Journal of Physics:
Complexity

PAPER • OPEN ACCESS

Scale-free to Pareto–Tsallis transitions in the distributions of waiting times: weather, sea-level, currency trading and automotive datasets

To cite this article: Paul-Adrian Gogîță *et al* 2026 *J. Phys. Complex.* **7** 035007

IOP Publishing J. Phys. Complex. 7 (2026) 035007 https://doi.org/10.1088/2632-072X/ae8aa5

Journal of Physics: Complexity

OPEN ACCESS

RECEIVED
2 April 2026

REVISED
22 June 2026

ACCEPTED FOR PUBLICATION
14 July 2026

PUBLISHED
24 July 2026

PAPER

Scale-free to Pareto–Tsallis transitions in the distributions of waiting times: weather, sea-level, currency trading and automotive datasets

Paul-Adrian Gogîță[1,2], **Tudor-Gabriel Dumitru**[1,3], **Florin-Ioan Constantin**[1,2], **Tudor-Andrei Diac**[1,4], **Alexandra-Florentina Neagoe**[1], **Mihaela-Carina Raportaru**[2] and **Alexandru Nicolin-Żaczek**[2,*]

[1] Faculty of Physics, University of Bucharest, Atomiştilor 405, Măgurele, România
[2] Institute of Space Science—INFLPR Subsidiary, Atomiştilor 409, Măgurele, România
[3] Reykjavik University, School of Science and Engineering, Menntavegur 1, Reykjavik, IS-101, Iceland
[4] Lund University, Faculty of Science, Box 118, 22100 Lund, Sweden
[*] Author to whom any correspondence should be addressed.

E-mail: alexandru.nicolin@spacescience.ro

Keywords: scale-free distributions, Pareto–Tsallis distributions, distribution of waiting times, terrestrial and space weather, fiat and crypto-currency trading, automotive data, Kullback–Leibler divergence

Abstract

We report a series of detailed statistical analyses on the distributions of waiting times pertaining to a diverse set of complex systems, including terrestrial and space weather, sea-level variations, currency trading (for both fiat and cryptocurrencies), and synthetic automotive data. Given a generic time series, we define a waiting time as the shortest time interval needed to find an entry of value of at least $A+\delta$, with δ a given threshold, after a certain entry of value A was observed. Going through the entire time series we obtain the complete set of waiting times for a specific value of δ and can determine their distribution. This distribution can be seen as a dynamic fingerprint of the process to which the time series pertains and is particularly useful to directly compare the dynamics of otherwise very different systems. To this end, we show that the aforementioned distributions have a prominent scale-free character for small values of δ for all datasets under scrutiny, while for large values of δ the observed distributions of waiting times converge to a Pareto–Tsallis distribution. We identify the threshold values δ at which this transition occurs using the goodness-of-fit indicators, and further substantiate these results by analyzing the behavior of the generalized Kullback–Leibler divergence. Our results are robust across all of the considered datasets.

1. Introduction

The surge of observational data that is now (almost) freely available both to the scientific community and to society at large has generated a subtle shift from the study of individual systems and specific phenomena to the investigation of fain-grained patterns of interconnections between seemingly unrelated systems (see [1] for an overview focused on physical sciences and [2] for a broader context). On a technical level, the existence of so many digital data sources of unprecedented volume and variability, to name only two of the defining v's of big data, coupled with the rapid developments on the side of data analysis tools has allowed for unprecedented investigations into systems that were traditionally considered to be disconnected from one another. The study of complex systems is a great example of how research topics as diverse as climate dynamics, linguistics, and financial markets, to give only a few examples, can be investigated within a single framework to determine previously unknown correlations. There is a visible asymmetry between the large number of articles dedicated to 'what is hidden inside the data' and the substantially fewer ones explaining 'why this is so' (see [3] for the overall framework). This reflects the asymmetry between the constantly increasing number of data hubs and the substantially fewer supporting models, which are generally harder to develop and validate.

The distribution of waiting times (DWTs) is one of the generic dynamical fingerprints that can be used to describe and study the properties of different physical systems. It was originally introduced for financial markets under the name of investment horizon, i.e. the shortest time interval needed for an index to vary by a given amount, and was later adapted in the context of seismic studies to determine the distributions of time intervals between quakes of specific properties using seismic data pertaining to earthquakes in Romania, Italy, the USA, and Japan [4–6], as well as data pertaining to moonquakes and marsquakes [7]. In this specific context, the waiting time was defined as the shortest time interval needed to find a quake with magnitude of at least $M+\delta$, with δ a given threshold, after a quake of magnitude M was observed. It should be noted that a closely connected distribution is that of the inter-occurrence times for seismic events, sometimes referred to as inter-event times or inter-nucleotide distances in the case of DNA structure, which provides the distribution of the times (or distances) between seismic events (or nucleotides) without imposing specific constrains on the triggering and terminal event (or nucleotide) [8, 9].

Quite remarkably, the DWTs observed for seismic events are distinctly scale-free for small values of δ, independent of the position of the epicenter, which, given the existing results on sea level fluctuations and the chaotic behavior of the logistic map, suggests a more generic behavior [10, 11]. In the case of sea-level variations the DWTs was computed using one of the largest existing data catalogs that covers more than a century of measurements in the port of Trieste and it was shown that for small-values of δ it has a distinct scale-free character [10]. Similarly, for the logistic map $x_{n+1}=r\cdot x_n\cdot(1-x_n)$, a similar scale-free like distribution was observed in the probability distribution of waiting iterations for values of r well into the chaotic regime [11].

The DWTs is, of course, just one of the relatively numerous methods which produce compact, interpretable descriptors that provide insight into the dynamics of a timeseries. While most commonly employed methods cannot be directly compared to the proposed distribution, as they do not capture both short and long-term correlations across multiple scales, we mention, however, the methods derived from the so-called recurrence plots, particularly the recurrence quantification analyses (RQA), which allow for the identification of transition points (e.g. chaos–chaos and chaos-order transitions) and anomalies (see [12] and references therein). RQA methods have been successfully used in deep-learning contexts such as the determination of emotions through convolutional neural networks [13], a topic on which we also expect future developments on side of DWTs.

Let us note that heavy-tailed statistics have a long history in complex systems, from Pareto's historic observations of non-algebraic tails [14] to nonextensive statistical mechanics, where q-exponential (Pareto–Tsallis) distributions arise beyond Boltzmann–Gibbs additivity [15]. A complementary dynamical perspective on the emergence of such tails is provided by self-organized criticality (SOC), which explains the ubiquity of scale-free behavior in driven dissipative systems [16], but similar distributions have been observed in cellular automata, continuous-time random walks, fractional kinetics, etc. Finally, the ubiquity of these heavy-tailed distributions comes with significant technical challenges on the side of fitting the empirical distributions, as direct probability distribution function (PDF) fits are biased toward the bulk and are quite sensitive to binning (see [10], particularly figure 1), whereas the gold-standard for fitting these heavy-tailed distributions relies on more refined approaches. These involve fits of the cumulative distribution function (CDF) with appropriate distance metrics and specific rules for tail selection [17, 18].

To check the previously mentioned generic behavior of the DWTs, we have carried out a series of large-scale statistical investigations over a diverse set of time series pertaining to terrestrial and space weather indicators, sea-level variations, currency trading of fiat and cryptocurrency and synthetic automotive datasets. For all these datasets we find that for small values of δ the DWTs exhibit clear scale-free like distributions over a few orders of magnitude, while for large values of δ the DWTs converge to a Pareto–Tsallis distribution, the transition being smooth. We show that a transition point can be effectively identified using the quality indicators for the scale-free and Pareto–Tsallis fits and show that around this point there is a jump in the derivative of the generalized Kullback–Leibler (KL) divergence, which is one way of measuring the differences between two given distributions and has been used in the past to signal phase-transitions [19].

The rest of the manuscript is structured as follows: in section 2 we describe the datasets used to compute the DWTs and the technical details regarding the computation of waiting times, while in section 3 we gather the bulk of the numerical results, including a brief discussion on the numerical efficiency of the current implementations. Lastly, in section 4 we present our conclusions and an outlook on future research.

2. Distributions of waiting times: datasets and implementations

2.1. Datasets

The statistical analyses reported in this article pertain to seven distinct datasets, six of which are observational datasets, while one is synthetic. The observational datasets cover the wind speed variations in the town of Geisenheim, where a weather station has been operating since 1884, the daily solar index published by the US National Oceanic and Atmospheric Administration, the so-called total electron content (TEC) in the ionosphere across Romania, which is an internal dataset of the Institute of Space Science—INFLPR Subsidiary (ISS), the sea level variations in the Port of Trieste (one of the relatively few so-called ultra-centennial time series), the BTC-USDT and EUR-USD exchange rates, and, lastly, synthetic automotive datasets internal to the TRUSTEE European Project (specifically, simulated instantaneous fuel consumption).

The dataset pertaining to wind speed variations in the town of Geisenheim is publicly available via the Meteostat project and consists of approximately 647 500 data points recorded every hour, covering the period from January 1951 to November of 2024. The first space weather dataset refers to the so-called daily solar index surveyed for a little under 38 years, from 1976 until 2014, each entry corresponding to one day. The dataset was aggregated from the data made public by the National Oceanic and Atmospheric Administration. Because this dataset is of relatively small volume (roughly 13 470 entries) with each entry having only two significant digits, it was used to calibrate the codes used in computing the DWTs on low-power devices and to check the implementations of the algorithm in the encrypted domain as part of the TRUSTEE European Project. It is also worth mentioning that a detailed analytical and numerical study of heavy-tailed distributions appearing in solar flare time series has been reported in [20], albeit with a slightly different definition of waiting times. The second space weather dataset is internal to the ISS and effectively refers to the so-called TEC in the ionosphere across Romania. TEC is of practical interest because it introduces delays in the propagation of electromagnetic signals due to parasitic reflections and refractions. TEC maps are calculated using the observable data sent by GNSS satellites to fixed ground receivers (see [21] for more in-depth technical details). This dataset covers approximately 45 430 entries recorded every 15 min.

The BTC-USDT exchange rate has been obtained with the help of the *CryptoDatum.io* API for a period of roughly 7 years. This dataset consists of more than 3 million values of minute-by-minute data, which is the largest set of data analyzed in this paper. The EURO-USD exchange rate was obtained from the financial market platform *investing.com*, a real-time data provider that covers information pertaining to more than 250 different exchange rates. This dataset provides daily recorded values for the end-of-day value on each business day from 1980 up to 2024, for a total of 11 572 entries.

A detailed description of the sea level variations in the port of Trieste can be found in [10]. While the likeness of the DWTs to scale-free distributions for low δ thresholds and to Pareto–Tsallis distributions for high δ thresholds has already been investigated for this dataset, the novelty of the analysis in the present paper concerns the generic behavior of the transition from one distribution to another.

The last dataset analyzed in the paper consists of synthetic driving data internal to the TRUSTEE project. The time series consists of a little under 1 million data points of instantaneous fuel consumption provided every second. There are certain peculiarities concerning the present data such as: heavy noise due to an inherent bias in the data generation process that may alter the Pareto–Tsaliss's heavy tail shape at high δ thresholds, and the very numerous presence of a value of 0 that naturally appears during engine breaking (or coasting while the car is in gear).

2.2. Data curation and preprocessing

In order to test the robustness of the algorithms used in generating the DWTs from the provided data, fitting heavy-tailed distributions to said DWTs and in computing the dynamical fingerprints of the transition, all datasets were first and foremost considered as is, in their raw form. This means that no further preprocessing took place, and that the curation process involved only handling missing values accordingly and rounding up the data to a common significant number of digits (for consistency with the chosen numerical granularity of the δ factor). It should be noted that all the DWTs generated from our raw datasets exhibit clear scale-free behavior at a small δ threshold.

However, as certain time series reflect underlying processes that deviate from a purely stochastic dynamics by encapsulating cyclical behaviors (e.g. day-night cycles or lunar phases), or by following an underlying monotonic trend, the data preprocessing techniques described in [10] are employed (e.g. detrending and pruning of extreme value outliers). The main reason is that, at low δ thresholds, low values of the waiting times dominate the statistic so that when the DWTs are renormalized to unity the noise and numerical artifacts present at long waiting times get smoothed out. At high δ thresholds, the

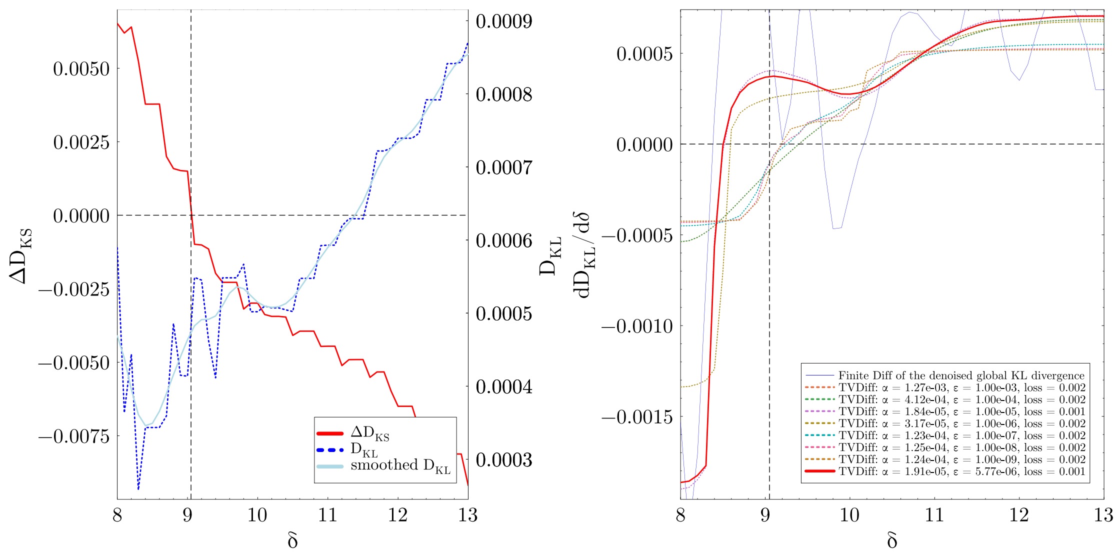

Figure 1. Transition in the DWTs for the wind speed dataset. **Left panel** shows (in solid red line) the difference in Kolmogorov–Smirnov quality of fit indicators between the two fits of the empirical distributions and (in dotted blue line) the generalized Kullback–Leibler (KL) divergence. Also shown (in light blue solid line) is the smoothed KL divergence by cubic spline interpolation. Please note that after the zero-crossing of the difference in goodness-of-fit indicators we notice an overall monotonic increase in the KL divergence. **Right panel:** showcases the derivative of the generalized KL divergence computed using the standard finite difference method over the smoothed KL divergence (see full blue line), as well as a series of derivatives computed using the TVDiff algorithm [27]. Please note that all of these derivatives have a jump-like structure in the local vicinity of the transition point.

presence of noise/numerical artifacts in the region of the heavy tail becomes more apparent and significantly influences the fitting procedure of the empirical CDF (this fact alone deems the smoothing of cyclical components by linear/logarithmic binning of the PDF unsatisfactory).

The impact of such phenomena will be shown in the figures pertaining to the port of Trieste sea level observations (preprocessed as described in [10] in order to smooth out cyclical components, eliminate the specific yearly increase in sea level, and prune out extreme outliers) and those describing the BTC-USDT exchange rate which has an inherent monotonic trend specific to long-term free-capital market exchanges (the dataset being used in its raw form for the purpose of this comparison).

The synthetic driving dataset also merits special mention. Although no rolling-average detrending is needed for this specific dataset, the inherent bias present in the data generator is responsible for producing a significant number of numerical artifacts in the form of extreme value outliers that do not decrease in statistical significance with the size of the dataset. Each present outlier in the dataset may be responsible for generating multiple erroneous waiting times at higher δ thresholds due to either prematurely ending the search between a value and its true correspondent or by providing a suitable correspondent to data points that would not generate a waiting time otherwise, on top of providing an additional data point from which a waiting time can be generated.

It is worth noting that even if the Pareto–Tsallis character of the DWTs generated from raw data at high δ thresholds may be obfuscated by the previously enumerated numerical artifacts embedded in the data, the chosen indicators (goodness-of-fit and derivative of the generalized KL divergence between the distribution tails) can still indicate a transition.

2.3. Numerical framework

The numerical framework developed for the purpose of analyzing the presented datasets covers three distinct aspects: generation of the DWTs from input data, fitting of the aforementioned heavy-tailed distributions (scale-free and Pareto–Tsallis) to the empirical DWTs and investigation of the transition from one distribution to another.

The concept of a waiting time may be described as follows: given the time series $\{A_n\}$ with associated timestamps $\{t_n\}$, with n running from 1 to N, where A_n represents an observable of interest, a waiting time is defined as the time interval elapsed between the observation of a value A_n and the following observation of a value A_m which satisfies the following condition: $A_m \geqslant A_n + \delta$, where δ is a given constant threshold. Formally, the waiting time τ corresponding to each index n can be written as follows:

$$\tau_n = t_m - t_n, \quad \text{where } m = \min\{m > n | A_m \geqslant A_n + \delta\}. \tag{1}$$

Iterating over the entire time series, i.e. spanning n from 1 to $N-1$, we obtain the complete set of waiting times for a specific value of δ and can determine their distribution. Please note that for some values of A_n there will be no associated waiting time, so the DWTs has fewer elements than the initial time series $\{A_n\}$.

This numerical recipe has been implemented in Julia [22], Python and C/C++ for benchmarking purposes, with the Julia programming language being most suited for the overall purpose of building a scientific computing framework in a composable manner. The major advantages of Julia are three-fold: the ease of prototyping and testing new ideas on the fly due to being a dynamically typed language, the optimizations benefiting compiled languages (in the present case, dealing with two nested for loops deems implementations in interpreted languages such as Python inherently slow), and finally the native capabilities for GPU acceleration or CPU multi-threading without major modifications in the established logic flow beyond avoiding race conditions, as the algorithm itself can be described as 'embarrassingly parallel'. The numerical codes used for generating DWTs have been deployed on devices of various computational capacities, from low-power devices such as the Raspberry Pi and Jetson Nano to powerful workstation PCs and distributed computing servers.

The normalized PDFs of the analyzed statistical distributions are given as:

$$p(x) = \frac{\alpha-1}{x_{\min}}\left(\frac{x}{x_{\min}}\right)^{-\alpha}; \quad x \in [x_{\min}, \infty], \ \alpha > 1, x_{\min} \geqslant 1, \tag{2}$$

for the scale-free distribution, while for the Pareto–Tsallis distribution we have considered:

$$p(x) = \frac{\alpha}{\lambda}\left[1 + \frac{x}{\lambda}\right]^{-(\alpha+1)}; \quad x \in [0, \infty], \ \alpha > 0, \lambda > 0. \tag{3}$$

It is worth acknowledging that from a strictly mathematical perspective, the scale-free and Pareto–Tsallis distributions exhibit distinct boundary constraints ($\alpha > 1$ versus $\alpha > 0$). This is a direct consequence of normalizability; the scale-free distribution possesses a singularity at $x = 0$ necessitating an $x_{\min}$ cutoff and generating a normalization factor of $\frac{\alpha-1}{x_{\min}}$, whereas the Pareto–Tsallis distribution avoids this pole via the characteristic scale λ, allowing integration over $[0, \infty]$. Furthermore, as observed in the empirical short-time behavior of several datasets, the Pareto–Tsallis form can be theoretically extended to a generalized distribution (the Generalized Beta Prime) of the form $p(x) \propto x^s[1 + (\alpha/\lambda)x]^{-(\alpha+1)}$. While this extension captures non-monotonic behavior at small scales, we strictly maintain the 2-parameter forms presented in equations (2) and (3) to preserve algorithmic stability. Computing a 3-parameter optimization over the empirical CDF introduces severe parameter coupling. In such scenarios, gradient-free solvers prioritize minimizing the dense probability mass at short waiting times, which destabilizes the algorithm and heavily compromises the accuracy of the asymptotic heavy-tail fits. Consequently, maintaining the isolated, standard analytical forms provides the most robust computational framework for evaluating the δ-driven transition in the distribution tails.

The inherent difficulty of fitting heavy-tailed distributions is discussed at length in [17], with some of the technical details concerning the implementation in the Julia language (at the time of writing) summarized in [10]. The crux of the problem lies in the fact that attempts of fitting the PDFs of such distributions lead to results that overemphasize the minimization of the loss between the data and the distribution fit at short waiting times to the detriment of the fit quality in the distribution's tail, in the region of long waiting times. This fact is clearly visible in figure 1 of [10] in the case of scale-free distributions where the influence of different types of binning procedures applied to the PDF is also shown.

Additionally, when visualizing heavy-tailed distributions on doubly-logarithmic plots, it is convenient to show the so-called complementary cumulative distribution function (CCDF) in order to better capture the tail region, that is:

$$F(x) = P(X > x) = \int_x^\infty p(X)\,\mathrm{d}X. \tag{4}$$

The numerical recipe employed for the fitting of scale-free and Pareto–Tsallis distributions of DWTs in the present paper applies the main ideas from [17, 18], namely the minimization of the maximal Kolmogorov–Smirnov distance between the data and model CDFs, and the proper renormalization of said CDFs in order to reconcile the differences between the span of the DWT data and the scale-free and Pareto–Tsallis distribution domains.

At the center of the numerical recipe lies the Julia `Optim` package [23] that is used to compute the distribution parameters with the black-box Nelder-Mead simplex algorithm described in [24–26]. We mention that in the case of the scale-free distribution a considerable sensitivity with respect to the scale parameter $x_{\min}$ is present, as $x_{\min}$ also represents the lower-bound cutoff that is required to eliminate the divergence of the normalization integral. Finally, we have checked the reliability of the fitting procedure with respect to various data pollution strategies and have found that neither changes in the sampling rate, nor the spurious random occurrence of gaps in the dataset, do not influence the computed scaling exponents in any significant fashion. The aforementioned data pollution strategies have proven the reliability of the fitting procedure even in cases when the dataset was effectively reduced down to the minimum size that would render the fitting procedure of heavy-tailed distributions numerically untractable.

Concerning the investigation of the structural transition occurring between the scale-free and Pareto Tsallis distributions, we employ two qualitatively different indicators in order to identify the δ threshold values that mark the transition for each dataset.

Firstly, we compute the difference $\Delta\mathrm{D_{KS}}$ between the goodness-of-fit indicators of the Pareto–Tsallis and scale-free distributions for a spectrum of δ values that spans from the scale-free dominant region (low δ values) to the Pareto–Tsallis dominant region (large δ values): $\Delta\mathrm{D_{KS}}(\delta) = \mathrm{D_{KS}^{PT}}(\delta) - \mathrm{D_{KS}^{s\text{-}f}}(\delta)$. The zero–crossing of the difference in KS distances provides a phenomenological marker that suggests the δ value where one model becomes a better empirical fit than the other, making this point a natural candidate for the transition threshold.

Secondly, in order to substantiate the transition identified by the difference in goodness-of-fit indicators, we employ the generalized KL divergence calculated between the tails of the best-fit scale-free and Pareto–Tsallis distributions for each considered value of δ where $\Delta\mathrm{D_{KS}}(\delta)$ has been previously obtained. Also known as relative entropy, the KL divergence $\mathrm{D_{KL}}(\mathrm{P}\parallel\mathrm{Q})$ quantifies the amount of information lost when a model distribution Q is used to approximate a true distribution P. While various other statistical distances exist in the literature, the KL divergence was specifically chosen for this framework due to its inherent sensitivity to asymptotic tail behavior. Supremum-based metrics, such as the Kolmogorov–Smirnov distance utilized for our global CDF fitting, evaluate the maximum absolute difference and are consequently dominated by the probability mass in the bulk of the distribution. By evaluating the logarithmic ratio of the distributions, the KL divergence strictly penalizes point-wise structural discrepancies in the extreme tails, making it uniquely suited to identify the phase-like transition between the scale-free and Pareto–Tsallis regimes. For continuous probability distributions $p(x)$ and $q(x)$ the KL divergence is defined as:

$$\mathrm{D_{KL}}(P\parallel \mathrm{Q}) = \int_{-\infty}^{\infty} p(x)\log\left(\frac{p(x)}{q(x)}\right)\mathrm{d}x, \tag{5}$$

while in the case of discrete data it can be computed as:

$$\mathrm{D_{KL}}(P\parallel \mathrm{Q}) = \sum_{i=1}^{n} p_i \log\left(\frac{p_i}{q_i}\right). \tag{6}$$

Furthermore, the generalization of the KL divergence extends its definition to include unnormalized probability distributions. This is particularly useful in our context because we consider that the tail region of the distributions at each δ begins at the scale-free $x_{\min}$ cut-off and ends at the last waiting time for which either of the model distributions begins to have a lower probability than the lowest probability found in the empirical DWT. In short, while each model distribution is normalized on its own corresponding domain (bounded by the empirical DWT), the common interval chosen as the tail region does not guarantee the proper normalization of each distribution at every threshold δ. In light of these remarks, we note that this generalization does not possess the propriety of being non-negative across the span of analyzed δ values. The generalized KL divergence was calculated for discrete data in the present paper as follows:

$$\hat{\mathrm{D}}_{\mathrm{KL}}(\mathrm{P} \parallel \mathrm{Q}) = \sum_{i=1}^{n} p_i \log\left(\frac{p_i}{q_i}\right) - p_i + q_i. \tag{7}$$

For ease of reading comprehension, the generalized KL divergence computed at each threshold value δ will be written from now on as simply $\mathrm{D}_{\mathrm{KL}}(\delta)$.

The numerically computed generalized KL divergence, when plotted as a function of δ, is inherently noisy (as can be seen in figure 1). Using a standard numerical differentiation method such as first-order finite differences can severely amplify noise, which would obscure any structural features of the derivative. In order to alleviate such issues, we employ the total variation regularized differentiation (TVDiff) algorithm described in [27]. In short, regularization of the differentiation process as a whole effectively suppresses spurious oscillations arising from noise, thus preserving jump discontinuities or other structural features from the differentiated function. This property is precisely what is required in order to accurately identify the jump-like structure in the derivative of the generalized KL divergence, namely $\mathrm{dD}_{\mathrm{KL}}/\mathrm{d}\delta$, which we identify as an additional, qualitatively-different indicator of the transition between the scale-free and Pareto–Tsallis distributions. The choice of this advanced algorithm is therefore not a minor technical detail, but a central pillar substantiating our main message on the position of the transition, ensuring that the observed jump-like structure is a real feature of the data and not the cumulative effect of numerical artifacts.

The hyper-parameters of the TVDiff algorithm (i.e. regularization parameter α and conditioning parameter ε) are determined in an iterative process (see figure 1) such that the maximum difference between $\mathrm{dD}_{\mathrm{KL}}/\mathrm{d}\delta$ calculated by applying the finite difference method on the smoothed $\mathrm{D}_{\mathrm{KL}}(\delta)$ and $\mathrm{dD}_{\mathrm{KL}}/\mathrm{d}\delta$ obtained from TVDiff over the whole span of δ values is minimized. All future plots in section 3 concerning the transition from a scale-free to a Pareto–Tsallis distributions will only showcase the derivative corresponding to the aforementioned parameter pair.

3. Results

In this section we gather our main numerical results on the statistical analyses of the previously discussed datasets, which we present, for clarity, one by one. Although the observed transition in the DWTs exhibits a generic character across all analyzed datasets, we purposefully maintain separate figures for each system to emphasize their distinct underlying dynamics. These domain-specific time series span vastly different orders of magnitude, exhibit diverse structural behaviors, and involve varying degrees of data curation—from raw financial data to heavily pre-processed sea-level records. Therefore, in figures 2–8 we show the individual distributions of waiting times in the limits of small- and large-values of δ along with all relevant fitting parameters, showcasing the PDF (referred to as $P_k(k)$, upper panels) and the CCDF (referred to as $F_k(k)$, lower panels). For each dataset, it is imperative to showcase both the PDF and the CCDF. The PDF representation is crucial for visually capturing localized structural behaviors, such as oscillatory features. Concurrently, the CCDF provides an integral representation that remains strictly unbiased by binning artifacts, which is the mathematically rigorous standard required for the reliable fitting of heavy-tailed distributions.

These results are summarized in a concise form in table 1 where we present the α exponent for different values of δ along with the D_{KS} quality-of-fit indicator. Given the relatively small number of considered datasets we will not speculate on the potential cross-domain grouping—akin to classes of universality—of the α exponents reported in table 1. The PDFs and the CCDFs provided in the aforementioned figures provide consistent results across the entire set of time series under scrutiny, for both small and large values of δ, with the observation that in the limit of large values of δ the PDFs are impacted by the poorer statistics (which, in turn, is due to the smaller number of waiting times) and we see small differences between the PDFs and the CCDFs at short waiting times.

We have found similar results for many other time series which we do not show here due to their close similarity with the ones that are depicted. The exchange rates pertaining to all major fiduciary currencies, for instance, share the behavior of the EURO-USD exchange rate, with the notable exception of the Chinese yuan and Russian ruble, the same way other major cryptocurrencies, like Ethereum, exhibit

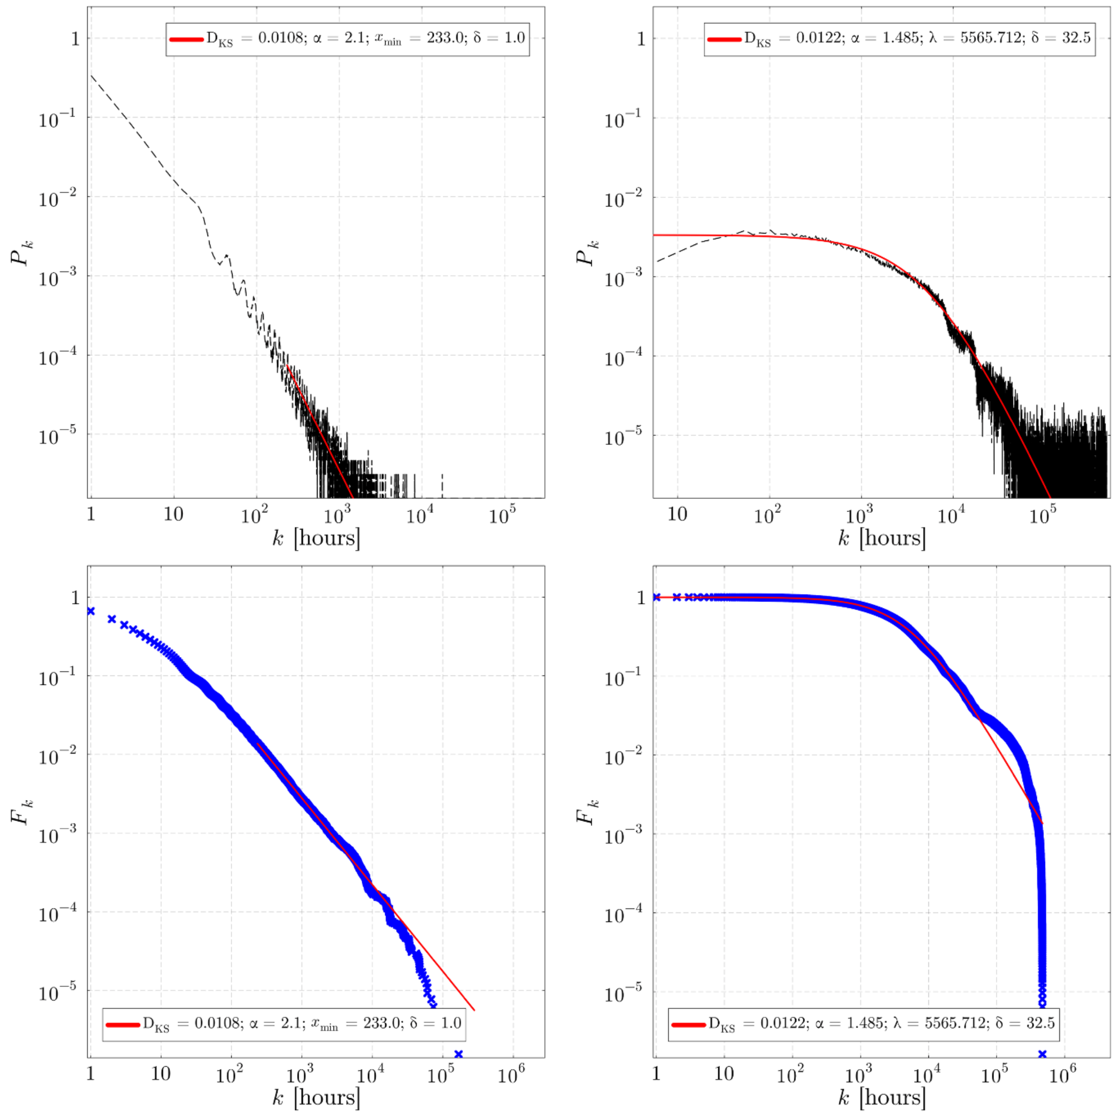

Figure 2. DWTs for wind speeds in the town of Geisenheim using the data published by Meteostat. The upper panels show the PDF results corresponding to $\delta = 1$ km h^{-1} (left side) and $\delta = 32.5$ km h^{-1} (right side), the black dashed line corresponding to the empirical data, while the red full line shows the distribution fit. The lower panels show the same DWTs but in the CCDF representation.

the same DWTs when priced against USDT. Similarly, for automotive data we seen the same DWTs for the time series pertaining to acceleration and velocities, though with slightly more noise in the distributions. In the case of terrestrial weather we have seen the aforementioned DWTs to be generic as well, though some time series like those pertaining to temperature have a very pronounced periodic day-night cycle that has to be first cleared during preprocessing. As we have shown in [10] for the sea-level variations in the port of Trieste, the DWTs are strongly impacted by the periodic components intrinsic to the time series, which is why without a careful detrending of the raw data one does not see a Pareto–Tsallis distribution in the DWT for large values of the threshold δ.

To the best of our knowledge, there is no general theory that describes the emergence of the two distributions observed in figures 2–8, though in the case of earthquakes the scale-free behavior of DWTs in the limit of small values of δ has been seen as a sign of the criticality of seismic zones [4]. This picture is supported by simulations on microscopic models like the celebrated Olami–Feder–Christensen model, which is known to exhibit SOC [28]. Moreover, there is a considerable literature that reports SOC behavior in systems pertaining to the domains from which we have extracted the datasets under scrutiny in this manuscript, but we have found no results that allow for a one-to-one comparison with our results. In the case of economy and finance the literature on self-organized processes is very broad [29], but for a direction related to our results we refer to [30] where stock market dynamics is discussed empirically. For space weather, for instance, SOC has been reported for substorm dynamics of

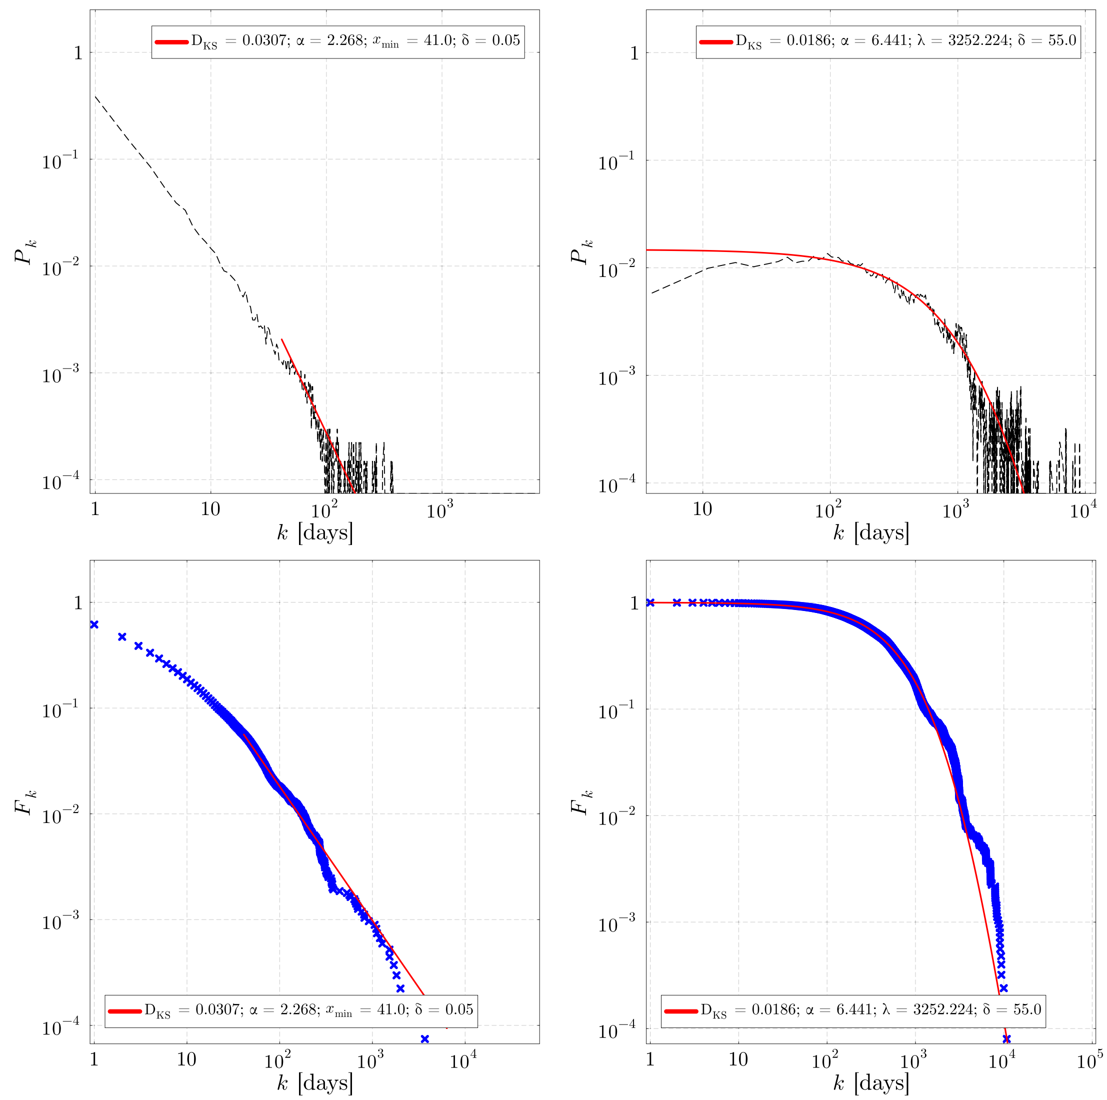

Figure 3. DWTs for the daily solar index published by the US National Oceanic and Atmospheric Administration. The upper panels show the PDF results corresponding to $\delta = 0.05$ (left side) and $\delta = 55$ (right side), the black dashed line corresponding to the empirical data, while the red full line shows the distribution fit. The lower panels show the same DWTs but in the CCDF representation.

the magnetosphere [31], while in the context of sea studies the soft-cliff retreat has been argued to also be a self-organized critical phenomena [32] and there are related results showing that tidal deltas also exhibit self-organization [33]. Complementary, atmospheric precipitation have been shown to exhibit self-organized critical features [34] and it has been shown that rain is analogous to natural relaxational processes such as earthquakes, which would suggest that the same scaling laws apply [35].

To gain a better understanding of how the DWTs change with δ, we show in figure 9 the transition from a scale-free to a Pareto–Tsallis distribution for the DWTs pertaining to wind speeds (upper-left panel), sea-level variations (upper-right panel), the total electron density in ionosphere (lower-left panel) and daily solar index (lower-right panel) datasets. In all four panels the full red curves show the difference between the KS goodness-of-fit indicators, namely ΔD_{KS}, while the blue dotted curves show the derivative of the generalized KL divergence, that is, $dD_{KL}/d\delta$. In figure 10 we depict the scale-free to Pareto–Tsallis transition of DWTs shown for the BTC-USDT (upper-left panel), the EUR-USD exchange rate (upper-right panel) and the instantaneous fuel consumption pertaining to the automotive synthetic dataset (lower-left panel), with all technical details as in figure 9.

It should be noted that around the value of δ when ΔD_{KS} crosses zero, i.e. when the two fits approximate the empirical DWT with equal accuracy, we also see a jump in the derivative of the generalized KL divergence, which is a strong indicator that this is the transition point. More importantly, the KL

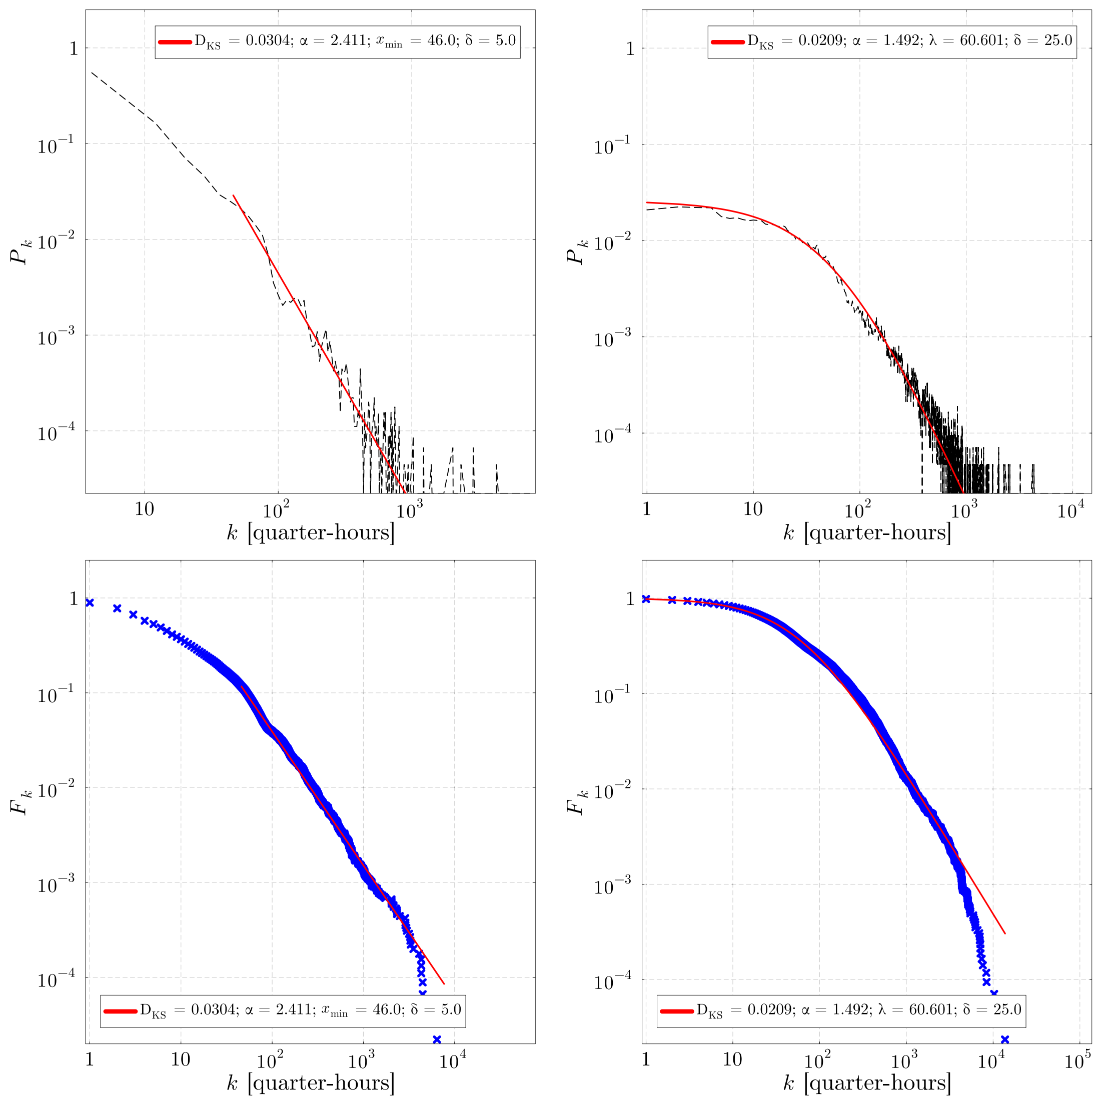

Figure 4. DWTs for the so-called total electron content in the ionosphere using a dataset internal to Institute of Space Science—INFLPR Subsidiary. The upper panels show the PDF results corresponding to $\delta = 5$ (left side) and $\delta = 25$ (right side), the black dashed line corresponding to the empirical data, while the red full line shows the distribution fit. The lower panels show the same DWTs but in the CCDF representation.

divergence was used to signal phase transitions (see, e.g. [19]), but without a supporting microscopic model we will abstain from speculations on the exact nature of this transition, referring to it simply as a structural change in the behavior of the system. As mentioned in the previous section, the numerical derivative of the generalized KL divergence is subject to some imprecision due to the numerical noise in the generalized KL divergence, but all numerical results show a clear jump in the derivative of the generalized KL divergence around the zero-crossing of ΔD_{KS}.

Knowing the $\delta-$regimes in which the observed distributions are scale-free like presents significant practical applicability, namely the clusterization of datasets based on their scaling exponents. In the case of automotive data, to give just one example, distinct datasets pertaining to distinct (real or virtual drivers) can be clustered into a few groups of driving styles, based on the values of the α exponents obtained for small values of δ. To showcase this, we depict in the lower-right panel of figure 10 the results for 25 distinct time series corresponding to synthetic automotive agents, clusterized into three distinct categories (aggressive, moderate, and calm). This classification was achieved utilizing a Fuzzy C-Means clustering algorithm applied to the scale-free exponent α evaluated at a small threshold ($\delta = 20$). Furthermore, this approach to identifying distinct agents provides a macroscopic statistical fingerprint that highly complements other non-linear time-series techniques commonly utilized for anomaly detection. For instance, while methods such as RQA are highly effective at diagnosing micro-scale structural deviations or isolated erratic maneuvers within specific, short time windows [12], the DWTs efficiently

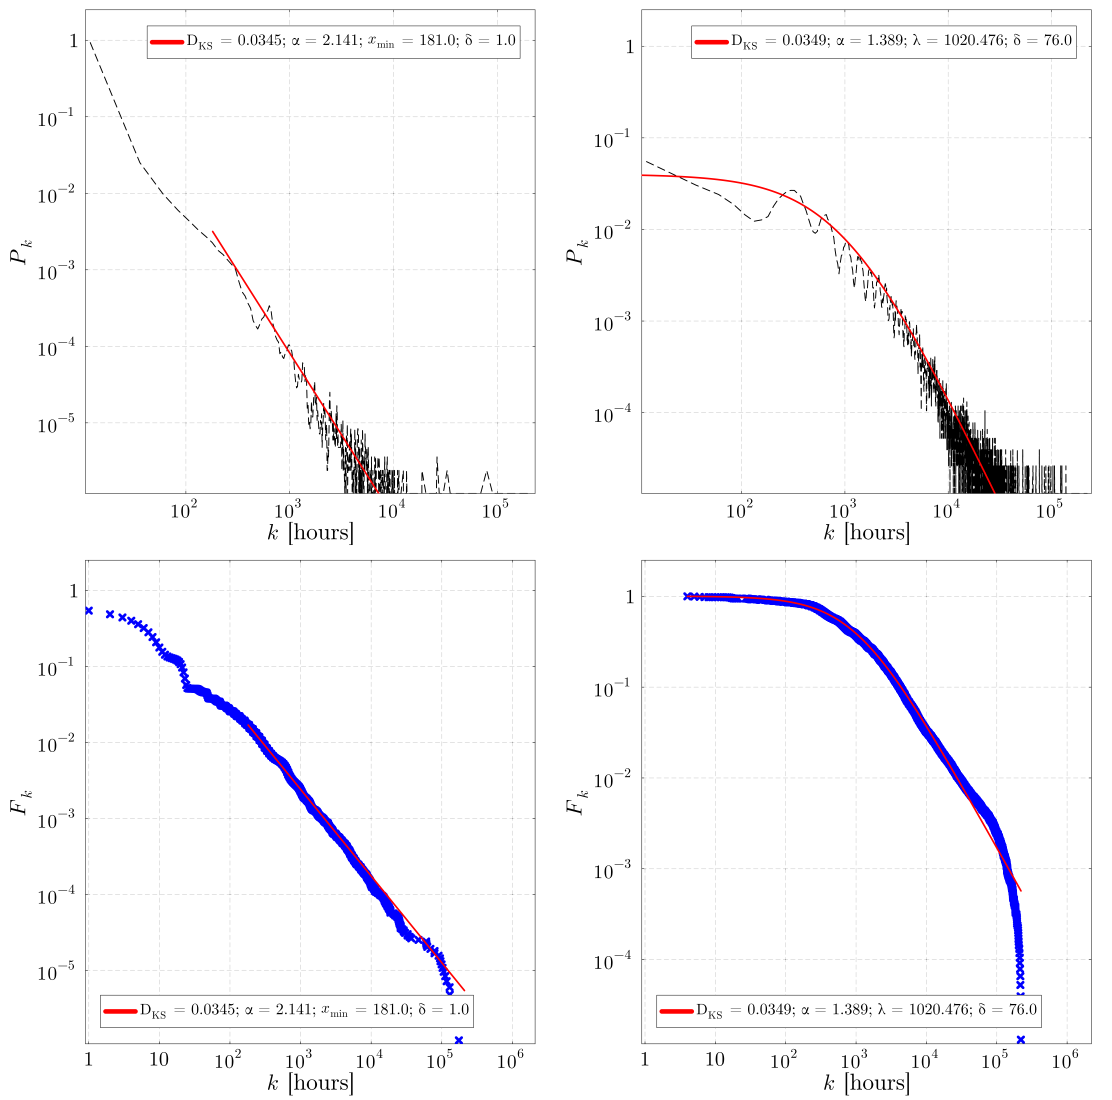

Figure 5. DWTs for sea-level variations in the Port of Trieste. The upper panels show the PDF results corresponding to $\delta = 1$ pps (left side) and $\delta = 76$ pps (right side), the black dashed line corresponding to the empirical data, while the red full line shows the distribution fit. The lower panels show the same DWTs but in the CCDF representation. For all panels δ is measured in percentage points and not centimeters, as in the raw dataset, due to the rolling average data preprocessing (see [10] where similar PDF and CCDF plots have been previously published).

capture the global degree of burstiness and long-range correlations inherent to the agent's overall profile. Consequently, leveraging the DWT for macro-scale behavioral classification alongside local, sliding-window anomaly detectors provides a robust and comprehensive framework for properly understanding the full dynamical complexity of time-series data.

More generally, one of the important questions to be answered in future studies is if the aforementioned distributions can be used not only to cluster agents based on specific indicators, as shown above, but to actually identify distinct agents in a manner similar to that in which the digital footprint can be used to identify users online. A positive answer to the aforementioned question would require one to redefine what constitutes personal data and reconsider the security provisions on the side of data processing for datasets which are not traditionally considered as personal data, such as automotive data. Lastly, on the side of forecasting, let us note that this is technically difficult in the case of small values of δ as any average which uses the DWT would rely highly on the underlying cut-offs at short and long waiting times, where the scale-free approximation does not hold.

The observed behavior with respect to δ of the aforementioned two qualitatively different indicators can be seen akin to a phase transition. Thus, if we consider δ as the control parameter and $\Delta \mathrm{D_{KS}}$ as the order parameter we have a behavior resembling that of a second-order phase transition in which $\Delta \mathrm{D_{KS}}$ crosses smoothly through zero, while if we take $\mathrm{dD_{KL}}/\mathrm{d}\delta$ as the order parameter we get an image somewhat similar to that of a first-order phase transition in which—even in the presence of numerical

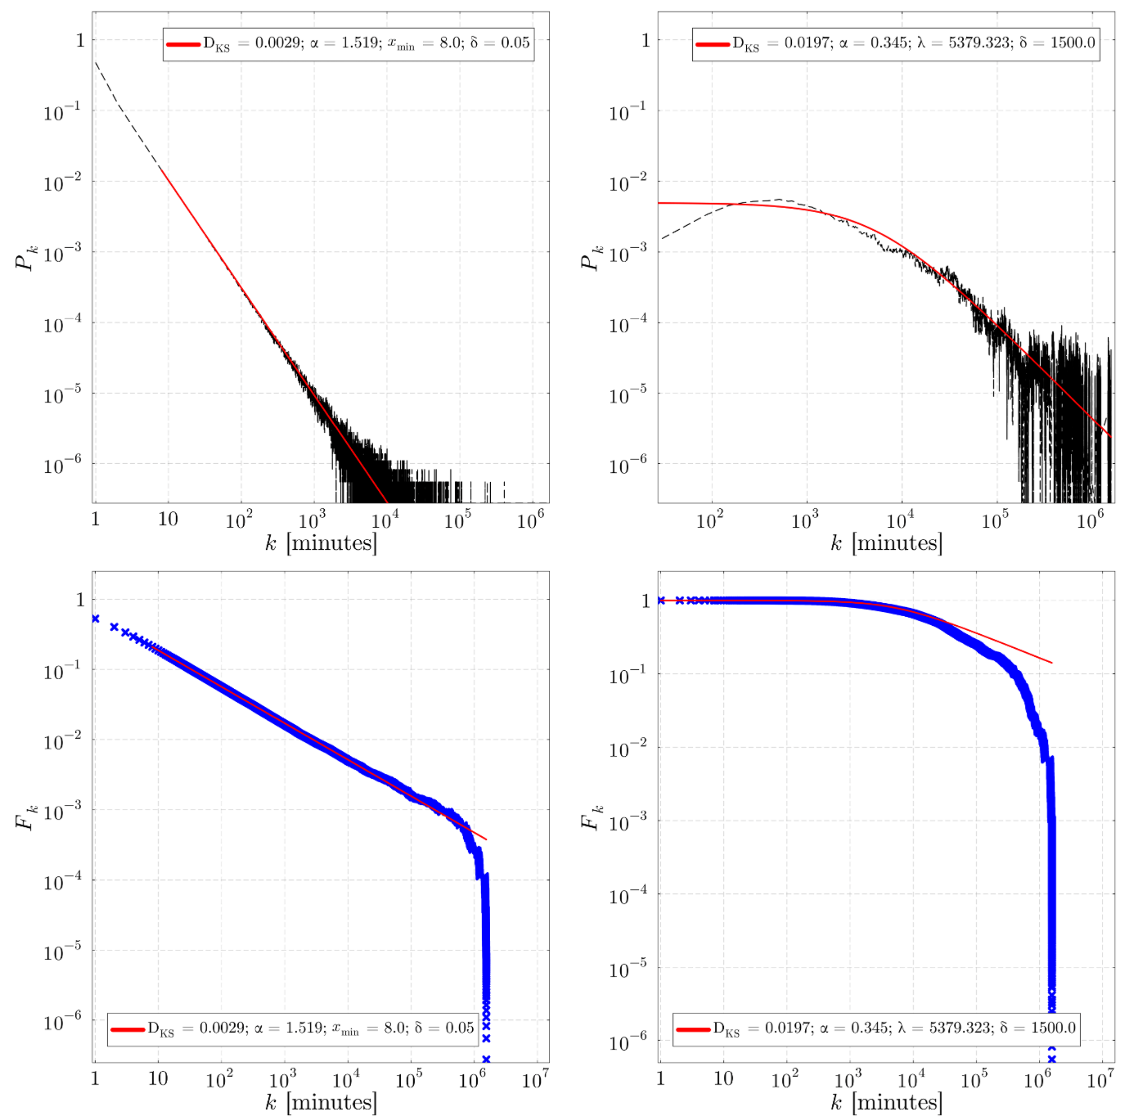

Figure 6. DWTs for BTC-USDT. The upper panels show the PDF results corresponding to $\delta = 0.05$ (left side) and $\delta = 1500$ (right side), the black dashed line corresponding to the empirical data, while the red full line shows the distribution fit. The lower panels show the same DWTs but in the CCDF representation. The Pareto–Tsallis fits (right side) computed for raw exchange data are to be contrasted with those in figure 5 where the preprocessing presented in [10] was applied. The comparison shows the importance of detrending in the case of serialized time series data with an underlying long-term monotonic trend.

noise—$dD_{KL}/d\delta$ exhibits a clear jump. The observed behavior in ΔD_{KS} and $dD_{KL}/d\delta$ suggests that the transition from a scale-free to a Pareto–Tsallis distribution is a structural one and gives rise to interesting questions about the underlying physics of the systems under scrutiny.

4. Conclusions and outlook

We have investigated by computational means seven time series pertaining to terrestrial and space weather, sea-level variations, fiat and cryptocurrency exchange rates, and synthetic data for autonomous driving. Our analysis has revealed that the DWTs exhibits two distinct and seemingly generic behaviors: for small values of the δ threshold the distributions are consistently of scale-free type, while for large values of the δ threshold they accurately reproduce a Pareto–Tsallis distribution.

We have developed and successfully applied a reliable, dual-indicator framework to identify the transition point between these distinct regimes. This point is marked by the confluence of two complementary indicators: the zero-crossing of the difference in Kolmogorov–Smirnov quality of fit ΔD_{KS} and the

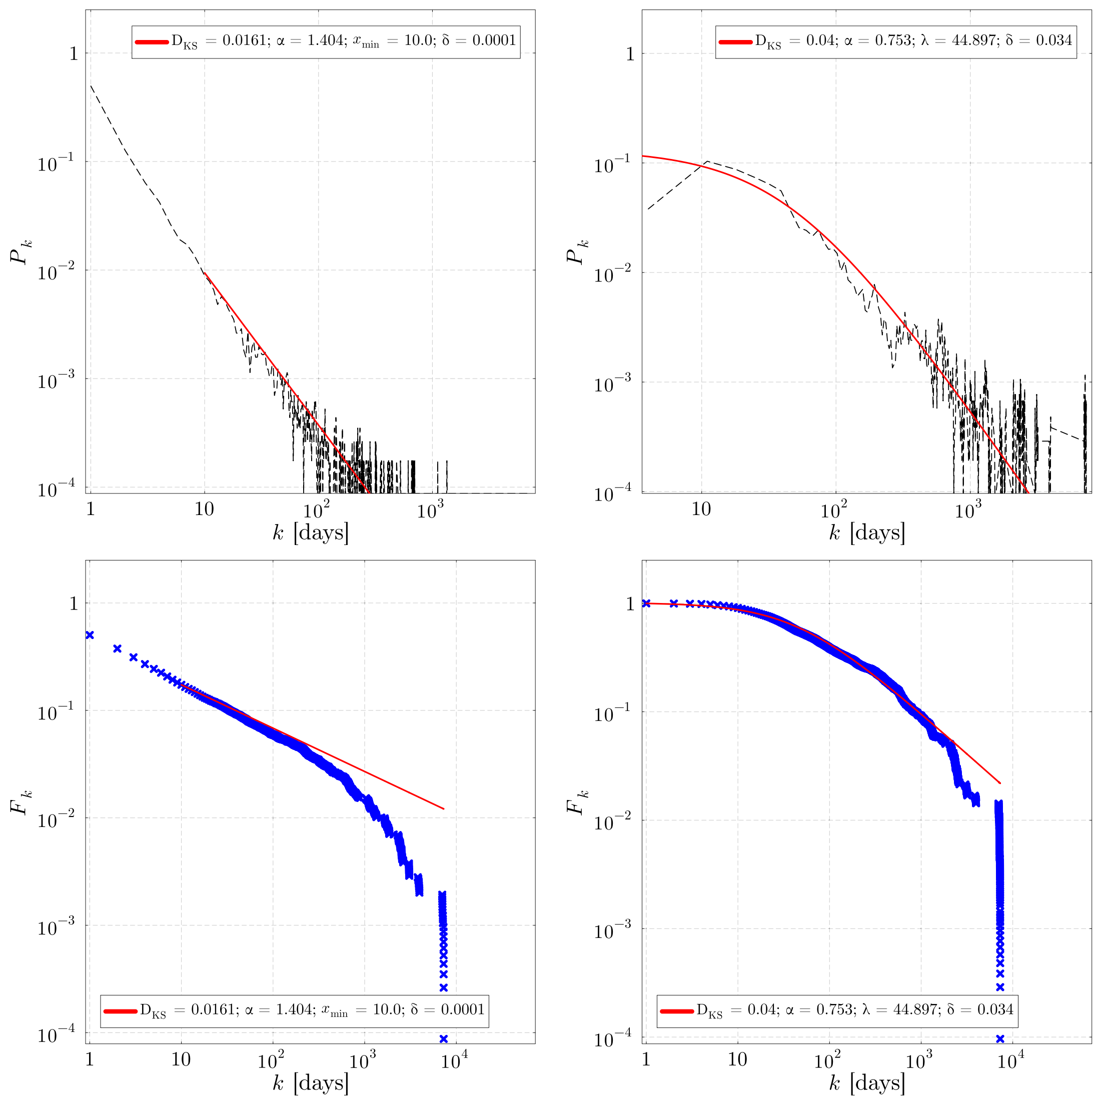

Figure 7. DWTs for EURO-USD exchange rate. The upper panels show the PDF results corresponding to $\delta = 0.0001$ (left side) and $\delta = 0.034$ (right side), the black dashed line corresponding to the empirical data, while the red full line shows the distribution fit. The lower panels show the same DWTs but in the CCDF representation.

sharp, jump-like feature in the derivative of the generalized KL divergence $\mathrm{dD_{KL}}/\mathrm{d}\delta$. This result is consistent across all datasets under scrutiny, which indicates that the δ threshold acts as an effective fingerprint of the underlying dynamics, which allows us to distinguish between different regimes of physical interest. As we have not seen a domain-specificity in the aforementioned transition from a scale-free to a Pareto–Tsallis distribution, apart from the particularities of the numerical data cleansing, we consider that one major area of future interest is represented by the development of microscopic (or potentially mesoscopic) models that are able to reproduce the regimes previously discussed together with the accompanying transition.

Concerning the practical future applications of our results, let us mention that for currency trading there is the prospect of identifying trading styles in one-agent datasets based on how the α exponents obtained for small values of δ cluster, while for datasets pertaining to the time-monitoring of emotional trackers one could define an indicator akin to the emotional volatility of the subject. Our preliminary investigations have shown that DWTs for main emotions (such as happiness, fear, anger, etc) extracted from recordings in the public domain using open-source tools share the same scale-free-like structure in the case of small values of δ. However, the accuracy of an emotional volatility indicator constructed using these DWTs is greatly impacted by how the emotional trackers are extracted, as there is no ground truth that one can compare to.

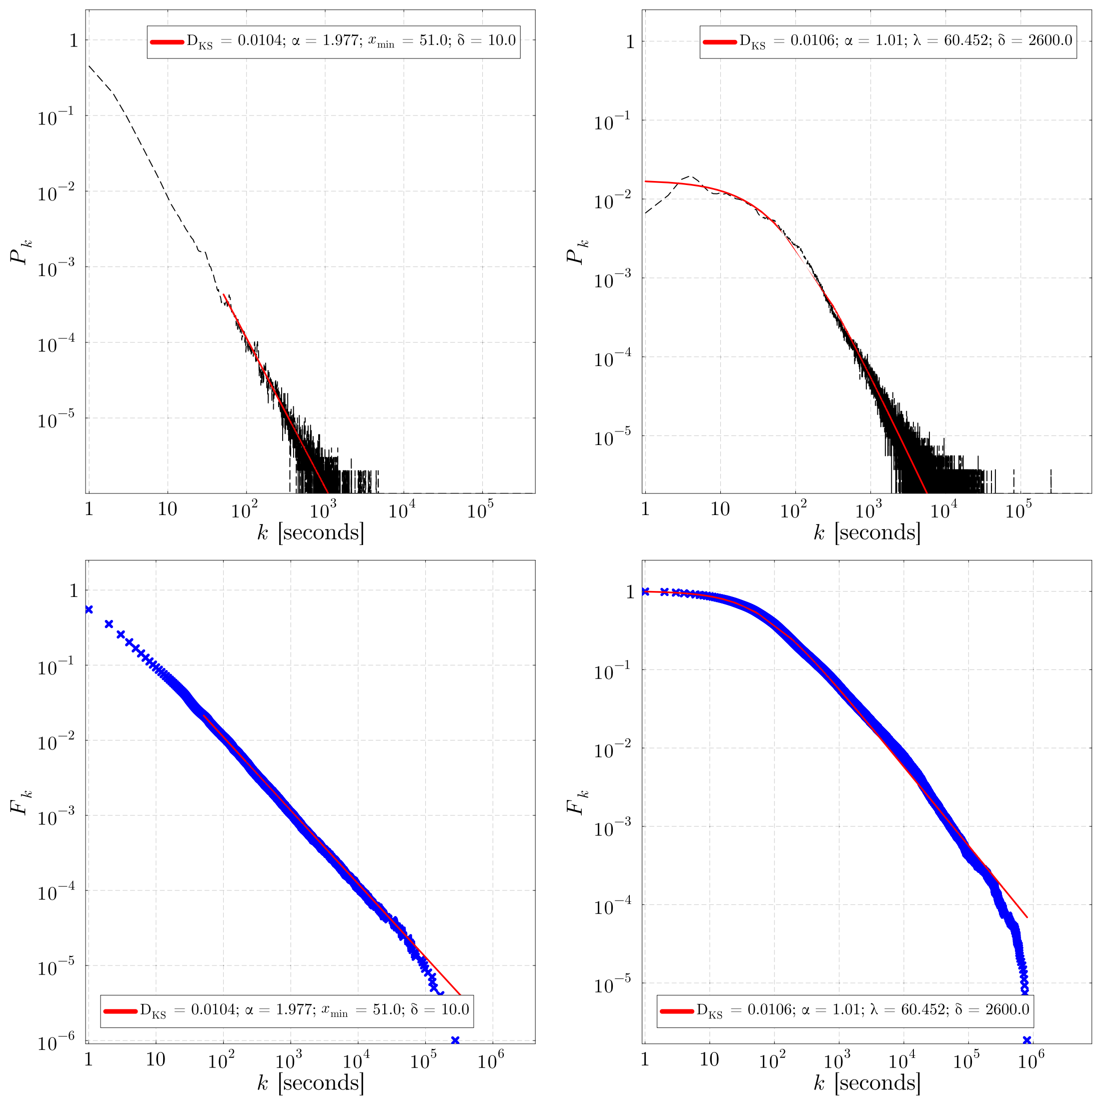

Figure 8. DWTs for instantaneous fuel consumption obtained from synthetic automotive data. The upper panels show the PDF results corresponding to $\delta = 10$ mg s^{-1} (left side) and $\delta = 2600$ mg s^{-1} (right side), the black dashed line corresponding to the empirical data, while the red full line shows the distribution fit. The lower panels show the same DWTs but in the CCDF representation.

Table 1. Exponents and goodness-of-fit D_{KS} obtained for the scale-free and Pareto–Tsallis distributions at corresponding small and large δ threshold values for all of the analyzed datasets.

	Scale-free distribution			Pareto–Tsallis distribution		
Dataset	δ	α	D_{KS}	δ	α	D_{KS}
Wind speed	1	2.100	1.08×10^{-2}	32.5	1.485	1.22×10^{-2}
Daily solar index	5×10^{-2}	2.268	3.07×10^{-2}	55	6.441	1.86×10^{-2}
TEC in ionosphere	5	2.411	3.04×10^{-2}	25	1.492	2.09×10^{-2}
Sea level variations	1	2.141	3.45×10^{-2}	76	1.389	3.49×10^{-2}
BTC—USDT	5×10^{-2}	1.519	2.88×10^{-3}	15×10^{2}	0.345	1.97×10^{-2}
EUR—USD	1×10^{-4}	1.404	1.61×10^{-2}	3.4×10^{-2}	0.753	4.00×10^{-2}
Fuel consumption	10	1.977	1.04×10^{-2}	26×10^{2}	1.010	1.06×10^{-2}

Finally, we expect that the DWTs will be useful for quality assessment and validation of synthetic data. Many of the classical statistical tools currently in use – going from distributional distance measures to the multitude of correlation coefficients – apply to the direct data, so they do not capture the similarity between long-range correlations in real and synthetic time series. Considering that the DWTs can be tuned (via the δ threshold) to capture long-range correlations across multiple scales, they can be employed as a more effective dynamical fingerprint of time series data. Consequently, DWTs have the

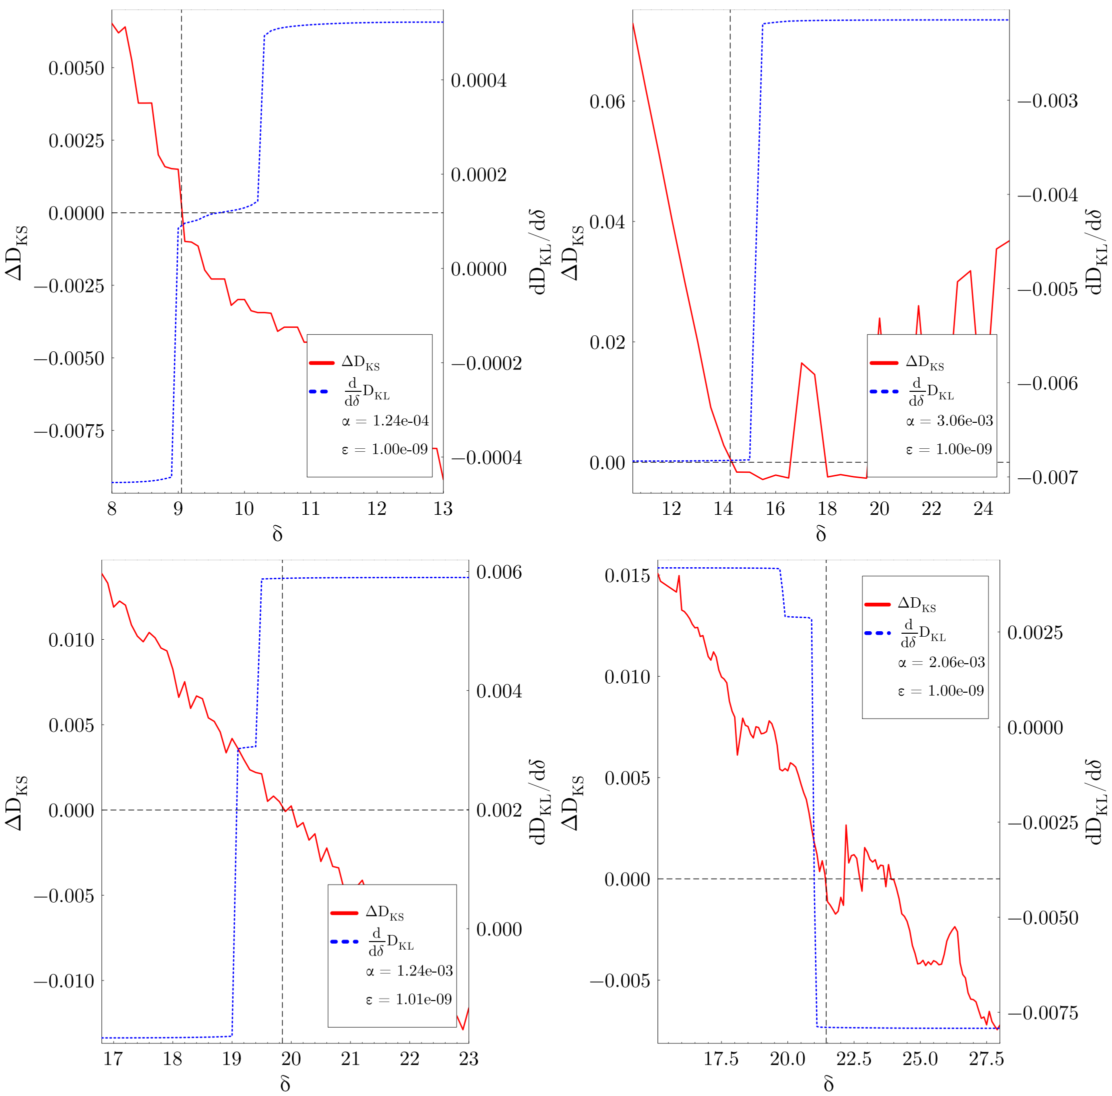

Figure 9. The scale-free to Pareto–Tsallis transition of DWTs shown for wind speeds (upper-left panel), sea-level variations (upper-right panel), the total electron density in ionosphere (lower-left panel) and daily solar index (lower-right panel) datasets. In all panels the full red line shows the difference in the Kolmogorov–Smirnov quality of fit indicators, while the dotted blue line shows the derivative of the generalized Kullback–Leibler divergence obtained using the method in section 2.3. Please note that the jump in $dD_{KL}/d\delta$ takes place in a close vicinity of the zero-crossing of ΔD_{KS}. Also note that unlike the other transitions reported here, in the case of the daily solar index the jump in $dD_{KL}/d\delta$ indicates a local decrease in the generalized Kullback–Leibler divergence after the zero-crossing of ΔD_{KS}, which is however followed by an overall monotonous increase at larger values of δ (not shown).

potential to offer a robust, multi-scale framework for the data-driven pipelines of bleeding-edge machine learning. They provide a stringent heuristic for curating massive, noisy datasets prior to the training of deep learning models, and can serve as an advanced benchmark that takes into account both short- and long-term dynamics to evaluate whether synthetic time series produced by generative AI successfully capture the heavy-tailed complexity of real-world phenomena.

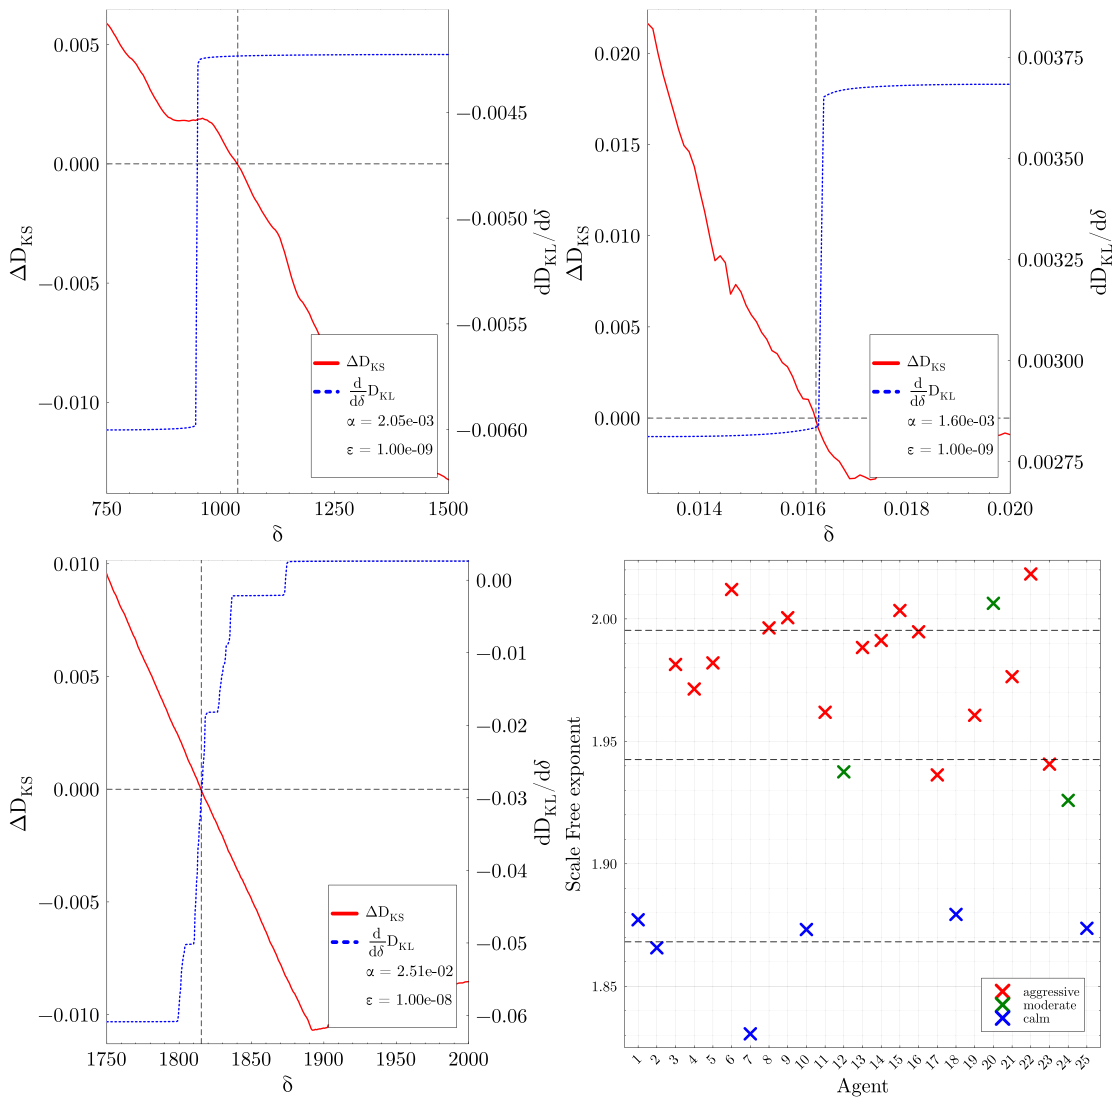

Figure 10. The scale-free to Pareto–Tsallis transition of DWTs shown for the BTC-USDT (upper-left panel), the EUR-USD exchange rate (upper-right panel) and the instantaneous fuel consumption pertaining to the automotive synthetic dataset (lower-left panel), with all technical details as in figure 9. Please note that the transition in the DWTs is clearly visible for the BTC-USDT dataset, despite working with raw (non-detrended) data, which is a sign of the robustness of this transition. In the lower-right panel, we show the categorization of 25 synthetic automotive agents into three distinct driving styles (aggressive, moderate, and calm). This classification relies on the scale-free exponent α evaluated at a small threshold ($\delta = 20$) utilizing a Fuzzy C-Means clustering algorithm, demonstrating the practical applicability of the DWT for extracting macroscopic behavioral fingerprints.

Acknowledgments

The authors acknowledge fruitful discussions with Mihai Dima on the side of financial time series and Dorin Dumitran on the side of emotional trackers. The numerical simulations reported here were performed in the computing center of the Institute of Space Science—INFLPR Subsidiary, Măgurele, Romania.

Data availability statement

All of the datasets used in this study are either publicly available or can be provided by the authors upon reasonable request. The Geisenheim wind speed data can be accessed via the Meteostat project at https://meteostat.net/en/station/10628. The daily solar activity index is available from NOAA's National Centers for Environmental Information at www.ngdc.noaa.gov/stp/space-weather/solar-data/solar-features/solar-flares/index/flare-index/. The ionospheric TEC dataset is an internal dataset; related data or derived waiting time distributions can be obtained from the authors upon request. Sea level data for the Port of Trieste has been published and is available as described in [10]. The BTC–USDT minute-by-minute data were retrieved using the CryptoDatum API from *cryptodatum.io*. Historical EUR–USD exchange rates can be obtained from financial data providers such as *investing.com*. The synthetic automotive dataset was generated as part of the TRUSTEE research project https://horizon-trustee.eu/; while not publicly posted, it can be shared by the authors for research purposes.

The data that support the findings of this study are available upon reasonable request from the authors. Part of the data is internal to the TRUSTEE Horizon European Project and is not cleared for public dissemination.

Funding

This work was supported by the Romanian Ministry of Education and Research under Romanian National Core Program LAPLAS VII—contract no. 30 N/2023. For this work ANŻ, MCR and PAG were partly supported by the TRUSTEE Horizon Project (Project Number: 10 107 0214).

Author contributions

Paul-Adrian Gogîţă 0009-0007-6716-165X
Conceptualization (lead), Data curation (lead), Formal analysis (lead), Investigation (lead), Methodology (lead), Software (lead), Visualization (lead), Writing – original draft (lead)

Tudor-Gabriel Dumitru 0009-0001-1970-0766
Conceptualization (equal), Formal analysis (equal), Visualization (equal)

Florin-Ioan Constantin 0009-0006-2034-6350
Data curation (equal)

Tudor-Andrei Diac 0009-0004-7940-8996
Formal analysis (equal), Visualization (equal), Writing – original draft (equal)

Alexandra-Florentina Neagoe 0009-0002-2349-208X
Writing – review & editing (equal)

Mihaela-Carina Raportaru 0000-0001-8604-1441
Conceptualization (equal), Methodology (equal), Supervision (equal), Writing – review & editing (equal)

Alexandru Nicolin-Żaczek 0000-0002-4940-4212
Conceptualization (equal), Funding acquisition (equal), Project administration (equal), Supervision (equal), Writing – review & editing (equal)

References

[1] Clarke1 P *et al* 2016 Big data in the physical sciences: challenges and opportunities ATI Scoping Report (available at: https://indico.cern.ch/event/449964/attachments/1253648/1849677/Bigdata_physicalsciences.pdf)
[2] Hassanien A E, Azar A T, Snasael V, Kacprzyk J and Abawajy J E 2015 *Big Data in Complex Systems. Challanges and Opportunities* (Springer)
[3] Kar A and Dwivedi Y 2020 Theory building with big data-driven research– moving away from the "What" towards the "Why" *Int. J. Inf. Manage.* **54** 102205
[4] Vivirschi B, Boboc P, Băran V and Nicolin A 2020 Scale-free distributions of waiting times for earthquakes *Phys. Scr.* **95** 044011
[5] Pană G T, Ivanoaica T, Raportaru C M, Băran V and Nicolin A 2021 Towards the implementation of FAIR principles on an earthquake analysis platform *2021 20th RoEduNet Conf. Proc.: Networking In Education And Research (RoEduNet)* pp 1–4
[6] Pană G T and Nicolin-Żaczek A 2023 Motifs in earthquake networks: Romania, Italy, United States of America and Japan *Physica A* **623** 129301
[7] Pană G T, Zgură S, Băran V and Nicolin A 2023 Waiting times distributions for moonquakes and marsquakes *AIP Conf. Proc.* **2843** 020004
[8] Antonopoulos C G, Michas G, Vallianatos F and Bountis T 2014 Evidence of q-exponential statistics in Greek seismicity *Physica A* **409** 71–77
[9] Bogachev M I, Kayumov A R and Bunde A 2014 universal internucleotide statistics in full genomes: a footprint of the DNA structure and packaging? *PLoS One* **9** e112534
[10] Pană G T, Gogîţă P-A and Nicolin-Żaczek A 2024 Waiting times for sea level variations in the Port of Trieste: a computational data-driven study *Rom. J. Phys.* **69** 111
[11] Baran V, Zus M, Bonasera A and Paturca A 2015 Quantifying the folding mechanism in chaotic dynamics *Rom. J. Phys.* **60** 1263–77
[12] Marwan M, Romano M C, Thiel M and Kurths J 2007 Recurrence plots for the analysis of complex systems *Phys. Rep.* **438** 237–329

[13] Yang Y-X, Gao Z-K, Wang X-M, Li Y-L, Han J-W, Marwan N and Kurths J 2018 A recurrence quantification analysis-based channel-frequency convolutional neural network for emotion recognition from EEG *Chaos* **28** 085724
[14] Pareto V 1897 Cours d'économie politique, F. Rouge Lausanne
[15] Tsallis C 2009 *Introduction to Nonextensive Statistical Mechanics: Approaching a Complex World* (Springer)
[16] Bak P 1996 *How Nature Works. The Science of Self-Organized Criticality* (Copernicus Press New York)
[17] Clauset A, Shalizi C R and Newman M E J 2009 Power law distributions in empirical data *SIAM Rev.* **51** 661–703
[18] Połoczański R, Wyłomańska A, Maciejewska M, Szczurek A and Gajda J 2017 Modified cumulative distribution function in application to waiting time analysis in the continuous random walk scenario *J. Phys. A: Math. Theor.* **50** 034002
[19] Gnosh J and Sanders B C 2016 Quantum simulation of macro and micro quantum phase transition from paramagnetism to frustrated magnetism with a superconducting circuit *New J. Phys.* **18** 033015
[20] Aschwanden M J, Johnson J R and Nurhan Y I 2021 The poissonian origin of power laws in solar flare waiting time distributions *Astrophys. J.* **921** 166
[21] Constantin F I, Popescu E M, Chiriţoi G, Popescu F A, Nicolin-Żaczek A, Năstase E I and Muntean A 2024 Local monitoring of TEC using GNSS services over Romania *Rom. J. Phys.* **69** 805
[22] Bezanson J, Edelman A, Karpinski S and Shah V B 2017 Julia: a fresh approach to numerical computing *SIAM Rev.* **59** 65–98
[23] Mogensen P K and Riseth A N 2018 Optim: a mathemathical optimization package for Julia *J. Open Source Softw.* **3** 615
[24] Nelder J A and Mead R 1965 A simplex method for function minimization *Comput. J.* **7** 308–13
[25] Lagarias J C, Reeds J A, Wright M H and Wright P E 1998 Convergence Properties of the Nelder-Mead simplex method in low dimensions *SIAM J. Optim.* **9** 112–47
[26] Gao F and Han L 2012 Implementing the Nelder-Mead simplex algorithm with adaptive parameters *Comput. Optim. Appl.* **51** 259–77
[27] Chartrand R 2011 Numerical differentiation of noisy, nonsmooth data *ISRN Appl. Math.* **2011** 164564
[28] Olami Z, Federm H J S and Christensen K 1992 Self-organized criticality in a continuous, nonconservative cellular automaton modeling earthquakes *Phys. Rev. Lett.* **68** 1244
[29] Stanley H E, Amaral L A N, Buldyrev S V, Gopikrishnan P, Plerou V and Salinger M A 2002 Self-organized complexity in economics and finance *Proc. Natl Acad. Sci.* **99** 2561–5
[30] Bartolozzi M, Leinweber D B and Thomas A W 2005 Self-organized criticality and stock market dynamics: an empirical study *Physica A* **350** 451–65
[31] Sitnov M I, Sharma A S, Papadopoulos K and Vassiliadis D 2001 Modeling substorm dynamics of the magnetosphere: from self-organization and self-organized criticality to nonequilibrium phase transitions *Phys. Rev. E* **65** 016116
[32] Paredes C, Godoy C and Castedo R 2015 Soft-cliff retreat, self-organized critical phenomena in the limit of predictability? *Fractals* **23** 1540009
[33] Fagherazzi S 2008 Self-organization of tidal deltas *Proc. Natl Acad. Sci.* **105** 18692–5
[34] Peters O and Neelin J D 2006 Critical phenomena in atmospheric precipitation *Nat. Phys.* **2** 393–6
[35] Peters O and Christensen K 2002 Rain: relaxations in the sky *Phys. Rev. E* **66** 036120