of nonstationary time series,” *Physica A: Statistical Mechanics and its Applications*, vol. 316, no. 1, pp. 87–114, 2002.
- [60] P. Castiglioni and A. Faini, “A Fast DFA Algorithm for Multifractal Multiscale Analysis of Physiological Time Series,” *Frontiers in Physiology*, vol. 10, p. 115, 2019.
- [61] E. A. F. Ihlen and B. Vereijken, “Multifractal formalisms of human behavior,” *Human Movement Science*, vol. 32, no. 4, pp. 633–651, 2013.
- [62] E. A. F. Ihlen and B. Vereijken, “Detection of co-regulation of local structure and magnitude of stride time variability using a new local detrended fluctuation analysis,” *Gait & Posture*, vol. 39, no. 1, pp. 466–471, 2014.
- [63] Z. Nagy, P. Mukli, P. Herman, and A. Eke, “Decomposing Multifractal Crossovers,” *Frontiers in Physiology*, vol. 8, p. 533, 2017.
- [64] H. E. Hurst, “The Problem of Long-Term Storage in Reservoirs,” *International Association of Scientific Hydrology. Bulletin*, vol. 1, no. 3, pp. 13–27, 1956.
- [65] D. T. Breeden and R. H. Litzenberger, “Prices of State-contingent Claims Implicit in Option Prices,” *The Journal of Business*, vol. 51, no. 4, pp. p.621–651, 1978.
- [66] G. Bakshi, N. Kapadia, and D. Madan, “Stock Return Characteristics, Skew Laws, and the Differential Pricing of Individual Equity Options,” *Review of Financial Studies*, vol. 16, no. 1, pp. 101–143, 2003.
- [67] P. Castiglioni, D. Lazzeroni, P. Coruzzi, and A. Faini, “Multifractal-Multiscale Analysis of Cardiovascular Signals: A DFA-Based Characterization of Blood Pressure and Heart-Rate Complexity by Gender,” *Complexity*, p. 4801924, 2018.
- [68] P. Carr and D. Madan, “Towards a theory of volatility trading,” *Option Pricing, Interest Rates and Risk Management, Handbooks in Mathematical Finance*, pp. 458–476, 2001.
- [69] K. Demeterfi, E. Derman, M. Kamal, and J. Zou, “A guide to volatility swaps,” *Journal of Derivatives*, vol. 7, pp. p.9–32, 1999.
- [70] Y. Ait-Sahalia, M. Karaman, and L. Mancini, “The term structure of equity and variance risk premia,” *Journal of Econometrics*, vol. 219, no. 2, pp. 204–230, 2020.

- [71] A. Neuberger, "The log contract," *Journal of portfolio management*, vol. 20, p. 74, 1994.
- [72] I. Martin, "Simple variance swaps," tech. rep., 2011.
- [73] B. B. Mandelbrot, *Multifractals and 1/f Noise: Wild Self-Affinity in Physics (1963–1976)*. Springer, 2013.
- [74] D. Grech and G. Pamuła, "On the multifractal effects generated by monofractal signals," *Physica A: Statistical Mechanics and its Applications*, vol. 392, no. 23, pp. 5845–5864, 2013.
- [75] G. Pamuła and D. Grech, "Influence of the maximal fluctuation moment order q on multifractal records normalized by finite-size effects," *Europhysics Letters*, vol. 105, no. 5, p. 50004, 2014.
- [76] D. Rak Rafałand Grech, "Quantitative approach to multifractality induced by correlations and broad distribution of data," *Physica A: Statistical Mechanics and its Applications*, vol. 508, pp. 48–66, 2018.
- [77] T. Schreiber and A. Schmitz, "Surrogate time series," *Physica D: Nonlinear Phenomena*, vol. 142, no. 3, pp. 346–382, 2000.
- [78] W. Schadner, "An idea of risk-neutral momentum and market fear," *Finance Research Letters*, vol. 37, p. 101347, 2020.
- [79] E. F. Fama, "Efficient Capital Markets: A Review of Theory and Empirical Work," *The Journal of Finance*, vol. 25, no. 2, pp. 383–417, 1970.
- [80] C. A. Los and R. M. Yalamova, "Multi-fractal spectral analysis of the 1987 stock market crash," *International Research Journal of Finance and Economics*, vol. 1, no. 4, pp. 106–133, 2006.
- [81] F. M. Siokis, "Multifractal analysis of stock exchange crashes," *Physica A: Statistical Mechanics and its Applications*, vol. 392, no. 5, pp. 1164–1171, 2013.
- [82] CME Group, "The Big Picture: A Cost Comparison of Futures and ETFs," tech. rep., 2016.
- [83] D. K. Backus, S. Foresi, and L. Wu, "Accounting for Biases in Black-Scholes," *SSRN Electronic Journal*, sep 2004.
- [84] J.-C. Duan and J. Wei, "Systematic Risk and the Price Structure of Individual Equity Options," *The Review of Financial Studies*, vol. 22, no. 5, pp. 1981–2006, 2009.
- [85] S. Aboura and D. Maillard, "Option Pricing Under Skewness and Kurtosis Using a Cornish–Fisher Expansion," *Journal of Futures Markets*, vol. 36, pp. 1194–1209, dec 2016.
- [86] B. Feunou, J.-S. Fontaine, and R. Tédongap, "Implied volatility and skewness surface," *Review of Derivatives Research*, vol. 20, no. 2, pp. 167–202, 2017.
- [87] S. Mixon, "What Does Implied Volatility Skew Measure?," *Journal of Derivatives*, vol. 18, no. 4, pp. p.9–25, 2010.

Appendix

A1 - Fluctuation Plots




Figure 8: Fluctuation functions F_q versus time scale τ per moment q for implied volatilities (left) and realized returns (right). The linear relationship holds well throughout τ , hence a single fractal regime properly describes the expected and realized return dynamics.

A2 - Rough Volatility

The methodology also applies to recent developments of rough volatility option pricing models, which are found to have a very high fit to empirical data at a low number of parameters. The key idea behind those models is that the process of log-variance is assumed to follow a fractional Brownian motion, which is the monofractal generalization of classic Brownian motion. The empirical evidence points into the direction that variance is typically highly anti-persistent at $H < 0.5$, hence the term "rough" evolved.¹⁰ Models that recently emerged from the literature are for example the rough Heston or the rough Bergomi. There are two ways to assess implied multifractality of the variance process. The first one is to use term-structures of volatility options or the CBOE Volatility-of-Volatility index to run the very same procedure as described above. The second routine presented here emerges from the scaling behavior of implied skewness. For this purpose let us focus on a metric that is frequently used to motivate rough volatility: the at-the-mones (ATM) volatility skew $\psi_t(\tau)$. It is defined as

$$\psi_t(\tau) = \left. \frac{\partial \sigma_{BS,t}(\tau)}{\partial k} \right|_{k=0}$$

with σ_{BS} as Black-Scholes implied volatilities and k as the log forward moneyness. Note that σ_{BS} represents an iid-annualized measure, hence $\sigma_{BS}\sqrt{\tau}$ is the Black-Scholes standard-deviation from 0 to τ , and $\tilde{\psi}_t(\tau) :=$

¹⁰When simulating fractal Brownian motions, the picture gets smoother the larger $H - 0.5$ and it becomes rougher for series $H < 0.5$. See Fig.1.

$\psi(\tau)_t \sqrt{\tau}$ the horizon specific ATM skew. As demonstrated in [51], if volatility is fractal, then $\psi_t(\tau)$ follows a power-law in the fashion of

$$\psi_t(\tau) \sim \tau^{H_t-0.5} \implies \bar{\psi}_t(\tau) \sim \tau^{H_t}$$

. What the above mentioned models all assume is that volatility is mono-fractal, which requires $\forall t : H_t = H$. The generalization to multifractality can now be made in a similar fashion as above. Consider

$$\langle \bar{\psi}_q(\tau) \rangle_q = \sqrt[q]{\langle \bar{\psi}_t(\tau)^q \rangle}$$

similarly to Eq.9, hence multifractality is determined by

$$\langle \bar{\psi}_q(\tau) \rangle_q \sim \tau^{h(q)} \tag{11}$$

when $h(q) \neq \text{const}$. The estimation of generalized Hurst exponents $h(q)$ can again be made by taking logs of Eq.11 and apply linear fits q -wise.

At this point, a deeper understanding for the coefficient $\bar{\psi}(\tau)$ should be emphasized. In particular, it corresponds to the slope of the strike-wise $\sigma_{BS} \sqrt{\tau}$ surface. Readers familiar with the estimation of risk-neutral moments might already guess the direction. It is an unspoken law that the slope of the implied volatility surface is the crucial driver of risk-neutral skewness. Following [66, 83], for example, approximating the implied volatility surface from Gram-Charlier expansion it evolves that $\psi(\tau) \approx \sigma_{BS}(\tau, k=0) \text{skew}(\tau)/6$ where *skew* denotes the risk-neutral Pearson's skewness (i.e., central and standardized). Testing the Gram-Charlier expansion, empirical support for the implied volatility surface being affine in its risk-neutral moments is given by [66, 84]. They report significant evidence that the slope is linear in *skew*. Note that the valid domain of *skew* is limited within a Gram-Charlier expansion, [85] therefore propose to use the more broader Cornish-Fisher expansion. Models correcting for the modest non-linearities in the slope-skew relationship are given by [86]. A deep discussion on the link between risk-neutral skewness and non-parametric slope proxies is given by [87]. All in all, evidence points into the direction that the slope is a deterministic function of *skew*,

$$\psi(\tau) \sim \text{skew}$$

. Hence, when using one of the above mentioned models that explicitly specify the link allows to estimate the variance process' multifractal behavior out of *skew* estimates. This fact forms the appealing connection to the rich body of literature on risk-neutral skewness. It might be of further interest for applications outside the discipline of finance.