

# A Physics – Based Model of Returns: Anomalous Scaling and Irreversibility

## Abstract

Every empirical study of financial returns since Mandelbrot (1963) has converged on the same short list of stylized facts: heavy polynomial tails, anomalous scaling of volatility with the aggregation horizon, power-law persistence of absolute returns, multiscaling across moment orders, and vanishing linear autocorrelation. These regularities are robust across asset classes, geographies, and decades, yet no single textbook model (Black–Scholes, GARCH, jump-diffusion) reproduces them simultaneously. Each captures one or two features and patches the others, producing an inelegant landscape of model-specific extensions.

We revisit a framework from statistical physics, the Baldovin–Stella (BS) model, originally developed to describe anomalous diffusion in out-of-equilibrium systems. Built on a single structural assumption about how memory propagates through returns, the BS model derives the full set of stylized facts as mathematical consequences rather than calibration targets. Its density ansatz is formally identical to Widom’s scaling hypothesis for critical phenomena; its memory kernel is a fractional integration operator, of the same family underpinning rough volatility and fractional Brownian motion; its time-inhomogeneity encodes the irreversibility of financial evolution in a mathematically precise way.

The article has two aims. The first is to present an elegant, physics-anchored framework for building financial models that account for all the relevant regularities at once, rather than any strict subset. The second is evidential: to argue that the language of statistical mechanics encompassing scaling, universality and irreversibility is not a metaphor but a working toolkit whose financial applications are only beginning to be charted.

## Stylized Facts and State of the Art

Each statistical model that aims at reproducing the behaviour and the evolution of financial returns must be calibrated such that the output is consistent with the stylized facts, meaning empirical properties which characterize the distribution of these returns. Before introducing our model, we will briefly summarize these well-known regularities and describe how the latter are expressed and treated in the most famous financial models.

- (SF1) **Fat Tails** are the expression of extreme returns happening far more often than a Gaussian assumption would allow.
- (SF2) **Anomalous scaling of the volatility**: empirically, we observe that it deviates of the “square root law” implied by a Standard Brownian Motion
- (SF3) **Volatility clustering** with long memory: the autocorrelation function of returns decays as a power law not exponentially, translating in persistent fingerprints following shocks.
- (SF4) **Multiscaling**: the distribution of returns actually changes shape as aggregating across horizons, not just rescale.
- (SF5) **Approximately zero linear autocorrelation of returns themselves**: today’s returns bear no information about tomorrow market’s behaviour.

---

*All the views expressed are opinions of Bocconi Students Investment Club members and can in no way be associated with Bocconi University. All the financial recommendations offered are for educational purposes only. Bocconi Students Investment Club declines any responsibility for eventual losses you may incur implementing all or part of the ideas contained in this website. The Bocconi Students Investment Club is not authorised to give investment advice. Information, opinions, and estimates contained in this report reflect a judgment at its original date of publication by Bocconi Students Investment Club and are subject to change without notice. The price, value of and income from any of the securities or financial instruments mentioned in this report can fall as well as rise. Bocconi Students Investment Club does not receive compensation and has no business relationship with any mentioned company.*

Given the generality of these properties, they must constitute the cornerstones of any analytical model for simulating or replicating financial scenarios of this kind, yet standard models usually fall short in incorporating all of these features simultaneously. For example, the renowned Black-Scholes model assumes a Gaussian distribution and cannot exhibit fat tails and clustering; GARCH family models capture volatility clustering but with an exponential decay and does not natively produce power-law memory or multiscaling and, finally, jump diffusions patch fat tails but again fails to model clustering and multiscaling.

In the following sections we will discuss the mathematical foundation of the model and the interpretation of its parameters and then proceed to calibrate the model on a sample containing 130 years of DJIA returns.

### Inhomogeneous scaling of the density function

Let's define  $r(t, T)$  as the log-return at time  $t$ , aggregated on a scale  $T$  and let  $p_T(r)$  be the pdf of these returns on the same scale. We can now reformulate (SF1-5) in a more formal way:

$$(SF1) \quad p_T(r) \sim C|r(t, T)|^{-(1+\alpha)} \text{ as } |r| \rightarrow \infty \text{ and } \alpha > 0$$

$$(SF2) \quad \sigma(T) = \sqrt{\mathbb{E}[r(t, T)^2]} \sim T^D \text{ with } D \neq \frac{1}{2}$$

$$(SF3) \quad \text{Corr}(|r_t|, |r_{t+\tau}|) \sim \tau^{-\beta} \text{ where } 0 < \beta < 1$$

$$(SF4) \quad \mathbb{E}[|r(t, T)|^q] \sim T^{qD(q)} \text{ with a non-constant } D(q)$$

$$(SF5) \quad \text{Corr}(r_t, r_{t+\tau}) \approx 0 \text{ for } \tau > 0$$

It is worth noting that a standard Gaussian model satisfies only (SF2) with  $D = 1/2$  and (SF5), whereas a GARCH model would be able to incorporate also (SF3), albeit with an exponential decay.

In order to create a model which is simultaneously able to meet these requirements and to provide enough degrees of freedom, we make a single strong assumption on the structure of the pdf, which is set to be described by the following expression:

$$p_T(r) = \frac{1}{\sigma(T)} F\left(\frac{r}{\sigma(T)}\right)$$

Where:

$$\sigma(T) = T^D g(t)$$

It is now crucial to understand the twofold assumption encompassed in this formalization and to explicit its financial interpretation. Basically, we can disentangle two separate sub-assumptions:

- (AS1) **Self-similarity:** the pdf on a  $T$  horizon is a scaled version of a universal function  $F$ . Pragmatically, this means that plotting  $\sigma(T)p_T(r)$  against  $r/\sigma(T)$  for different values of  $T$  the curves collapse on the single  $F$  curve (see Fig.1 below).

---

*All the views expressed are opinions of Bocconi Students Investment Club members and can in no way be associated with Bocconi University. All the financial recommendations offered are for educational purposes only. Bocconi Students Investment Club declines any responsibility for eventual losses you may incur implementing all or part of the ideas contained in this website. The Bocconi Students Investment Club is not authorised to give investment advice. Information, opinions, and estimates contained in this report reflect a judgment at its original date of publication by Bocconi Students Investment Club and are subject to change without notice. The price, value of and income from any of the securities or financial instruments mentioned in this report can fall as well as rise. Bocconi Students Investment Club does not receive compensation and has no business relationship with any mentioned company.*

(AS2) **Anomalous scaling:** the scaling factor is inversely proportional to  $\sigma(T)$ , which is, on turn, factored as a power-law and a time-scaling factor  $g(t)$ .

As it is written, the process is self-affine but not stationary. The introduction of the time scaling factor in the expression for the volatility is crucial, making the model time inhomogeneous. In other words, the explicit presence of the variable  $t$  breaks homogeneity and code the irreversibility of the process. The mathematical expression of this function is particularly complex and irrelevant for the purpose of this article, but it models the fact that the markets “ages”, meaning that response change depending on how much time has passed from the initial shock.

### The Universal Shape Function: Truncated Levy

We now face the challenge of choosing  $F$  such that our (SF1-5) can be satisfied. Of course, the standard Gaussian would fail on fat tails and multiscaling. We opt for a truncated Levy curve: a curve that behaves like a Levy-stable distribution in the middle tails (exhibiting power-law decay), but transitions to Gaussian exponential decay in the extreme tails, ensuring finite variance. We can talk about a “Gaussian-Lévy crossover”, modelled by idiosyncratic parameters. This solution is the only one which can simultaneously satisfy the tail, the finite variance and the self-similar scaling requirements.

Unfortunately, there exist no simple closed expression for the aforementioned curve, which is indeed often expressed as a characteristic function in the Fourier space, whose specific form is given in the Appendix.

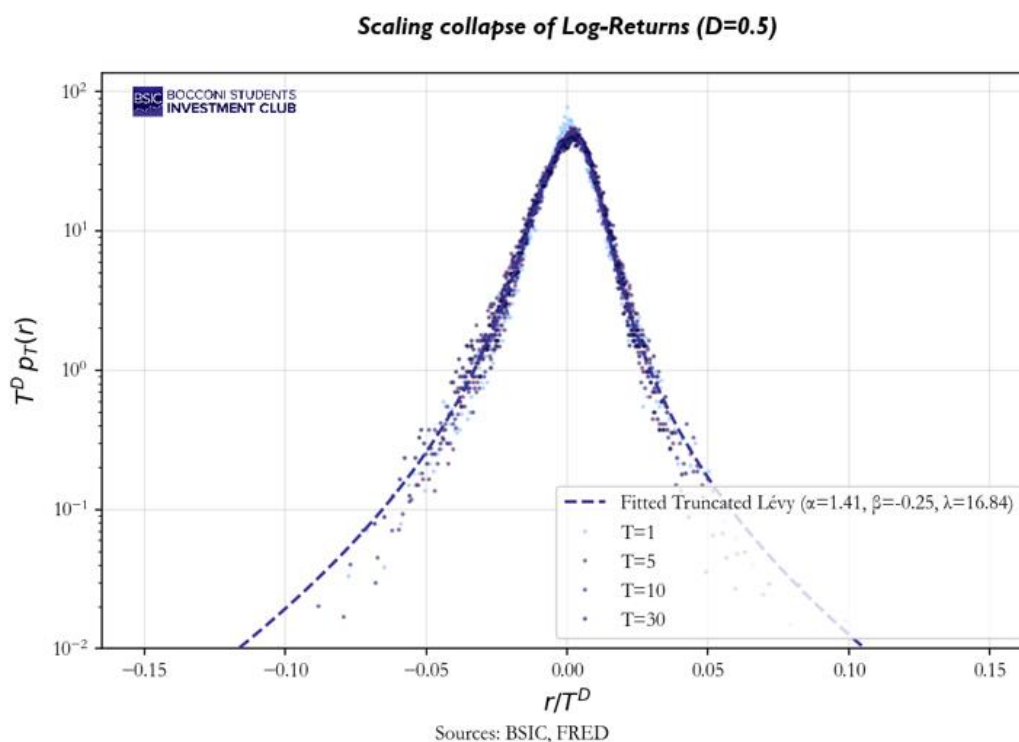


Figure 1: Scaling collapse of log-return densities at different horizons. Plotting  $T^D p^T(r)$  vs  $r/T^D$ , all curves fall onto a single universal shape  $F$ , fitted here by a truncated Lévy distribution. Data: DJIA 1897–2025.

All the views expressed are opinions of Bocconi Students Investment Club members and can in no way be associated with Bocconi University. All the financial recommendations offered are for educational purposes only. Bocconi Students Investment Club declines any responsibility for eventual losses you may incur implementing all or part of the ideas contained in this website. The Bocconi Students Investment Club is not authorised to give investment advice. Information, opinions, and estimates contained in this report reflect a judgment at its original date of publication by Bocconi Students Investment Club and are subject to change without notice. The price, value of and income from any of the securities or financial instruments mentioned in this report can fall as well as rise. Bocconi Students Investment Club does not receive compensation and has no business relationship with any mentioned company.

## The core idea

In this section our objective is to provide an intuition of the mechanism underlying the model, sufficient to grasp how and why it produces the desired output.

A strong assumption has been made on the marginals of the return distribution at each horizon, but these do not exhaust the description of the stochastic process. We still must specify how returns at different times depend on each other. In very general words, by accounting for the additivity of the log-returns and their non-independence (which along with the scaling ansatz would lead to a Levy-stable process, not suitable for our purpose), the actual process can be described in terms of increments which share some innovations, introducing the desired time-dependence and clustering components.

Delving deeper into this matter, today's return is computed as a weighted sum of past innovation (i.i.d truncated-Levy), whose weights decay as a power-law. It can be proved (Peirano and Challet, 2012) that, in distribution, the process is equivalent to the following linear filter, where  $\xi$  represents the i.i.d shocks and  $a_k$  the decaying weight of the k-th term (k represents the time "lag"). It is crucial to notice the transition from a family of marginals of the same distribution, to a unique stochastic process, compatible with the required properties.

$$|r(t)| = \sum_{k=0}^m a_k \xi(t - k)$$

The expression for  $a_k$  is not arbitrary but stems from requiring that the variance scales as  $T^{2D_e}$ . The explicit formula is now reported for completeness:

$$a_k = [k^{2D_e} - (k - 1)^{2D_e}]^{\frac{1}{2D_e}}$$

The key parameter is  $D_e$ , which modulates the structure of the memory of the process. We could think of it as the "texture" of the memory kernel, meaning how shocks are distributed across time. Notice how this is fundamentally different from ARMA or GARCH processes, where the exponentially decaying weights allow old shocks to be forgotten within weeks. Low  $D_e$  means a compressed memory, where only the very recent past matters strongly, though older shocks retain a non-negligible residual. Vice versa, high values of  $D_e$  mean a flat memory where all past shocks weigh similarly. This latter case converges to the limit of  $D_e = 1/2$ , which is the pure Gaussian case, in which every past event has the same significance, as well shown in the graph below.

What is now left to explore is the link between the value of this parameter and the empirical distribution. This connection represents the cornerstone of the whole calibration. Asymptotically, the ACF of the absolute returns decays as  $\tau^{-\beta}$  and one can show that  $\beta = 1 - 2D_e$  in the infinite memory, gaussian innovation limit. Practically, finite m, truncated Levy innovations and the time-reset mechanism require a numerical approach. The solution is a Monte Carlo mapping, which simulates the process for a selected array of  $D_e$  values, estimates  $\hat{\beta}$  for each of them and tabulates the resulting curve. The empirical value of  $\beta$  is then inverted through the curve to yield  $\widehat{D}_e$ .

---

*All the views expressed are opinions of Bocconi Students Investment Club members and can in no way be associated with Bocconi University. All the financial recommendations offered are for educational purposes only. Bocconi Students Investment Club declines any responsibility for eventual losses you may incur implementing all or part of the ideas contained in this website. The Bocconi Students Investment Club is not authorised to give investment advice. Information, opinions, and estimates contained in this report reflect a judgment at its original date of publication by Bocconi Students Investment Club and are subject to change without notice. The price, value of and income from any of the securities or financial instruments mentioned in this report can fall as well as rise. Bocconi Students Investment Club does not receive compensation and has no business relationship with any mentioned company.*

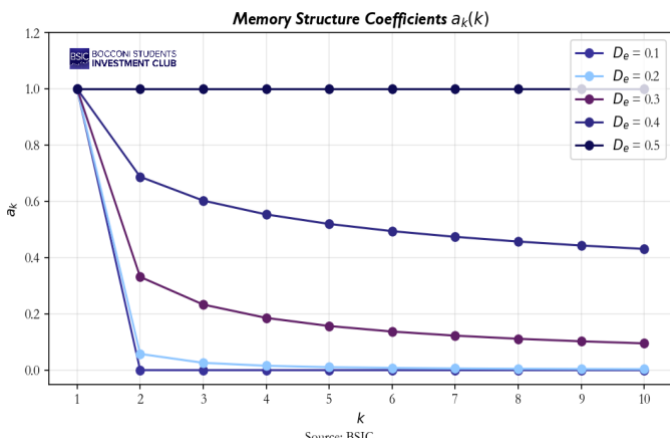


Figure 2: Memory kernel coefficients  $a_k$  for different  $D_e$ . Low  $D_e$  produces a sharply decaying kernel; high  $D_e$  produces a flat kernel. The DJIA empirical value,  $D_e \approx 0.16$ , sits in the compressed regime

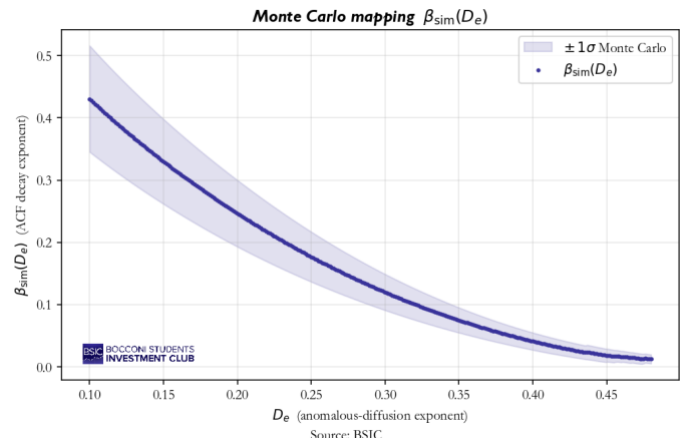


Figure 3: Monte Carlo calibration curve  $\beta_{sim}(D_e)$ . Each point is the average  $\beta$  at a given  $D_e$ , with  $\pm 1\sigma$  bands. The empirical DJIA value  $\hat{\beta} \approx 0.24$  inverts to  $\hat{D}_e \approx 0.16$ .

### In depth look at the Sawtooth Function $g(t)$ and its reset time $\tau_c$

We now return to the time-scaling function  $g(t)$ . Recall that the model specifies that the scale of log-returns at horizon  $T$  depends not only on the horizon itself but also on the absolute time of origin  $t$ . It is precisely this dependence that breaks stationarity and makes the process irreversible.

A natural question arises: if  $g(t)$  keeps accumulating memory forever, should we not expect the variance of returns to grow without bound over long samples? Empirically, the variance of the DJIA today is not systematically larger than fifty or a hundred years ago, so something must cap this accumulation. The BS model handles this through a structural reset  $\tau_c$  controlled by the reset-time parameter  $\tau_c$ . Concretely,  $g(t)$  grows between successive resets, but every  $\tau_c$  trading days the clock “restarts” and  $g(t)$  returns to its baseline value. The resulting shape is sawtooth.

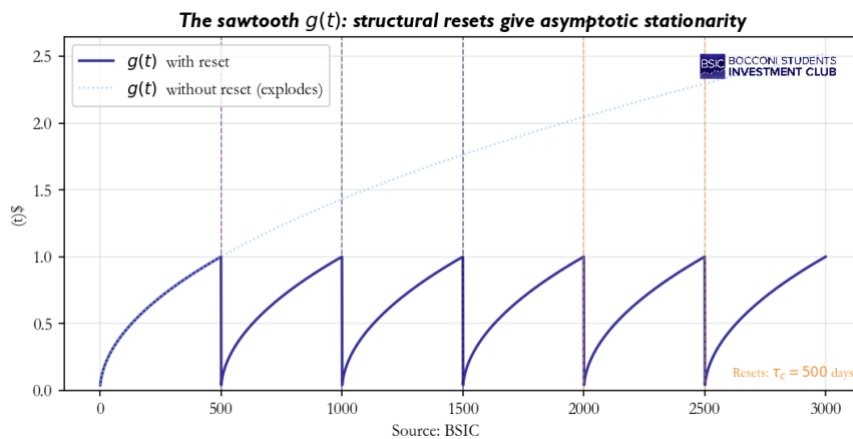


Figure 4: The sawtooth  $g(t)$  with reset time  $\tau_c = 500$  days compared with the same process without reset (light-blue dotted), which would grow unbounded.

All the views expressed are opinions of Bocconi Students Investment Club members and can in no way be associated with Bocconi University. All the financial recommendations offered are for educational purposes only. Bocconi Students Investment Club declines any responsibility for eventual losses you may incur implementing all or part of the ideas contained in this website. The Bocconi Students Investment Club is not authorised to give investment advice. Information, opinions, and estimates contained in this report reflect a judgment at its original date of publication by Bocconi Students Investment Club and are subject to change without notice. The price, value of and income from any of the securities or financial instruments mentioned in this report can fall as well as rise. Bocconi Students Investment Club does not receive compensation and has no business relationship with any mentioned company.

## Summary of the involved variables and parameters

Before delving into the empirical results, it is useful to collect all model parameters and their financial interpretation.

- 1) The parameter  $D$  measures how quickly the total volatility increases with aggregating horizon. A standard Brownian motion would lead exactly  $D = 0.5$ . For different values we talk about super- or sub- diffusion. In very simple terms, whereas normally the monthly volatility turns out being  $\sqrt{22}$  times the daily one, here it corresponds to  $22^D$  times the same benchmark.
- 2)  $D_e$  controls the memory architecture and can be named under effective diffusion exponent. It modulates the weights  $a_k$ , determining how the past is taken into account in the process. Low values tend to a “Dirac” memory, in which only today’s returns matter, whereas for  $D_e \rightarrow 0.5$  an homogeneous memory is obtained. Financially, this is the parameter most relevant for pricing path-dependant derivatives and understanding how quickly the market “forgets” a volatility regime.
- 3) The exponent  $\beta$  measures how quickly the ACF of the absolute returns decays. It is deeply interconnected with the volatility clustering, which gets more evident as the parameter decreases.
- 4) The memory is set to be finite and this directly translates into the parameter  $m$ , which serves as memory cutoff. In simpler words, it sets after how many trading days the market forgets.
- 5) Finally, the hidden parameter  $\tau_c$  modules the time scaling function  $g(t)$  by imposing a structural renewal in the process. Basically, the function  $g(t)$  “accumulates” memory for a finite period before being reset and starting over. More formally, our time scaling function represents a sawtooth function which grows between resets, whose frequency is regulated by  $\tau_c$ . If this limit did not exist, the process could not be asymptotically stationary since  $g(t)$  would explode. In financial terms, this assumption makes sure that today’s variance is not systematically higher than fifty years ago, which is empirically evident.

## From model to Stylized Facts: Empirical Validation

We calibrate the model on 130 years of DJIA daily returns (1897 – 2025, approximately 32800 trading days). For each stylized fact, we show the empirical evidence, explain which inner mechanism of the model produces it and present the estimated parameters.

### Fat tails

Fig. 7 plots the complementary CDF of standardized absolute returns on a log-log scale. Both the positive and negative tails exhibit the signature of a power law, before tapering off for extreme values. The actual exponent (Pareto exponent) lies well within the range documented in the literature (usually 2 – 5 for equity indices). On the contrary, the Gaussian benchmark clearly fails, falling away far more steeply.

In our model, these fat tails arise because the innovations are drawn from a truncated-Levy distribution, whose tails are polynomial by construction. Underlying this result is another crucial, albeit subtle, aspect. We could be

---

*All the views expressed are opinions of Bocconi Students Investment Club members and can in no way be associated with Bocconi University. All the financial recommendations offered are for educational purposes only. Bocconi Students Investment Club declines any responsibility for eventual losses you may incur implementing all or part of the ideas contained in this website. The Bocconi Students Investment Club is not authorised to give investment advice. Information, opinions, and estimates contained in this report reflect a judgment at its original date of publication by Bocconi Students Investment Club and are subject to change without notice. The price, value of and income from any of the securities or financial instruments mentioned in this report can fall as well as rise. Bocconi Students Investment Club does not receive compensation and has no business relationship with any mentioned company.*

tempted of employing a Central Limit Argument stating that our linear filter will eventually Gaussianise the shocks, which is empirically inaccurate. The pivotal point here is the slow decay of our memory weights, which makes the sum of a finite number of heavy-tailed terms remain heavy-tailed (Generalized Central Limit Theorem).

### Volatility Clustering with Power-Law Memory

Fig.6 shows the ACF of the absolute returns on a log-log scale. The empirical points align on a straight line with slope  $-\beta \approx -0.24$ , confirming power-law decay behaviour over more than two decades of lag. Interestingly enough, the GARCH(1,1) benchmark (persistence 0.97) dramatically undershoots at medium and long lags.

Our linear filter produces absolute returns which share multiple innovations, whose weights  $a_k$  decay as a power-law, hence the correlation asymptotically tends to this behaviour too. It is worth noticing that the model does not fit the clustering: it predicts it given  $D_e$ .

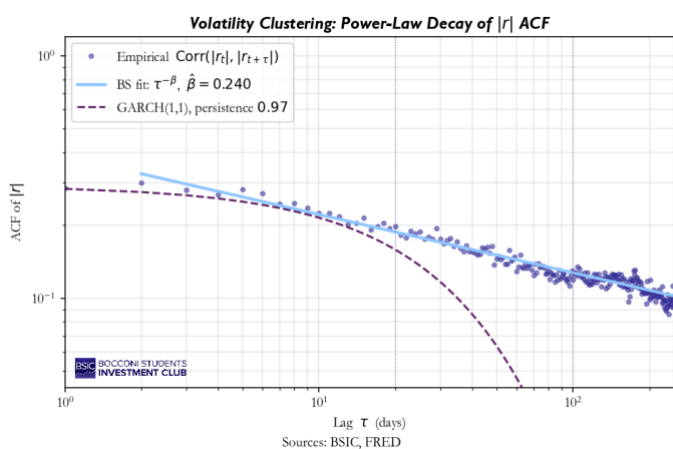


Figure 6: ACF of absolute returns in log-log scale. Empirical points follow the BS power-law fit  $\tau^{-\beta}$  with  $\hat{\beta} = 0.240$ . The GARCH(1,1) exponential benchmark (dashed) undershoots dramatically at medium-to-long lags.

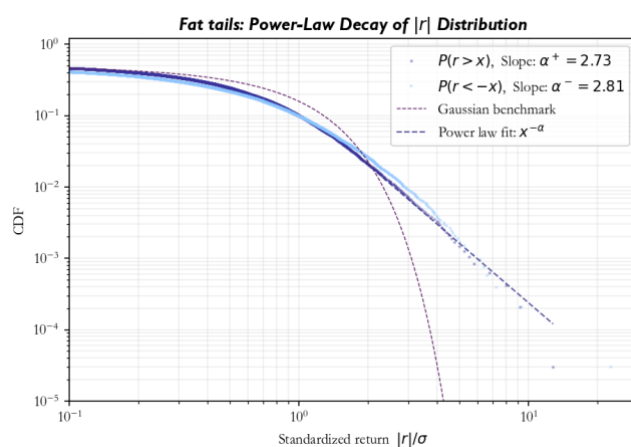


Figure 5: Complementary CDF of standardised returns. Both tails follow a power law with  $\alpha \approx 2.7-2.8$  (solid line), far heavier than the Gaussian prediction (dashed).

### Anomalous Scaling of Volatility

Figure 8 plots the empirical standard deviation of aggregated returns  $\sigma(T)$  against the aggregation horizon  $T$  on a log-log scale. The fit yields  $D = 0.516 \pm 0.004$ , statistically different from the gaussian benchmark  $D_{gauss} = 0.5$ . The inset shows the ratio  $\sigma(T)/(\sigma_1 \sqrt{T})$ , which would be constant if diffusion were normal and which shows a statistically significant super diffusion instead.

All the views expressed are opinions of Bocconi Students Investment Club members and can in no way be associated with Bocconi University. All the financial recommendations offered are for educational purposes only. Bocconi Students Investment Club declines any responsibility for eventual losses you may incur implementing all or part of the ideas contained in this website. The Bocconi Students Investment Club is not authorised to give investment advice. Information, opinions, and estimates contained in this report reflect a judgment at its original date of publication by Bocconi Students Investment Club and are subject to change without notice. The price, value of and income from any of the securities or financial instruments mentioned in this report can fall as well as rise. Bocconi Students Investment Club does not receive compensation and has no business relationship with any mentioned company.

## Multiscaling

Figure 7 shows the scaling spectrum  $\zeta(q) = qD(q)$ , empirically estimated from the regression of  $\log \mathbb{E}[|r(T)|^q]$  against  $\log(T)$  for different moment orders  $q$ . If diffusion were purely simple scaling (same exponent for all moments),  $\zeta(q)$  would be a straight line through the origin with slope  $D$ . Instead, the empirical points curve below this benchmark being strikingly well captured by the mathematical framework of the model. Basically, multiscaling arises because our innovations are characterized by non-trivial higher moments. By aggregating heavy-tailed variables with persistent weights, different moment orders “see” different parts of the tail, producing this moment-dependent scaling. This matters specially for option pricing: non-linearity means that IV across strikes will not obey simple Black-Scholes transformation, giving birth to the well know volatility smile.

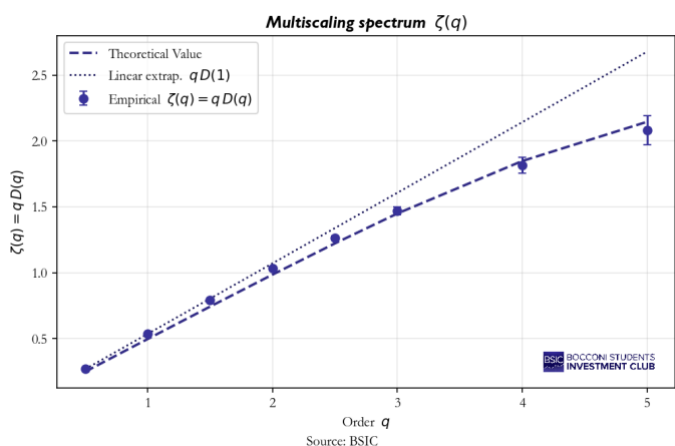


Figure 7: Multiscaling spectrum  $\zeta(q)$ . Empirical points (with bootstrap error bars) curve below the linear extrapolation (dotted), confirming non-constant  $D(q)$ . The BS-model theoretical prediction (dashed) tracks the empirical curve closely.

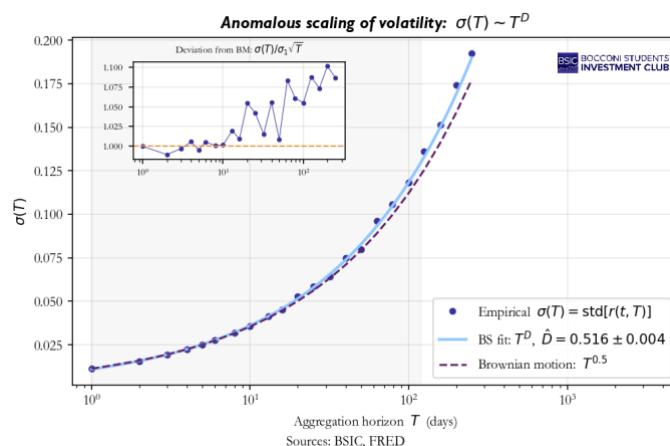


Figure 8: Anomalous volatility scaling:  $\sigma(T) \sim T^{\hat{D}}$  with  $\hat{D} = 0.516$ . The inset shows the ratio to the  $\sqrt{T}$  benchmark, systematically  $> 1$  at medium horizons. Data: DJIA 1897–2025.

## Concluding Remarks

The presented framework is valuable not only because of its mathematical elegance, but also because one compact specification produces the five pillars of empirical return behaviour simultaneously. They all emerge from the density ansatz and the very same memory kernel, parametrized by  $D_e$  and  $\beta$  alone. Furthermore, the model offers an ulterior degree of freedom, not addressed by classical models: irreversibility is indeed implemented through the sawtooth function.

Most risk and pricing infrastructures in the financial industry are built on models that explicitly assume away these features. Treating them as core and not as anomalies to patch around may significantly change the perspective about tail risk and pricing. The pivotal parameters of the model emerge as economically meaningful state variables, describing market’s memory architecture. Tracking them down through time may even reveal regime shifts, invisible to simpler classical model.

*All the views expressed are opinions of Bocconi Students Investment Club members and can in no way be associated with Bocconi University. All the financial recommendations offered are for educational purposes only. Bocconi Students Investment Club declines any responsibility for eventual losses you may incur implementing all or part of the ideas contained in this website. The Bocconi Students Investment Club is not authorised to give investment advice. Information, opinions, and estimates contained in this report reflect a judgment at its original date of publication by Bocconi Students Investment Club and are subject to change without notice. The price, value of and income from any of the securities or financial instruments mentioned in this report can fall as well as rise. Bocconi Students Investment Club does not receive compensation and has no business relationship with any mentioned company.*

The broader message of this article aims to be methodological. The language of statistical mechanics turns out to be a real working toolkit to model complex systems even in the financial world, carrying with it a precise, structured but flexible mathematical framework.

From this perspective, financial returns can be naturally interpreted as the “outcome” of a complex system. This opens the door to a deeper form of syncretism between finance and hard sciences, which would encompass the adoption of a full conceptual and technical apparatus, rather than employing just some isolated tools or concepts. In this framework, which is not at all limited to this specific paper, models’ parameters may also be linked to the nature and structural properties of the models themselves, rather than just being considered as calibrations’ outputs.

## Mathematical Appendix

Throughout the main body of this article, we have deliberately kept the mathematical details in the background, to favour the communication of the intuition and the empirical validation. Given the general complexity of the theoretical topics, several were introduced only verbally. For completeness, this Appendix aims to provide a deeper insight into these key constructions by filling them in with the minimum rigour to make each step transparent. Nonetheless, we do not attempt a full technical treatment for which one can refer to our references.

### The Truncated- Levy Characteristic Function

The universal shape function  $F$  is most compactly specified through its Fourier transform, the so-called characteristic function:

$$\widetilde{g}(k) := \langle e^{ik\xi} \rangle := \int F(\xi) e^{ik\xi} d\xi$$

In our specific example we would have:

$$\tilde{g}(k) = \exp\left(\frac{-Bk^2}{1 + C_\alpha |k|^{2-\alpha}}\right)$$

Effectively, for  $|k| \rightarrow 0$  the denominator tends to one and we get the characteristic function for a Gaussian distribution (variance would be equal to  $2B$ ). Vice versa, for extreme values of the variable  $k$  we get the characteristic function of a Levy-stable curve, whose tails hence decay polynomially:

$$\tilde{g}(k) \xrightarrow{k \rightarrow \infty} \exp\left(-\frac{B}{C_\alpha} |k|^\alpha\right)$$

### From Marginal Densities to a stochastic Process

The scaling ansatz we encountered fixes the marginal density of the return  $r(t, T)$  at every horizon. However it does not fix the joint distribution of returns at different times, which would be the same as determining the

---

*All the views expressed are opinions of Bocconi Students Investment Club members and can in no way be associated with Bocconi University. All the financial recommendations offered are for educational purposes only. Bocconi Students Investment Club declines any responsibility for eventual losses you may incur implementing all or part of the ideas contained in this website. The Bocconi Students Investment Club is not authorised to give investment advice. Information, opinions, and estimates contained in this report reflect a judgment at its original date of publication by Bocconi Students Investment Club and are subject to change without notice. The price, value of and income from any of the securities or financial instruments mentioned in this report can fall as well as rise. Bocconi Students Investment Club does not receive compensation and has no business relationship with any mentioned company.*

stochastic process itself. It is quite easy to prove that two different processes can have same marginals and yet completely different time-dependence structure.

The requirement that allows to close our construction is the additivity of log-returns, which link returns at different horizons and times:

$$r(t, T_1 + T_2) = r(t + T_1) + r(t + T_1, T_2)$$

It is now crucial to investigate two possible extreme cases. If the two summands were completely independent, their characteristic function would hence multiply and be consistent with the scaling form:

$$\widetilde{p}_{T_1+T_2}(k) = \widetilde{p}_{T_1}(k) \cdot \widetilde{p}_{T_2}(k)$$

$$\widetilde{p}_T(k) = \widetilde{g}(\sigma(T)k)$$

For these to be simultaneously satisfied would force  $\widetilde{g}$  to be stable and  $\sigma(T)$  to be a simple power, which would fail in reproducing the volatility clustering. Independence is hence incompatible with our empirical sample. In the opposite case, strongly dependent returns would grant volatility clustering but gives birth to an unconstrained problem, that does not fix a unique process.

Our solution sits in the middle of these two extremes. Log-returns are expressed as a sum (integral) of innovations which share support across different horizons. Formally:

$$r(t, T) = \sum_{s=0}^{T-1} \left[ \sum_{k=0}^m a_k \cdot \xi(t + s - k) \right]$$

It is worth noticing that this expression is manifestly consistent with additivity and allows shared innovations and hence some degree of correlation. The work of Peirano and Challet (2012) then provides the key representation result as a linear filter mentioned in the main body of the article.

### Derivation of the Memory Weights

Another interesting aspect concerns the structure of the memory weights. They are indeed uniquely determined by the requirements that the variance of aggregated return scales as  $T^{2De}$ . We now explicitly derive the formula. Recalling the linear filter representation of the returns, we can easily compute the aggregated version:

$$r(t, T) = \sum_{s=0}^{T-1} r(t + s) = \sum_{s=0}^{T-1} \sum_{k=0}^m a_k \xi(t + s - k)$$

Through some algebraic manipulation, the aggregate coefficient ends up being:

$$A_T(j) = \sum_{k=\max(0, t+T-1-j)}^{\min(m, t-j)} a_k$$

---

*All the views expressed are opinions of Bocconi Students Investment Club members and can in no way be associated with Bocconi University. All the financial recommendations offered are for educational purposes only. Bocconi Students Investment Club declines any responsibility for eventual losses you may incur implementing all or part of the ideas contained in this website. The Bocconi Students Investment Club is not authorised to give investment advice. Information, opinions, and estimates contained in this report reflect a judgment at its original date of publication by Bocconi Students Investment Club and are subject to change without notice. The price, value of and income from any of the securities or financial instruments mentioned in this report can fall as well as rise. Bocconi Students Investment Club does not receive compensation and has no business relationship with any mentioned company.*

Now we exploit the fact that these innovations are i.i.d and with fixed variance  $\sigma_{\xi}^2$ .

$$\text{Var}[R(t, T)] = \sigma_{\xi}^2 \cdot \sum_j A_T(j)^2$$

We now impose the mentioned scaling requirement:

$$\text{Var}[R(t, T)] \propto T^{2D_e} \quad \text{for every } T$$

Solving this constraint boils down to recognising that the "telescoping" implicit in  $A_T(j)$  is satisfied by any sequence of the form:

$$a_k \propto [A(k+1) - A(k)]^{1/2}$$

for some function  $A(\cdot)$ . Imposing the scaling, pins down  $A(k) = k^{2D_e}$  which yields:

$$a_k = [k^{2D_e} - (k-1)^{2D_e}]^{\frac{1}{2D_e}}$$

This is exactly the formula used in the main text. A final technical remark: the formula is valid in the asymptotic, infinite-memory limit and under Gaussian innovations. When the memory is truncated at  $m$  and the innovations follow the truncated-Lévy distribution, the relation between  $D_e$  and the empirical ACF decay exponent  $\beta$  is corrected at finite lags; this is the reason the main text inverts  $\beta$  to recover  $D_e$  via a Monte-Carlo mapping rather than the asymptotic formula  $\beta = 1 - 2D_e$ .

## References

- [1] Mandelbrot, B. (1963). The Variation of Certain Speculative Prices. *Journal of Business*, 36(4), 394–419.
- [2] Baldovin, F. & Stella, A.L. (2007). Scaling and Efficiency Determine the Irreversible Evolution of a Market. *PNAS*, 104(50), 19741–19744.
- [3] Baldovin, F. & Stella, A.L. (2007). Central Limit Theorem for Anomalous Scaling. *Physical Review E*, 75, 020101(R).
- [4] Peirano, J. & Challet, D. (2012). Baldovin–Stella Stochastic Volatility Process and Wiener Process Mixtures. *European Physical Journal B*, 85, 276.
- [5] Sokolov, I.M., Chechkin, A.V. & Klafter, J. (2004). Fractional Diffusion Equation for a Power-Law-Truncated Lévy Process. *Physica A*, 336, 245–251.

TAGS: Anomalous Scaling, Stylized Facts, Volatility Clustering, Multifractal Processes, Long-Memory Process, Statistical Mechanics, DJIA, Truncated Levy, Rough Volatility

---

*All the views expressed are opinions of Bocconi Students Investment Club members and can in no way be associated with Bocconi University. All the financial recommendations offered are for educational purposes only. Bocconi Students Investment Club declines any responsibility for eventual losses you may incur implementing all or part of the ideas contained in this website. The Bocconi Students Investment Club is not authorised to give investment advice. Information, opinions, and estimates contained in this report reflect a judgment at its original date of publication by Bocconi Students Investment Club and are subject to change without notice. The price, value of and income from any of the securities or financial instruments mentioned in this report can fall as well as rise. Bocconi Students Investment Club does not receive compensation and has no business relationship with any mentioned company.*