

# How to Avoid Bankruptcy?:

## Monte Carlo Simulation of Three Financial Markets, using the Multifractal Model of Asset Returns

Author: Rostislav SIBIRTSEV (41241)      Tutor: Michael HALLING

26,351 words

### **Abstract**

This paper has been an effort to apply fractal mathematics to understanding the general behaviour of financial markets. Fractals are special shapes that look similar at various scales. The specific model used is called the Multifractal Model of Asset Returns (MMAR) — the first ever model used for multifractal financial analysis. I apply the MMAR-model to analyse approximately 30 years of data for three distinct financial markets: (1) the US dollar/Norwegian krona (USD/NOK) currency market; (2) the Swedish stock market index — the OMXS30; and (3) the 12-month LIBOR-rate in British pounds.

First, I outline the basic concepts behind the application of fractals to finance. The point is that (multi-)fractal models should be able to capture three stylized facts about financial markets: high-kurtosis in price changes, non-independent price movements and the tendency for highly volatile days to come in clusters. The advantage of the MMAR is that the same parameters can generate simulations that fit different timescales (months, years etc.), unlike the GARCH-family or random-walk models.

Next, I applied the model to three markets and ran 10,000 Monte Carlo simulations on each market to see what the model predicts. I compared these to simulations from an ordinary random-walk model.

I found that all three markets appear to exhibit multifractal characteristics. To the naked eye, the simulations look quite realistic and exhibit the behaviour outlined above. However, to my knowledge there is no statistical method to check the model's accuracy (although here we run into a philosophical issue — the characteristics of a "realistic simulation" are not well-defined). Some simulations pass the Kolmogorov-Smirnov test for equal distributions, but the number is fewer than half. The model also appears to underestimate (or "under-simulate") kurtosis if the kurtosis of the original data was very high. Nevertheless, it is still capable of generating simulations with a wide range of kurtosis values, and these are often consistent with the data. A simple simulation suggested that, under high-kurtosis, high leverage could lead to very rapid bankruptcy, even in one day.

Finally, I discuss some implications of the MMAR-model, including some methodological, practical and philosophical considerations. I conclude that the model certainly looks better than the simple random-walk, and can be used for stress-testing portfolios. Despite this, I am unsure of its general utility for risk management, mainly because the model takes a long time to calculate, because it requires a lot of human judgement, and because its predictions are unclear and its interpretability uncertain. Also, it seems that the model can "break" after a high-kurtosis event actually occurs.

## Table of Contents

<b>How to Avoid Bankruptcy?: Multifractal Simulation of Three Financial Markets, using the Multifractal Model of Asset Returns .....</b>	<b>1</b>
<b>I. Introduction .....</b>	<b>4</b>
1.1. Motivation for study .....	5
1.2. What is a “fractal”? .....	7
1.3. Multifractal finance vs. Standard finance .....	8
1.4. Three markets — stocks, bonds and currencies — Description of data and data sources	9
<b>II. Theory .....</b>	<b>11</b>
2.1. Large price swings — Fat tails, kurtosis and Stable distributions .....	11
2.2. Dependence and clustering — the H-coefficient.....	17
2.3. Multifractals and “trading time” .....	19
2.4. Summary of theory.....	21
<b>III. Literature Review .....</b>	<b>22</b>
3.1. Mainstream finance .....	22
3.2. Mandelbrot .....	24
3.3. Other studies .....	26
3.4. Summary of Literature Review .....	28
<b>IV. Methodology .....</b>	<b>30</b>
4.1. 20+ Steps for Generating a Multifractal Simulation.....	30
4.2. Part 1: Estimating the multifractal parameters .....	31
4.3. Part 2: Constructing an MMAR simulation .....	40
4.4. Summary of Methodology .....	47
<b>V. Results.....</b>	<b>48</b>
5.1. Empirical measurements.....	48
5.2. Results from Norwegian currency market.....	51
5.3. Results from Swedish stock markets .....	52
5.4. Results from British LIBOR market.....	54
5.5. Analysis of simulated moments (final price, mean, standard deviation and kurtosis)	56
5.6. Are the simulations realistic? Kolmogorov-Smirnov testing.....	60

5.7.	Results from a basic random walk model .....	62
5.8.	Monsters — the simulations with the highest kurtosis .....	63
5.9.	Using the MMAR, what can we say about how to avoid bankruptcy? — A very basic model of payoffs from trading .....	65
5.10.	Summary of results .....	70
VI.	Analysis — Limitations and motives for future study .....	71
6.1.	Methodological considerations.....	71
6.2.	Practical considerations.....	73
6.3.	Philosophical considerations .....	75
6.4.	So is any of this useful? .....	76
6.5.	Summary of limitations .....	79
VII.	Conclusions .....	81
7.1.	Summary — What can fractals do and not do? .....	83
	Glossary .....	86
	Bibliography .....	90
	Appendix .....	97

# I. Introduction

Fractals are a concept from mathematics, which can be used in a non-mainstream way to model financial markets. The aim of this paper is to study the first ever (multi-)fractal model of financial markets — the Multifractal Model of Asset Returns (the “MMAR”). This model was developed by Benoît Mandelbrot, a Polish-born, French mathematician and a pioneer of fractal mathematics, together with two of his graduate students in the late 1990’s. I use fractals to model prices for three markets — one in equities, one in interest rates, and one in foreign currency exchange. I study the outputs of the model and show some basic statistics on what it predicts.

The **research questions** of this thesis can be summarized as follows:

- What is the MMAR-model? What does it show, and how does it work?
- What are the empirical MMAR-parameters for our three very different financial markets?
- What do MMAR simulations look like?
- How do they compare visually and statistically to the real data?
- What’s the difference between the MMAR and standard financial models, such as GARCH or a basic random walk?
- Does the model have any important limitations? What are they?
- What is the model useful for? And can it help avoid bankruptcy?

The paper is structured in seven parts as follows: [Part 1](#) introduces the idea of fractals and why they may be useful for financial modelling; [Part 2](#) outlines the theory behind the MMAR-model, including the specific features of financial markets that the model is meant to capture; [Part 3](#) provides a short literature review on the main works in fractal financial modelling; [Part 4](#) outlines the methodology that I use for this paper, including a beginner’s guide to performing MMAR simulations; [Part 5](#) displays the results of my 10,000 simulations, along with some comparisons to a simple random-walk model and some thoughts on the interpretation of results; [Part 6](#) turns a critical eye to the model and my results, where I criticise the model’s methodology, practicality and philosophical underpinnings; finally, [Part 7](#) provides the conclusion.

There are also *summary sub-sections* at the end of most of the seven parts, useful for fast reading and as a reminder of the topics that were covered. A [glossary](#) is also included at the end to help with unfamiliar concepts and terminology. The [appendix](#) contains relevant mathematical derivations and proofs, along with some further interesting insights.

## 1.1. Motivation for study

*"All I want to know is where I'm going to die, so I'll never go there."*

Charlie Munger, vice chairman of Berkshire Hathaway, as quoted in *Poor Charlie's Almanack* (Munger 2005)

Since the large-scale deregulation of the finance industry in the 1980's, the financial sector has grown rapidly, and consequently the canonical models of financial theory are being taught and used with ever-increasing frequency. And yet experience has shown that these models are incomplete — events keep occurring which, in their framework, should be impossible.<sup>1</sup> Critics of modern financial theory often cite the same limitations: independent and identically distributed events; Normal distributions; time-static volatility; perfect information; perfectly rational agents; the impossibility of booms and busts; and asset pricing models with definite values rather than ranges. Because of this, one critic — the mathematician Benoît Mandelbrot — has argued loudly that these models underestimate the risks involved in financial markets.

Today, the flaws of these assumptions are widely acknowledged by both practitioners and academic intellectuals alike, but they remain in very common use nonetheless. Their simplicity and widespread acceptance makes them convenient for making financial decisions. But with the role of finance now growing evermore predominant in the economy, it is perhaps more important than ever to develop a realistic understanding of how the dynamics of financial markets really work. Models have to be simplified, yes; after all, maps scaled one-to-one are useless. But is a map of the sea useful if it makes you believe that whirlpools don't exist?

This is not a mere academic quibble. For one thing, public confidence in financial firms has fallen to new lows. Gallup polls show that Americans' confidence in their banks has been

---

<sup>1</sup> Some examples include: (1) the 1987 Black Monday crash of world stock markets, including a 22.61% drop in the U.S., the odds of which should be less than one in  $10^{50}$  (according to Mandelbrot 2004); (2) the 1992 burst of the Japanese asset bubble, when the Nikkei stock market was at an extraordinarily high level which it hasn't attained since (see Shiller et al 1992); (3) the Russian financial crisis of 1998, which played a heavy role in the collapse of Long Term Capital Management; (4) the 2000 "dot-com" bubble, in which many Internet stocks collapsed completely; (5) the Global Financial Crisis of 2008, caused by a meltdown in the U.S. subprime mortgage market and leading to the worst economic crisis since the Great Depression of the 1930's; and (6) the 2010 trillion-dollar "flash crash" of U.S. stock markets, which lasted about 36 minutes before rebounding.

falling steadily since 1979, and has barely recovered from the 2008 crisis.<sup>2</sup> A 2017 study by the Cato Institute found that 48% of Americans have “hardly any confidence” in Wall Street and that 64% think bankers “get paid huge amounts of money” for “essentially tricking people”<sup>3</sup>. But a resurgence of public disdain for bankers is not the only issue — in some sense, the financial sector may be performing worse than before. John Kay (2015) is an economist whose book offers a very critical look at the modern financial industry and its tools, writing: “The risk models that were employed were essentially irrelevant to understanding the impact of extreme events (the situation, of course, for which risk models ought to be designed).” Kay also argues that the world has become riskier than before the deregulations. He offers a graph, based on calculations from Reinhardt & Rogoff (2010), that suggests that the incidence of banking crises has been increasing for all countries:




Figure 1 — *Are financial crises going up?* The economist John Kay (2015) argues that the 1980's deregulation of banks (red) has increased the incidence of banking crises as compared to after the Glass-Steagal legislation of 1933 (blue), and that modern financial theory was partly responsible for this increase.

Are the standard models really that bad? What is the essence of these critiques, and what are the alternatives? In the spirit of scientific exploration, this paper will look at a very different set of theories from what may be called “standard finance” — specifically, it will examine those theories pertaining to the mathematical science of *fractals*. Fractal finance is the framework that Mandelbrot passionately advocated, and in fact, it is the one he himself developed. One might say that we are trying to answer two questions — what was Mandelbrot trying to say, and did he have a point?

<sup>2</sup> See [appendix 5](#) for a graph of confidence over time.

<sup>3</sup> Source: <https://www.cato.org/survey-reports/wall-street-vs-regulators-public-attitudes-banks-financial-regulation-consumer>

## 1.2. What is a “fractal”?

A “fractal” (from the Latin word *fractus*, meaning “broken” or “fragmented”) is an object whose main mathematical property is that it is *self-similar* at different scales — fractals look similar whether you zoom in or zoom out. The term was coined by the mathematician Benoît B. Mandelbrot in 1975. Much of this thesis will focus on his work. Basically, a fractal is a specific type of shape — namely, one with “roughness” in its edges. Roughness means that the edges are not neatly straight or curved; instead, they exhibit a complex pattern, which can repeat at increasingly small scales.

Many people have an intuitive understanding of fractals — a common fractal shape in nature is a cauliflower: a branch of cauliflower looks very similar to a miniature version of the larger cauliflower tree.



Figure 2 — *Fractals in nature*. The cauliflower is: a typical example of a fractal shape.

Fractal geometry is found throughout nature, and is often much better at describing natural phenomena. In 1982, Mandelbrot released a seminal book — *The Fractal Geometry of Nature* — whose introduction famously declared: “Clouds are not spheres, mountains are not cones, coastlines are not circles, and bark is not smooth, nor does lightning travel in a straight line.”

In mathematics, a fractal is an abstract object. It can be a curve (such as the Weierstrass function) or a geometrical figure (such as the Sierpinski Gasket). When working with mathematical fractals, one is literally gazing at infinity. This is because patterns in mathematical fractals are *never-ending* — where a natural object is limited in its fractal scaling (the trees of the cauliflower can’t become as small as atoms), artificial fractals can exhibit the same shape even as zoom approaches infinity. Because of this, similar statistical tools can be used to analyse fractal shapes even at different zoom levels. Artificial fractals can thus be useful as models for reality when describing and simulating naturally occurring objects. Today, fractal mathematics

are used regularly in fields as diverse as astronomy; medicine; genetics; computer graphics systems; moviemaking; and earthquake prediction.




Figure 3 - **Fractal example — the Weierstrass function.** Published by Karl Weierstrass in 1872, and often cited as one of the first examples of fractals. The function is continuous, but cannot be differentiated at any point — it has no derivative in the form  $\frac{dy}{dx}$ . (Source: Wikipedia)




Figure 4 - **Fractal example — the Sierpinski Gasket.** The blue triangle has a smaller triangle punched inside, dividing it into three blue triangles, and so on forever. (Source: Wikipedia)

### 1.3. Multifractal finance vs. Standard finance

How do fractals have anything to do with finance? Simple: financial data also shows scaling behaviour. Mandelbrot (2012) writes: “All price charts look alike. [...] Strip off the dates and price markers and you cannot tell which is which. They are all equally wiggly.” In principle, the same formulas should be applicable to daily, monthly or yearly timescales.

Fractal modelling is meant to simulate three key features of financial markets, which we will keep coming back to. These are: 1) large price swings; 2) non-independence in price movements; and 3) the tendency for high-volatility periods to come in clusters. These ideas are explained in detail in the [theory](#) section.

In finance, fractal theories are mostly long-forgotten. But in the 1960's, there was high enthusiasm for fractal models of markets. The stakes are high — P. H. Cootner, an American economist, claimed that “if he [Mandelbrot] is right, almost all our statistical tools are obsolete [...] Almost without exception, past econometric work is meaningless.” Cootner then called for more empirical evidence to check Mandelbrot's work.<sup>4</sup> This paper aims to be a contribution to that evidence.

---

<sup>4</sup> Mandelbrot (1997), Pg. 9

For years, Mandelbrot has vehemently argued that most standard models underestimate the risks involved in financial markets, and how much money one really stands to lose<sup>5</sup>. The concepts of scaling, high kurtosis, volatility clustering, and long-term dependence are grave news for risk-averse investors. This is especially true for those investors who are strongly exposed to the short term.

The main point of using fractals in financial modelling is to incorporate the “anomaly” events — those wild, unpredictable swings that cannot be explained by standard financial theory — which seem to occur so often. Awareness of these concepts is a first step in preparing our portfolios for the worst. It can help us follow Charlie Munger’s advice — to know where we (or our portfolios) might die, so we could never go there.

It is important to note that the aim of these models is to be more *descriptive* rather than *predictive* in nature. The predictions are made to be qualitative first, and quantitative second. The aim is to get a better understanding of how financial markets generally work, and what kind of market dynamics one can expect; rather than to predict future price movements. In essence, one probably cannot use these models to become rich — but one may find them helpful in avoiding becoming poor.

#### [1.4. Three markets — stocks, bonds and currencies — Description of data and data sources](#)

The purpose of studying three completely different markets is to test Mandelbrot’s claim that fractal analysis is widely applicable — Mandelbrot (2004) argues that the same equations can easily be transformed to fit completely distinct market environments, whether be they stocks, bonds, currencies, commodities, or almost any other market. The main criterion is that the price data should show multifractal scaling.

The choice of the three markets is completely arbitrary — but, to my knowledge, none of them have ever been studied using the MMAR-model or any other type of fractal analysis.

The Norwegian krone has been the official currency of Norway for over 100 years. Meanwhile, the United States dollar is the most traded currency in the world. Thus, data for the

---

<sup>5</sup> The German fund manager Keppler (1990) offers a short and convincing argument that risk is not the same as volatility, and thus the terms should not be used interchangeably, as is often done in finance.

daily USD/NOK exchange rate is readily available from the Federal Reserve Economic Data database (the “FRED”). I use data from 25<sup>th</sup> January 1989 to 1<sup>st</sup> March 2019.

The Swedish OMXS30 index is probably the most important market index in Sweden. It represents the 30 most actively traded stocks on the Stockholm Stock Exchange. Its composition is revised twice per year. I use data from 18<sup>th</sup> November 1986 to 2<sup>nd</sup> January 2017. Notably, this is the only data I use which has not been taken from the FRED — instead, the OMXS30 data was taken from the National Data Center of the Swedish House of Finance.

LIBOR is an acronym for the London Inter-Bank Offered Rate. The LIBOR is a global index which indicates the benchmark rate at which major global banks would lend to each other. Each day, the Intercontinental Exchange asks the world’s major banks how much they would ask other banks for short-term loans — thus, the LIBOR is recalculated every day. In our case, we are using the rate for a 12-month loan in British pounds. But the LIBOR does not only concern banks — it is also important because its fluctuations can affect the rates on various consumer loans, such as those on credit cards, car loans or mortgages. I use data from 21<sup>st</sup> March 1989 to 1<sup>st</sup> March 2019.

Peters (1991) argues that 30 years of market data is needed to define a period for good fractal analysis. This is approximately the timespan that we will consider for all three markets. For all three datasets, I worked with both the raw price data and their logarithmic returns, as is common in economics. For the MMAR, no other type of data is necessary.

Today, the application of fractal tools is immeasurably easier than it was in the 1960’s or the 1990’s, owing to a tremendous increase in computing power and its availability. All simulations were performed on an ordinary personal computer using the Python programming language and the JupyterLab web interface.

The next section outlines the three important features of the MMAR-model, and explains why they may be relevant for financial modelling.

## II. Theory

### 2.1. Large price swings — Fat tails, kurtosis and Stable distributions

*“We were seeing things that were 25-standard deviation moves, several days in a row.”<sup>6</sup>*

David Viniar, former CFO of American investment bank Goldman Sachs, and current Board member, 2007

Many observers have noted that standard financial models significantly underestimate the likelihood of large — and particularly, extreme — price movements. Dowd et al (2008) point out that, under a normal distribution, the expected occurrence of a 25-sigma event should be once every  $1.3 \times 10^{135}$  years. By comparison, the Universe has only existed for about 14 billion years, i.e.  $1.4 \times 10^{10}$ , or about 125 zeroes short. Thus, it must’ve been one really, *really* unlucky day. And yet Goldman reported *several* 25-sigma days, in a row.

A simpler explanation may be that extreme price swings are more common than standard models presume. This section introduces a theoretical aspect for incorporating extreme events into financial models.


Most standard financial models employ the Normal distribution (also known as the Gaussian) as one of their core assumptions for modelling asset returns. Yet this assumption has been rejected repeatedly by many empirical researchers (more on this in [part three](#) of the literature review section). In fact, we can quickly check it on our own using two standard statistical tests. Below are the P-Values of two typical tests for Normality:

Normality Test / Country	Norway (USD/NOK)	Sweden (OMXS30 index)	Britain (12-month LIBOR in £)
Kolmogorov-Smirnov	0.0	0.0	0.0
Shapiro-Wilk	$3.228 \times 10^{-39}$	$2.522 \times 10^{-44}$	0.0

*Table 1 — Returns are not normally distributed. A table of P-values. Here, we test whether our returns could have been drawn from a Normal distribution. Since all P-values are practically zero, it is quite safe to reject the hypothesis that price returns are Gaussian.*

<sup>6</sup> As reported in the *Financial Times* on August 13, 2007. <https://www.ft.com/content/d2121cb6-49cb-11dc-9ffe-0000779fd2ac>

The reason the data fails the tests is because it exhibits a statistical property known as *excess kurtosis*. Kurtosis should be a relevant idea for investors, especially those who are exposed to short-run fluctuations. Kurtosis means that two data sets can have the same mean and the same standard deviation, while still looking completely different. Consider the demonstration below, in which both graphs have the same mean and the same variance (respectively, zero and 0.0072). The only difference is the kurtosis — the first was made with a Normally distributed random-walk, so it has an excess kurtosis of zero. Meanwhile, the second was made using the MMAR model, and has a kurtosis of 5.6:



*Figure 5 — High kurtosis explained.* Both graphs have the same mean and the same standard deviation, but look completely different. The left graph shows returns simulated with a Normal distribution, and thus its excess kurtosis is practically zero. The right graph shows simulations with high kurtosis, made with the MMAR model.

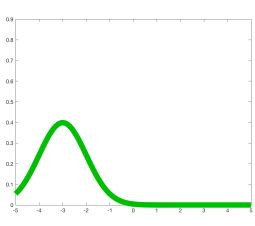
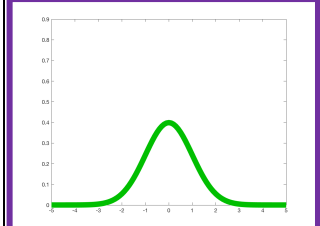
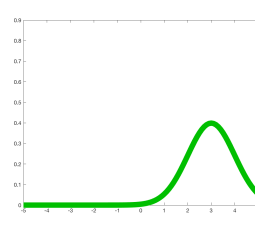
Kurtosis is a counterintuitive idea. Loosely speaking, it means that the data is very concentrated around the average — but when an outlier does occur, then it could in principle be almost any size. Human height is Normally distributed, with an average height for men around 175 centimetres. If height had high kurtosis, we might expect almost every man to be within a centimetre of 175 — but if we met a tall man, then he could almost equally likely be a head taller, or have the height of a house, or conversely, the height of an ant.

But where does excess kurtosis come from? To explain that, we will use a useful conceptual tool known as the *family of stable distributions*, also known as Levy distributions. Here, things get a bit technical. Say we have a distribution  $L(x_i)$  — a distribution is said to be *stable* if a combination of two independent random variables with this particular distribution (here two variables distributed by  $L(x_i)$ ) yields the same distribution, after rescaling. This is a generalized form of the probability density function, in which the standard Gaussian distribution is a special case. An analogy would be that playing a coin-toss game against ten people (of, say, ten tosses) would have the same expected value as simply playing one game against one person that was ten times longer (here, one hundred tosses), although the scale of possible outcomes would differ.

Stable distributions are described by four parameters —  $\alpha$ ,  $\beta$ ,  $\gamma$ , and  $\delta$ . These are more easily understood back-to-front: Firstly,  $\delta$  is the *locational parameter*. If  $\alpha$  is greater than one, then  $\delta$  is equal to the mean, as in the Normal distribution. Secondly,  $\gamma$  is the *scale parameter*, which measures the dispersion of data<sup>7</sup> — like the variance, although the two are not the same. Thirdly,  $\beta$  is the “*skewness*” parameter —  $\beta$  measures if the distribution is more concentrated on the left, right, or around the median. Its value ranges from -1 to 1, where 0 means perfect symmetry around the mode.

Finally,  $\alpha$  is the variable that affects the distribution’s kurtosis. It measures the “tail thickness” and peak height of the distribution. The value of  $\alpha$  ranges from 0 to 2. If  $\alpha = 2$ , we have a typical Normal distribution with skinny tails<sup>8</sup>, and events that are several standard deviations from the mean are extremely unlikely. However, if it  $\alpha$  is less than 2, then we have more concentration in the middle of the distribution, but the tails also decrease more slowly. This is what makes extreme events become more probable. Mathematically, it is as though the area under the probability curve were transferred away from the two middle-sections of both sides, and over to the centre and the tails instead. Formally, we can say that except for the Normal distribution, stable distributions are all *leptokurtic* — they all have “fat tails” and excess kurtosis compared to the Normal, which by definition has an excess kurtosis of zero.

Below is a table of illustrative examples to summarize the effects of each parameter. The Normal distribution is used as a reference case, signified by the purple boxes ( $\alpha = 2$ ;  $\beta = 0$ ;  $\gamma = 1$ ; and  $\delta = 0$ )<sup>9</sup>.

	<u>Range</u>	<u>Examples</u> (as compared to $\square$ = Normal distribution, where $\alpha = 2$ ; $\beta = 0$ ; $\gamma = 1$ ; and $\delta = 0$ )		
$\delta$  location (if $\alpha > 1$ , then this is the mean)	$\in \mathbb{R}$	 <p><math>\delta = -3</math></p>	 <p><math>\delta = 0</math></p>	 <p><math>\delta = 3</math></p>

<sup>7</sup> For  $\delta = 0$  and  $\gamma = 1$ , given that  $\beta = 0$  and  $\alpha = 2$ , we have the standard Normal distribution with mean 0 and a variance of 1.

<sup>8</sup> Other things being equal, i.e. given that  $\beta = 0$ .

<sup>9</sup> Note that when  $\alpha = 2$ , i.e. in a Normal distribution, the  $\beta$ -variable has no effect. Thus, I used the example of  $\alpha = 1$  here arbitrarily, merely for illustrative purposes.

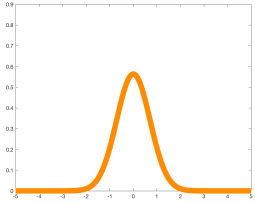
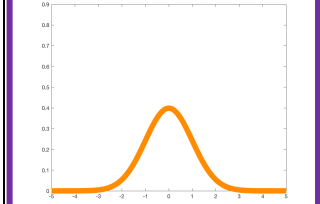
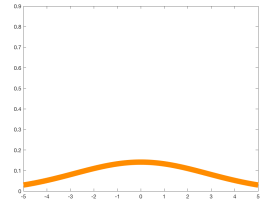
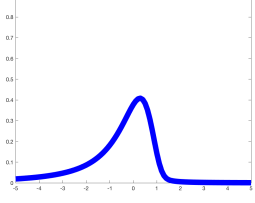
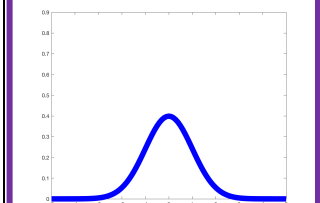
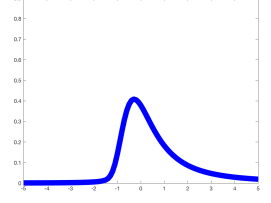
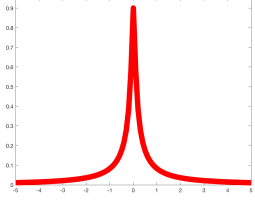
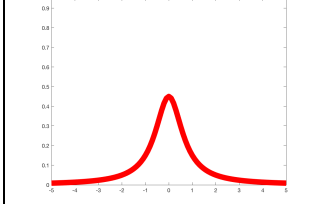
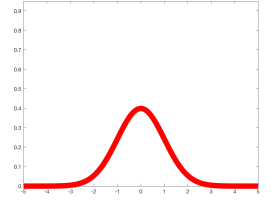
Scale (similar to the variance)	$\in \mathbb{R}^+$	 $\gamma = 0.5$	 $\gamma = 1$	 $\gamma = 2$
$\beta$ skewness (if $\beta = 0$ , then the distribution is symmetric; note that in the $\alpha=2$ case, $\beta$ has no effect)	$-1 \leq \beta \leq 1$	 $\beta = -0.9$ (given $\alpha = 1$ )	 $\beta = 0$	 $\beta = 0.9$ (given $\alpha = 1$ )
$\alpha$ tail thickness and peak height (kurtosis) (if $\alpha < 2$ , then the variance becomes undefined; if $\alpha \leq 1$ , then so does the mean)	$0 < \alpha \leq 2$	 $\alpha = 0.5$	 $\alpha = 1$	 $\alpha = 2$

Table 2 — *The monstrous family of Stable Distributions explained.* Apart from a few particular cases, these distributions can only be estimated — they cannot be written down in a generalized equation. Graphs outlined in purple are equivalent — they are all examples of the Normal distribution, also known as the Gaussian.

Technically, when data comes from a stable distribution, the empirical distribution function should show straight lines in a log-log plot of absolute price changes plotted against frequency, in which the slope of the line is equal to  $-\alpha$ . This is because low- $\alpha$  distributions follow what's known as a *power law* — an idea that's known in the popular conscience as Pareto's "80/20 principle". Mandelbrot (1963) was a controversial paper that showed that cotton prices followed a power-law distribution (I include the graph in [appendix 5](#)). His argument was that price returns should be understood as following a tall, pointy distribution with low  $\alpha$  and high kurtosis.

In a book on economic history, MacKenzie (2006) tells how these Levy distributions were originally branded as "monsters". Except for a few special cases like the Normal or the Cauchy, it seems that these distributions can't even be written down as a single equation — they do not have a distinct probability density function. They can only be estimated. This impractical limitation explains why they were not adopted by theorists or practitioners. For our purposes,

they are mainly useful as a conceptual tool. The model which we will be using to simulate markets — the MMAR — is designed to simulate returns with high kurtosis without the need to calculate one of these distributions.

But there is more to be said — high kurtosis has implications for traditional asset pricing. Besides fat tails, stable distributions have the curious property of having *undefined variance*. For any value of  $\alpha < 2$  — in other words, any non-Normal distribution — the variance generally increases with sample size, varies with different samples and never settles on a single constant value.<sup>10</sup> How would this look in the data? With a Normal distribution, as more and more data is added to the sample, the sample standard deviation ( $\sigma$ ) should converge with the real standard deviation of the population. This is not the case with other stable distributions (when  $\alpha$  is anything but 2). A consequence of the infinite variance assumption of stable distributions is that *the sample variance ( $\sigma^2$ ) never converges* — thus, neither does its square root, the standard deviation ( $\sigma$ ).

We can obtain evidence for the type of distribution that exists in markets by looking at a graph of *sequential standard deviation*. Take a data set and measure the variance. Measure it again as you add another observation, and then again, and so on. If market returns were Normally distributed, we would expect the variance to converge as we added more and more observations (when normalized, a variance of 1 and a mean of 0). Yet in fact, Peters (1994) shows that this is not the case for the Dow Jones Industrial index — one of the oldest indices on record — even over 100 years.

---

<sup>10</sup> In fact, if  $\alpha \leq 1$ , then even the mean of the distribution becomes undefined — the function ceases to converge on a single average.

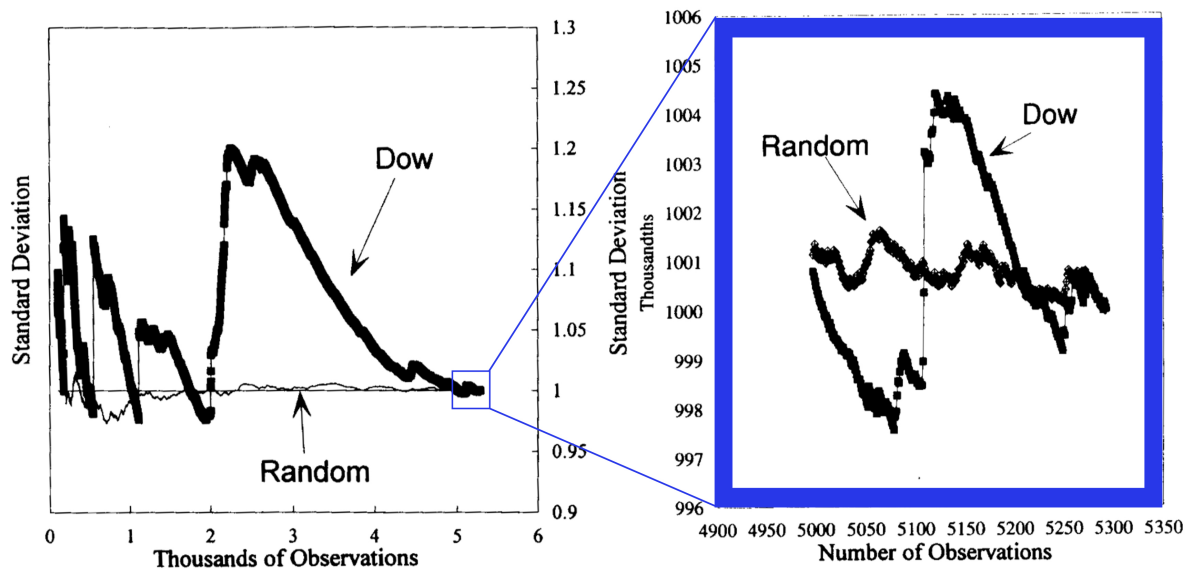


Figure 6 — *The standard deviation does not readily converge.* Here we see the sequential standard deviation of five-day returns for the Dow Jones Industrial Average (1888–1990), versus a simulation of a normally distributed variable (the “random” line). The blue-boxed, right diagram is a zoomed-in section of the left diagram. Source: Peters (1994).

Simulate a normalized Gaussian time-series, then measure its standard deviation each time as you add more and more observations — it should converge towards 1. Do the same with the Dow and compare it to your simulation. We can see that the Dow is more volatile than it should be if it were Gaussian. Furthermore, the sample variance can exhibit wild jumps — here, the fat scatter-points (the observations) do not form a continuous line, but rather a disjointed one. This is happening because in non-Normal Stable distributions, where  $\alpha < 2$  (meaning kurtosis is high), the population variance is theoretically infinite.

What’s more, the returns data exhibits large jumps — there are “discontinuities” in asset returns. This is evidence that — even in the long run — the Dow is characterized by a different type of stable distribution, with a different  $\alpha$ -value. The sample variance can never converge and will always exhibit discontinuities, because the (theoretical) population variance is infinite. A few extreme events — where the *squared* distance from the mean is very high — is enough to massively change the variance ( $\sigma^2$ ) for hundreds of observations. In [appendix 3](#), I show that the same non-convergence is visible in the returns of the three markets of this study.

Why does it matter if the variance never converges? The implication is this: if we are using standard deviation to calculate the asset’s price, then our price can change considerably if we just choose a different set of dates to calculate it. Since taking the variance of the last 5 years could give a considerably different price than if we took the variance of the last 10 years, then how do we know the right price? Should variance even be incorporated into the price? And how do we incorporate kurtosis, which is another great source of risk?

Note that, obviously, the real variance of the Dow is certainly not infinite. Nothing in real life resembles an infinite fractal structure, and broccoli buds cannot become smaller than atoms.

I merely argue that the assumption is conceptually useful to explain some market dynamics — namely, for explaining why the variance does not converge, even after 100 years.<sup>11</sup>

## **2.2. Dependence and clustering — the H-coefficient**

In standard finance, price movements are assumed to be independent events, but this does not necessarily have to hold. In fact, we might wonder whether returns are actually more likely to be *dependent*, i.e. where the auto-correlation is *not* zero? Real investors are often prone to fashions, bubbles, overconfidence and irrational exuberance, and this should be portrayed in the model.

Dependence can be measured by a tool known as the Hurst- or  $H$ -coefficient — a value between 0 and 1. The way to think about dependence is as follows:

Assume that price movement magnitudes are normally distributed — most moves are nearly zero and large moves are exceedingly rare. If  $H = 0.5$ , then the movements are perfectly independent: the correlation between movements is zero, and there is no way to predict the next price move; prices thus follow the canonical “random walk.”

However, if  $H$  is big ( $H > 0.5$ ), then prices can have persistent motion: a movement upwards, whether big or small, makes another upwards movement more likely, and the same holds for downwards movements. We would tend to see many up-movements in a row before they were reversed into down-movements.<sup>12</sup> In other words, up- and down-movements will *congregate* together.

Conversely, if  $H$  is small ( $H < 0.5$ ), then price movements will tend to cancel each other out quickly — an up-movement would be more likely to be followed by a down-movement. We would thus see fewer patterns in price movements than would be expected by chance and less congregation.

Differences in  $H$  can cause trends to appear in the price path — even though the data is really random, and at any point the trend might get reversed in the long run. We will see that the

---

<sup>11</sup> In fact, the multifractal model does not assume infinite variance. But this is a complicated technicality.

<sup>12</sup> More technically — for all time-lags,  $H > 0.5$  means the process exhibits significant positive and non-summable auto-covariance; whereas if  $H < 0.5$ , then it exhibits significant negative summable auto-covariance.

$H$  can change the range of price paths that a simulation is likely to take — high  $H$ 's *amplify* trends, and low  $H$ 's *dampen* them.




	<u>Value</u>	<u>Example</u>	<u>Trendiness</u>
<b>Big <math>H</math>;</b> persistent movements	$0.5 < H < 1$	 $H = 0.9$	Many apparent trends
<b>Random-walk case;</b> Perfectly independent movements	$H = 0.5$	 $H = 0.5$	Some trends (by chance)
<b>Small <math>H</math>;</b> Anti-persistent, cancellation movements	$0 < H < 0.5$	 $H = 0.1$	Few trends

Table 3 — **The  $H$ -coefficient explained.** If  $H > 0.5$ , then up-movements are more likely to be followed by another up-movement — the up- and down-movements “congregate” together, before eventually reversing in the other direction. But this reversal can take a very long time. (Image source: Mandelbrot 2004) These graphs assume Normally distributed price movements.

As we will see in [the literature review](#), the  $H$ -coefficient was originally developed for measuring rainfall and for engineering dams — when you want to build a high enough dam to prevent a flood, it matters if rain is completely random or if droughts and heavy rainfall usually occur for several days in a row.

This leads to one of Mandelbrot’s most famous ideas — high  $H$ 's might be a sign of *long-term dependence*. Correlations might even persist in the long run. For instance, price movements from many months or years before may be correlated with price movements today. In such cases, correlations start off strong (say, at 0.7), and then decrease slowly, but without ever quite getting to zero. However, interpreting the Hurst as a sign of long-term dependence in the markets is controversial (see [section 3.3](#)), so I will ignore the idea. I mention it for completeness.

Some technical notes: Note that although  $H$  implies autocorrelation — the statistical property where movements in a time series resemble (i.e. correlate with) earlier movements — the two concepts are not the same. The  $H$ -exponent measures and suggests a causal relationship — e.g. for a high  $H$ , up-movements propagate more up-movements. And crucially, we can

simulate data with a given  $H$ . Calculating the  $H$ -exponent by hand is very tedious, but many programming languages come with simple software packages for doing it. Fortunately, however, measuring the  $H$ -exponent is built into the mathematical derivation of the MMAR-model.

### 2.3. Multifractals and “trading time”

The trouble with assuming Levy distributions is that they imply infinite variance, which can manifest randomly. The trouble with simulating a high Hurst exponent is that it on its own, it'll still give you a basic, Normally distributed random walk. The trouble with both of these is that they don't solve another key feature found in market data — namely, the clearly discernible of periods of high volatility, when prices appear to go wild for prolonged periods and show “25-sigma moves for several days in a row”, and then go back to low volatility afterwards. It's not enough to have a model that generates high-kurtosis events every once in a while, we also want those days to *cluster together*. In other words, we also want to simulate occasional “market panics”.

This is where Mandelbrot's final key innovation comes in: the concept of “*trading time*”, represented by *multifractals*.

As with fractals, there are very few rigorous and formal definitions of a multifractal. But the following explanation may be helpful: A multifractal is like any ordinary fractal in that it shows obviously similar patterns at different scales. The difference is that those patterns will also be different in some obvious way. The cliché example is a mountain — both the top and bottom of mountain ranges will show a similar, unpredictable up-and-down pattern. But the low foothills of a mountain will be much smoother than the jagged, spiky mountain ridges at the top.




Figure 7 — **Multifractality in nature.** A mountain is a common multifractal found in the real world. It's foothills on the bottom are smooth and regular while the ridges at the top are sharp and spiky, whether they be big or small. Thus, both areas — and everything in between — are self-affine.

What does this have to do with finance? The answer is Mandelbrot’s concept of “trading time”. The idea comes from the simple experience that not every day is the same in financial markets — some days just “feel” faster than others. On some days, trading is very slow and boring, as though very little was happening. On these days, trade volumes are probably low. On the other hand, some days feel like a mess — prices fluctuate wildly, trade volumes go up, phones ring non-stop, and fortunes are made or lost. These are the fast days. They are similar to the slow days, except that prices change faster — they are more *volatile*. And usually, these fast days come in groups, rather than randomly out-of-the-blue. Multifractal “trading time” is the final component of the MMAR-model that we will use to simulate markets, and this is the component that gives us volatility clustering.

The technical version of “trading time” is explained in the [methodology](#). For now, I can say that the particular way in which multifractals are used in the MMAR model is by generating a *multiplicative cascade*. A cascade is a concept that is often used in geophysics. In the graphs below, the simplest version of a cascade is shown.

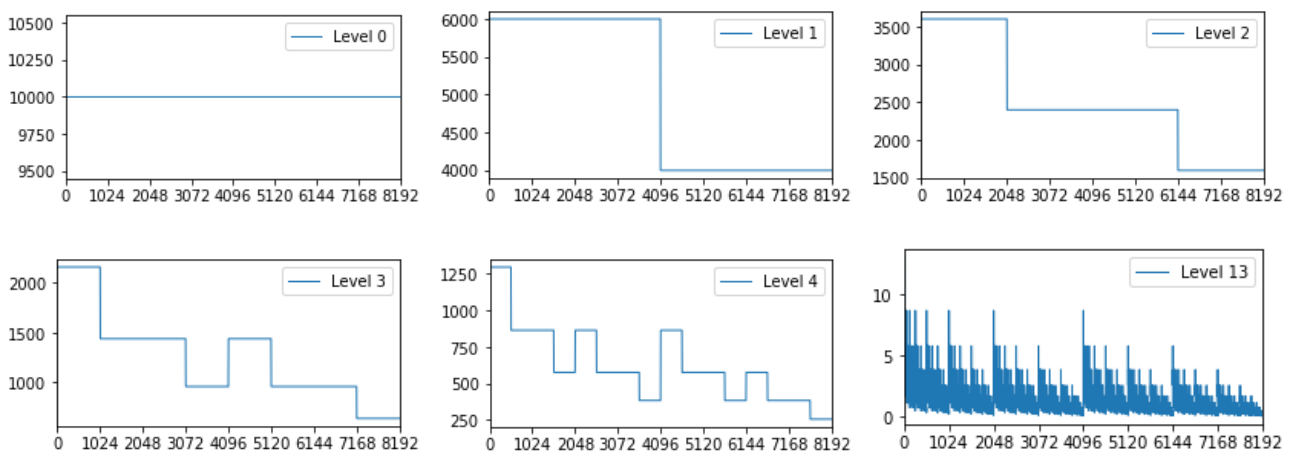


Figure 8 — **Step by step.** A basic multiplicative cascade. Here, we have  $b = 2$ ,  $m_0 = 0.6$  and  $k = 13$ . At each step, each section is divided in two parts. The left height side’s height is multiplied by the “mass”  $m_0$ , and the right side gets  $1 - m_0$ . Since  $m_0$  is bigger, the left section is taller than the right, and the leftmost section gets relatively taller and taller.

*This particular cascade is completely deterministic. Later we will introduce some randomness to the mass allocations.*

To understand this, imagine that the line at step 0 starts out with a “mass” of 10,000. At each step, the line is divided into two sections, one of which will have its mass multiplied by 0.6, and the other by the left over 0.4. At the next step, these two sections are themselves divided into two sections, for a total of four, and the same multiplication is applied. We can continue as long as we want, until we have eight, or sixteen, or thirty-two sections or more. Now, one doesn’t have to be a mathematician to see that this process is completely deterministic — it will be the same every time we do this, and we will always end up with a very tall peak on the left side. Later, we will introduce some randomness into generating these cascades.

## 2.4. Summary of theory

- Real returns have high kurtosis. We can think of this as the data having a low- $\alpha$  Stable distribution, with a tall peak and fat tails.
- These probability distributions imply that the sample variance never converges, because theoretical variance is infinite. The sample variance for the Dow Jones or our three markets does not converge. This suggests asset risks might be mispriced.
- The  $H$ -index can capture non-independent price movements, and can be simulated easily enough. High  $H$  means persistence; low  $H$  means anti-persistence.
- Although real returns can show high-kurtosis events, these events are not completely random — they usually come in clusters.
- To simulate these high-volatility, we can use the concept of “trading time”. In essence, this is the idea that trading can be fast or slow. We can construct a “trading time” simulation using fractal cascades.

### III. Literature Review

Research on fractal methods in finance is rare. It is unusual to see a paper with even 100 citations; by contrast, papers in science that use fractal methods can often gather thousands. This is despite Mandelbrot (2005) arguing that there is an “inescapable need for fractal tools in finance.”

The three original papers on the Multifractal Model of Asset Returns were by Mandelbrot, Calvet and Fisher in 1997, with each of the papers including the authors names in different order. These are related to other literature in [section 3.3](#).

This section will aim to summarize how fractal methods differ from those in standard finance, and to point out some published research in the field that is relevant to this study. We will start by considering some key assumptions in mainstream finance, from which the MMAR — and the fractal approach to finance in general — is meant to differentiate itself.

#### 3.1. Mainstream finance

Perhaps the first researcher to apply mathematical models to financial markets was the French mathematician Louis Bachelier (1900), in a PhD. dissertation titled “*Théorie de la Spéculation*” (“the Theory of Speculation”)<sup>13</sup>. The work was a study of price movements on the Paris exchange, the *Bourse*. The aim was to determine how derivative instruments could be priced using mathematical formulas. Despite being unimpressive to its contemporaries, the work was very original and soon became the core of most modern financial theories. Bachelier was perhaps the first to model price movements statistically, and to include many of the most common assumptions in today’s mainstream theories, including:

1. That price movements for financial instruments are statistically independent events, with a specific probability of moving up or down, in the style of the “random walk” hypothesis.
2. That the magnitude of price movements follows a Gaussian, Normal distribution, where the probability of large movements is extremely small.

---

<sup>13</sup> I admit that my interpretation of this story comes from Mandelbrot (2004), who may be somewhat biased. However, the historical account appears to be quite fair. Embrechts et al (1997) tell a similar story.

Much of modern finance has followed from this first paper and its assumptions, including the theories of the Efficient Market Hypothesis (the “EMH”) (Fama 1970); the efficient frontier and efficient portfolio theory (Markowitz 1952); the Capital Asset Pricing Model, or “CAPM” (Sharpe 1964); and the Black-Scholes model for option pricing (Black & Scholes 1973). Modern financial theory has largely been built on these earliest works.

Fractal models are often contrasted to those of the GARCH family. This is a sophisticated group of newer models, the first of which was developed by the American statistician Robert Engle, and for which he won the Nobel Prize in economics in 2003. Like the MMAR, GARCH tries to capture the features of volatility clustering and high kurtosis when it simulates market prices. Nevertheless, Mandelbrot, Calvet & Fisher (1997) have pointed out a number of differences between GARCH and multifractal models. One problem is that they imply a random-walk<sup>14</sup>. But the main issue seems to be that GARCH is *scale-invariant*. Simply put, you can’t use the same parameters for a long run simulation as you do for the short run. This is inefficient, and it suggests that forecasts might differ if we look at data with different timescales. Parameters have to be recalculated again and again. If we ignore this problem, then our simulated returns should start to look Gaussian if we increase the timescale.

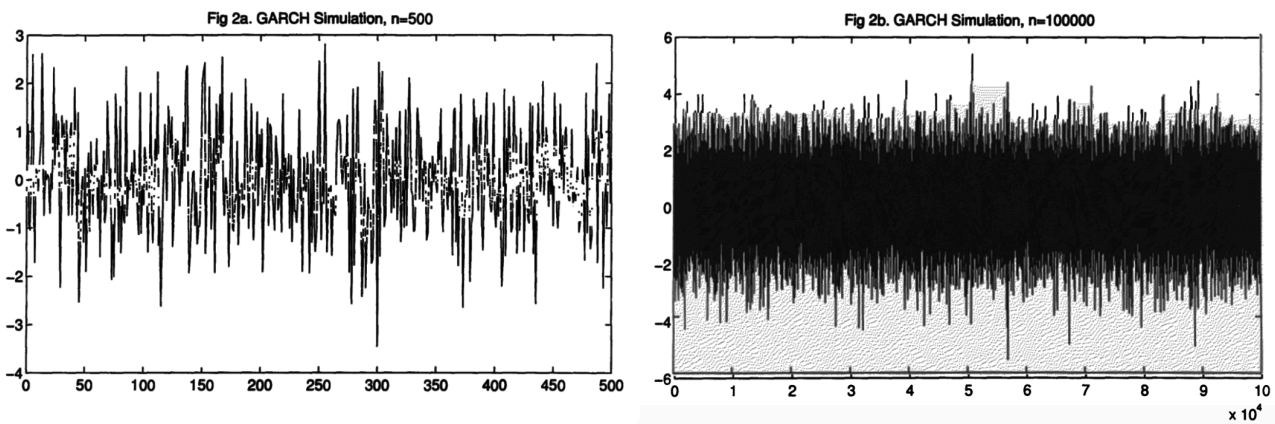


Figure 9 — **Apparently, GARCH doesn't scale well.** Mandelbrot et al's (1997) criticism seems to be that GARCH requires you to recalculate parameters each time when you work with a different timescale. Otherwise, the simulation looks practically Gaussian. Source: Mandelbrot, Calvet & Fisher (1997).

Fisher, Calvet & Mandelbrot (1997) run Monte Carlo simulations and find that GARCH and MMAR simulations are clearly distinguishable. Many of the other authors cited [in section 3.3](#) seem to agree. Lux (2001) looked into the topic, and concluded that the MMAR created a more realistic representation of market returns than GARCH, based on the Kolmogorov-Smirnov test.

<sup>14</sup> With some technically-extended exceptions, such as FIGARCH.

Empirical rejections of standard theories seem to be abundant. The empirical analysis of Cont (2001) rejects most of the tenets of the EMH, and cites volatility clustering as a “stylized fact” of financial timeseries. Lo & MacKinlay (2002) notably titled their book *A Non-Random Walk Down Wall Street* to parody the famous title by Malkiel (1985). They present mounds of technical evidence for non-independent price movements; and also argue that multi-factor models do not explain why prices deviate from the CAPM. Provocative titles are not limited to outsiders: Shiller (2000) offers *Irrational Exuberance* and Reinhardt & Rogoff (2009) offer *This Time Is Different* for examples of massive financial bubbles and poor theorizing leading to crises. Bundt & Murphy (2006) show that many economic variables — including unemployment, inflation, exchange and interest rates — follow a stable distribution with high kurtosis, rather than the Normal<sup>15</sup>. Stanley (2003) — an “econophysicist” — called for finance to recognize outlier events and gave the analogy of earthquakes, whose understanding was improved when they were shown to fit a log-log plot (previous theories of earthquakes could not account for the rare, devastating shocks.) Log-log plots are a staple of fractal financial analysis, and we shall see them when we calculate partition functions for the data.

What about standard theory in practice? As for usefulness, even Harry Markowitz himself seems to doubt his own theory — when asked about how he allocated his retirement, he reportedly answered: “I should have computed the historic covariances of the asset classes and drawn an efficient frontier. Instead ... I split my contributions fifty-fifty between bonds and equities.” (Zweig 1998) Meanwhile, the researcher Spyros Makridakis (2000, 2009 & 2018) has run several economic forecasting competitions, and concluded that most expert forecasters can predict very little, whether using theory, statistical prediction methods or machine learning algorithms.

### **3.2. Mandelbrot**

Benoît Mandelbrot is the “father of fractals” who we introduced in the introduction. Interestingly, much of Mandelbrot’s early academic work was in finance before entering the other sciences. Already in the 1960’s, Mandelbrot proposed a massive overhaul of financial

---

<sup>15</sup> Apparently, they also support the general Austrian Economics position of being “very sceptical of mathematical economics in general”.

theory and the methods used to study financial markets. Later, he argued that a more realistic framework could be made by using tools developed in fractal mathematics.

Much of the key literature on fractal finance was written by Mandelbrot himself. Perhaps the most famous and influential paper is “The Variation of Speculative Prices” (1963), which drew several price records and in particular those on cotton prices. This paper proposed that empirical data on price changes behaved as though it had infinite variance and a finite mean — as in a Stable distribution with  $1 < \alpha < 2$ . This was supported by some economists, including Eugene Fama, a graduate student of Mandelbrot and a later Nobel laureate; Fama (1965) wrote that “the Mandelbrotian hypothesis fits the data better than the Gaussian [normal distribution] hypothesis.”

To this day, Mandelbrot is probably the source of most of the relevant papers in the field: (1965) introduced long-term dependence; (1972) introduced the notion of a *multifractal*; although the word *fractal* was coined for the first time in (1975); (1989) explained how to form multifractal measures. Finally, his (1997) book compiled all these most of these papers and combined them in one work. Also, although the  $H$ -coefficient was developed by H.E. Hurst (1951) when he was studying the floods of the Nile as a water engineer in Egypt, it was Mandelbrot & Wallis (1968) who augmented the concept and showed its relevance to fractals.

Mandelbrot seems to have quite the well-deserved reputation as a maverick scientist, with a somewhat large ego<sup>16</sup>. His books are littered with tales of his personal arguments with economists. For instance, in his 1997 book, *Fractals and Scaling in Finance*, Mandelbrot rejects the standard use of the word “model.” Economists apply the term model to refer to a mathematical expression of an economic relationship, like “supply equals demand.” Instead, Mandelbrot sees models as statistical algorithms, whose aim is to produce simulations that are largely indistinguishable from actual data records (whether by eye or by statistical tools). Meanwhile, his 2004 book — written for a popular audience together with Richard Hudson, and also the main inspiration for this thesis — seems to show little patience for standard financial theories. His chapter titles include “The case against modern finance” and “10 heresies of finance”. Meanwhile, standard theories are described as a “sickly tree” grown from “a bad seed”; furthermore, “that people still lose money on these models should come as no great surprise.” This is quite the indictment!

---

<sup>16</sup> A reputation that he readily accepts, even calling himself a “scientific maverick” in the title of his memoirs.

But big ego or not, it is undeniable that Mandelbrot has made many contributions to science. His work, including on finance, is summarized clearly in *Chaos* by James Gleick (1989) — a popular science book on the history of chaos theory.

### **3.3. Other studies**

I like to compare multifractal analysis to standard academic finance by drawing an analogy to pop music and Heavy Metal. This is a clearly distinct subculture of technically skilled individuals, who appear to be intimately aware of each others' work, toiling valiantly (and unapologetically) outside the mainstream of their field, and occasionally entering the conventional spotlight. And like Metal, multifractal analysis ain't dead. A healthy dose of papers seems to be published each year in quantitative finance (or physics) journals. There even appears to be a whole website dedicated to outlining this research: <http://www.long-memory.com>.

Unfortunately, the field appears to be quite restrictive to those without the necessary mathematical training. Topics include technical analyses of Hurst parameter estimation<sup>17</sup>, Markov Chain analysis<sup>18</sup>, wavelet analysis<sup>19</sup>, and the more recent high dimensional MF-DFA ("multifractal detrended fluctuation analysis")<sup>20</sup>, and many other difficult mathematical concepts. Most of these methods build on the original assumptions of the MMAR-model. Ruipeng Liu (2008), a researcher in the field, provides a clear overview of multifractal models in his PhD thesis. I suspect that a key reason for why the field has been largely ignored is precisely because it is so punishing to those without a strong mathematical background, and its authors have mostly avoided trying to reach a popular audience. To my knowledge, Mandelbrot's 2004 book is the only such attempt.

One interesting researcher to note is Edgar E. Peters (1991, 1994 & 1999). Peters is an asset manager who employs methods from fractals and Chaos theory in his financial analyses, and occasionally publishes in academic journals. His first two books are often used as introductions to fractals and chaos in finance, and his name recurs throughout this thesis.

---

<sup>17</sup> E.g. Barunik & Kristoufek (2010); Kristoufek (2010).

<sup>18</sup> E.g. Calvet & Fisher (2008); Zhong & Zhiao (2012).

<sup>19</sup> E.g. Muzy et al (1994) ; Kristoufek (2013).

<sup>20</sup> E.g. Stanley (2003).

Despite its mathematical complexity, the field still seems to leave plenty of room for interpretation. For one thing, measuring the Hurst exponent is an art in itself, and discussing its interpretation and measurement is a whole technical field of its own. It may be interesting to note that there doesn't seem to be any consensus on how to measure the  $H$ -exponent. One would think that there would be some progress in this field, seeing as it's one of the most important factors underlying the whole theoretical framework! Regarding methodology, the classic method was the R/S analysis method, which Mandelbrot used in the 1960's and which was used to motivate research in fractal finance. R/S analysis is the original and most common way is to measure  $H$  — although, notably, this is not the method used for the MMAR! In fact, I've been unable to find a single other methodology that measures  $H$  in the same way.

Measuring  $H$  is straightforward, but data-intensive — it requires a large number of data points over a large timescale. Furthermore, calculations increase exponentially. Peters (1991) points out two potential under-sampling problems when performing R/S analysis — short time horizons, which don't allow fractal processes to reveal themselves; and using data whose frequency is too low, which misses cycles.

But it seems that the original estimations of  $H$  should have been accurate. Barunik & Kristoufek (2010) compared five different methodologies for measuring  $H$  by running several extensive Monte Carlo simulations on Stable distributions, varying the  $\alpha$  parameter from 1.1 to 2 (from very heavy tails to Normal distribution). Although they found that the least biased and least variable method was the Generalized Hurst Exponent approach (meaning it had the narrowest confidence interval), this method is more complicated than the simple R/S analysis, which was still the second most robust method even under heavy tails.<sup>21</sup>

I must mention here one of Mandelbrot's most famous ideas — namely, long memory and long-term dependence. This is the idea that price movements show autocorrelation even in the very long run, which suggests that they may be in some sense predictable. A high  $H$  is meant to imply long memory. As a technical finding, the long memory of time series has been reported in

---

<sup>21</sup> The *pracma* package of the *R* language includes five common methods for measuring  $H$ . One method is that of Peng et al (1994), a team of biologists, created a new method for measuring the  $H$ -exponent while working with DNA nucleotides. Using Monte Carlo simulations, Weron (2002) found this method to be the most effective. However, Barunik & Kristoufek (2010) found that most of these methods are not effective for heavy-tailed distributions, such as those in finance. For those datasets, common R/S analysis is robust enough.

the form of hyperbolic decay rates for the autocorrelation function of asset returns, first reported by Taylor (1986) and “now well documented” (according to Cont 2001). Or is it? In fact, the literature seems to be mired by contradictions and semantic arguments. For instance, Lo & MacKinlay (2002) find non-independent price movements, but reject this as a sign of long-term dependence. Cheung & Lai (1995) found little support for long memory as well. Because of this controversy, this thesis will ignore the concept of long-term dependence and its implications.

This paper restricts itself to one of the original and simpler forms of multifractal analysis — the first ever multifractal financial model, known as the Multifractal Model of Asset Returns (“MMAR”). The model was developed by Benoît Mandelbrot, Laurent Calvet and Adlai Fisher when the latter two were graduate students. The model was exposed in a compendium of three papers in 1997. Each co-author got to have his name listed as the first author for one of the three papers — Mandelbrot’s paper introduced the model, Calvet’s showed many of the mathematical derivations, and Fisher’s looked at the USD-Deutschemark rate as the first empirical study using the MMAR-model. This is the model that Mandelbrot (2004) advocates in his popular book.

Are these MMAR simulations realistic? Lux (2001) finds that simple multifractal measures not only make realistic representations of the volatility-clustering behaviour of markets, but also make the right type of unconditional distribution for asset returns. So the simulations seem to fit quite well. In fact, the author concludes that, in a double-blind test, he would be unable to distinguish between real financial data and that which was simulated by a multifractal cascade process. His metric was the Kolmogorov-Smirnov test — and we will use the same test in [section 5](#). Spoiler: my interpretation will be *less* optimistic than his.

### **3.4. Summary of Literature Review**

- Much of financial theory is built on assumptions of independence and Gaussian returns. GARCH is a notable exception, but it is scale-invariant.
- These assumptions seem to be undermined by real economic data, which often shows evidence of dependence, high kurtosis and market crashes. The usefulness of standard models appears uncertain.
- Mandelbrot has loudly denounced this for decades. He is the source of most of the original literature on fractal finance, but his work also shows signs of a big ego. The MMAR model was developed by him and his two students as a solution to these criticisms.

- Today, fractal finance is a stable field, albeit a small one. Its members argue that fractal assumptions fit the data better. However, their work is very technical and mathematically challenging, which may be restricting newcomers.
- The KS-test can be used to measure whether the MMAR is realistic.
- Measuring  $H$  is complicated business, and there doesn't seem to be a consensus on the method nor its interpretation. The MMAR has its own unique method.

## IV. Methodology

We will be exploring the predictions of the Multifractal Model of Asset Returns (the “MMAR”). To do this, we are going to generate 10,000 Monte Carlo simulations for each of our three markets.

We can see that this is a study of the long-term — we will be simulating 8,192 days of data. As with R/S analysis, the MMAR-model is different from standard statistical analysis in that having a lot of data points is only a secondary consideration — what’s more important is that the data encompasses *as many economic cycles as possible*. In other words, it’s more important to have data that goes back several years (or preferably, decades), rather than having data points for every minute of the day. This gives the best chance of estimating realistic parameters.

In fact, daily data is probably the easiest to work with. The short time-increments of high-frequency data are likely to produce statistical noise, or other complicated elements of market microstructure (for a discussion, see Peters 1991). In the third of the three original MMAR papers, Fisher, Calvet & Mandelbrot (1997) needed to apply complicated filtering for their high-frequency study of the Deutschemark.

We will be working with daily data, although in principle our data could have been monthly, minute-by-minute or millisecond-by-millisecond.

Below I offer a step-by-step way to generate MMAR simulations. This is perhaps an unusual way of describing a methodology, but I think this is probably the easiest way to understand the process. However, before we can proceed, we must define some key ideas.

### 4.1. 20+ Steps for Generating a Multifractal Simulation

#### Definition of a multifractal stochastic process:

A stochastic process  $X(t)$  is defined as multifractal if it has stationary increments, and also satisfies:

$$\mathbb{E}[|X(t)|^q] = c(q)t^{\tau(q)+1} \text{ for } t \in \mathcal{T}, q \in \mathcal{Q}$$

where  $\mathcal{T}$  and  $\mathcal{Q}$  are intervals on the real line. We also assume that  $\mathcal{T}$  and  $\mathcal{Q}$  have positive lengths, and that  $0 \in \mathcal{T}$ ,  $[0,1] \subseteq \mathcal{Q}$ .

This is a technical definition that is of little interest to non-mathematicians, so I won't discuss it much. I include it here for completeness.

### Definition of the MMAR-model:

The MMAR is a model for simulating growth over time. It consists of two components — a multifractal cascade and a fractional Brownian motion. The model aims to capture three important “stylized facts” of financial markets: 1) fat-tailed distributions; 2) non-independence in price changes (i.e.  $H \neq 0.5$ ); and 3) highly volatile days tending to cluster together (which is made using the concept of “trading time”).

If we define a stochastic process  $X(t)$  as  $X(t) = \ln P(t) - \ln P(0)$ , for the arbitrary time period  $P(t): \{0 \leq t \leq T\}$ , then the MMAR proposes that  $X(t)$  is a compound process in the form:

$$X(t) \equiv B_H[\theta(t)]$$

where  $B_H[t]$  is a Fractional Brownian Motion process with the Hurst exponent  $H$ , and  $\theta(t)$  is a cumulative distribution function of a multifractal cascade measure defined on the time period  $[0, T]$ . Thus, in effect, the trading time  $\theta(t)$  acts as the time parameter  $t$  for the  $B_H$  function. The relevant Hurst exponent  $\hat{H}$  is estimated before beginning the simulation.

The MMAR is estimated using four parameters —  $H$ ,  $\alpha$ ,  $\lambda$  and  $\sigma^2$ . These must be estimated empirically before we can construct a market simulation.

These definitions are complicated, but they might make more sense as we go through the simulation process step by step.

## 4.2. Part 1: Estimating the multifractal parameters

### Step 1:

**Get your data. Preferably with a highly composite number of data points.** It should be data on *prices* (closing price makes the most sense) for some type of asset — e.g. stocks, bonds, currency etc.

Daily frequency is good. Higher frequencies (minute-by-minute or millisecond-by-millisecond prices) can also work, but then you would have to do some complicated readjusting.

For the number of observations, I advise to choose a *highly composite number* (plus one, as we will be calculating rates of return and thus need an extra day). This is an integer that can be divided by a lot of factors. For example, the ancient Greeks famously chose 360 as the number of degrees in a circle, because this number can be divided by many factors and still yield an integer number. Specifically, 360 has 24 integer factors, including 180, 120, 90, 60, 45, 30, 15, and so on. Some good numbers to use are 5040 (sixty factors); 7560 (sixty-four factors); 10080 (seventy-two factors). The reason for this choice will be apparent when we calculate the partition function in step 5.

I will use 7560 +1 (meaning 7561) days of daily closing-price data for each market (equivalent to around 30 years of data with about 250 trading days per year). The factors<sup>22</sup> for 7560 also grow in a reasonably linear way in a log graph.




Figure 10 — **Highly composite numbers make life easy.** When analysing returns, will need to consider the data in many different-sized increments. Highly composite numbers let us look at 1-day periods, 7-day periods, 30-day periods etc. without running into problems with remainders. The factors of 7560 also grow almost linearly in a logarithmic graph. This is why I look at 7560 days of data for each country.

## Step 2:

**Calculate the log-returns between every observation.** This is sort of like the %-change between every observation, except we are using natural logs instead (as is common in economics). The formula is:

$$return = \ln \frac{P_{t+1}}{P_t}$$

---

<sup>22</sup> The following are the sixty four factors of 7560 — 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 12, 14, 15, 18, 20, 21, 24, 27, 28, 30, 35, 36, 40, 42, 45, 54, 56, 60, 63, 70, 72, 84, 90, 105, 108, 120, 126, 135, 140, 168, 180, 189, 210, 216, 252, 270, 280, 315, 360, 378, 420, 504, 540, 630, 756, 840, 945, 1080, 1260, 1512, 1890, 2520, 3780, 7560.

Where  $P_t$  is the price of the asset at time  $t$  for  $0 \leq t \leq T$ .

For instance, if you have daily data, then your first log-return will be  $\ln \frac{P_1}{P_0}$ , the second will be  $\ln \frac{P_2}{P_1}$ , and so on until the end.

### Step 3:

**Imagine a stochastic process  $X(t)$  and its increments  $X(t, \Delta t)$ .** The formulas are:

$$X(t) = \ln[P(t)] - \ln[P(0)]$$

and

$$X(t, \Delta t) \equiv X(t + \Delta t) - X(t)$$

for  $0 \leq t \leq T$ , where  $t$  is some particular instance of time since “time zero” and  $T$  is the last time instance under consideration. (By this definition  $X(0) = 0$ .)

Interpreting in words,  $X(t)$  is the total growth of our price since the beginning of time (where  $t = 0$ ) — not in percentages, but rather in powers of  $e$ .<sup>23</sup> For example, if  $X(t) = 2$ , that means that the price at time  $t$  is  $e^2 \approx 7.389...$  times bigger than it was in the beginning.

Meanwhile,  $X(t, \Delta t)$  is our next increment of price growth for some given change in time  $\Delta t$  — how much is the price going to change between times  $t$  and  $t + \Delta t$ ? — at every time instance  $t$ .

### Step 4:

**Choose some values of  $q$  (the raw statistical moments) which will be used for your empirical estimations.** Some notes:

---

<sup>23</sup> The mathematical constant  $e$ , equal to 2.71828..., is used as the base for the natural logarithm  $\ln()$ . It is sometimes called Euler's number.

Mandelbrot et al (1997) used 99 values of  $q$  between 0.01 and 30 (mostly concentrated at values around 2.0) to estimate  $\hat{\tau}(q)$ . Computers are now more powerful and versatile than they were two decades ago, so we will use 121 values of  $q$ , similarly concentrated around 2.0.<sup>24</sup>

### Step 5:

**Find the partition function.** This is done by splitting the data into  $N$  non-overlapping subintervals, whose length  $\Delta t$  is less than our total time  $T$ :

$$S_q(T, \Delta t) = \sum_{i=0}^{N-1} [|X(i \times \Delta t + \Delta t) - X(i \times \Delta t)|^q]$$

where  $N = \frac{T}{\Delta t}$ , i.e. the total number of time increments for that particular increment size.

This function is quite complicated to understand, but simple to compute. The partition function is the sum of the absolute values of all increments of price change, raised to some given power  $q$  (which is the raw statistical moment), for every size of time increment  $\Delta t$ . The time increments  $\Delta t$  are the dependent variable that we manipulate (the x-axis), and the partition function  $S_q(t, \Delta t)$  is the resulting dependent variable (the y-axis). This function is computed several times for different values of  $q$ .

How should we choose the  $\Delta t$ 's? We can't just choose all the integers less than  $T$  — if we have 1000 days of data, it doesn't make sense to choose a  $\Delta t$  increment of 999 and then have a remainder of 1 day. Nor can we really use 998, or 997 etc. This is why it was helpful to choose a highly composite number of data points, in which case it will be easy to divide the dataset into a high number of factors while also avoiding any complicated issues of overlapping increment<sup>25</sup>.

---

<sup>24</sup> The following is the list of  $q$ 's which we will use for estimating  $\hat{\tau}(q)$  — 0.01, 0.1, 0.2, 0.3, 0.4, 0.5, 0.55, 0.6, 0.65, 0.7, 0.75, 0.8, 0.85, 0.9, 0.95, 1.0, 1.05, 1.1, 1.15, 1.2, 1.25, 1.3, 1.35, 1.4, 1.45, 1.5, 1.55, 1.6, 1.65, 1.7, 1.75, 1.8, 1.81, 1.82, 1.83, 1.84, 1.85, 1.86, 1.87, 1.88, 1.89, 1.9, 1.91, 1.92, 1.93, 1.94, 1.95, 1.96, 1.97, 1.98, 1.985, 1.99, 1.991, 1.992, 1.993, 1.994, 1.995, 1.996, 1.997, 1.998, 1.999, 2.0, 2.001, 2.002, 2.003, 2.004, 2.005, 2.006, 2.007, 2.008, 2.009, 2.01, 2.015, 2.02, 2.025, 2.03, 2.04, 2.05, 2.06, 2.07, 2.08, 2.09, 2.1, 2.15, 2.2, 2.25, 2.3, 2.35, 2.4, 2.45, 2.5, 2.6, 2.7, 2.8, 2.9, 3.0, 3.1, 3.2, 3.3, 3.4, 3.5, 3.6, 3.7, 3.8, 3.9, 4.0, 4.5, 5.0, 6.0, 7.0, 8.0, 9.0, 10.0, 12.5, 15.0, 17.5, 20.0, 22.5, 25.0, 27.5, 30.0

<sup>25</sup> It is possible to do this if really necessary, but it's an unnecessary complication that doesn't help the interpretability of results.

**Step 6:**

**(Optional) Plot a log-log graph of the expectation of the partition function versus the time increments (meaning  $\ln[S_q(T, \Delta t)]$  against  $\ln[\Delta t]$ ) for different values of moments  $q$ .** The values of the partition function should be normalized (meaning divided by the first value for that  $q$  i.e.  $\ln[S_q(T, 0)]$ ). This is because we are not studying the behaviour of the intercepts (the  $c(q)$ 's).

According to the definition of (multi-)fractality, a random process  $X(t)$  is multifractal if: 1) its increments are stationary; and 2) it satisfies the equation:

$$\mathbb{E}[|X(t)|^q] = c(q)^{\tau(q)+1}$$

This definition means that the movements of the  $X(t)$  process have to satisfy a clear scaling relation, which can be modelled by the function  $\tau(q)$ . The partition function  $S_q(T, \Delta t)$  should have a similar shape for all values of  $q$ , and show an increase in slope as  $q$  goes up.

$$\mathbb{E}[S_q(T, \Delta t)] = c(q)(\Delta t)^{\tau(q)+1}$$

This idea is explained visually in the diagram below:

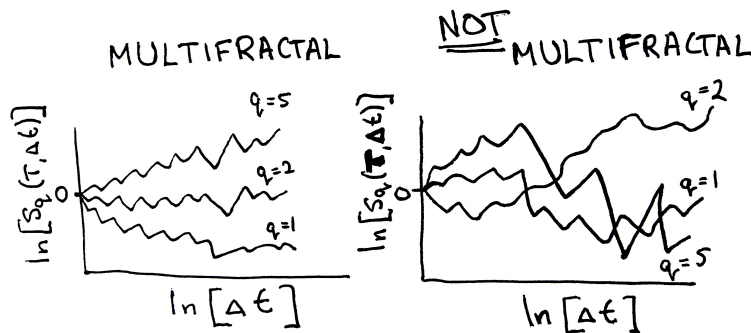


Figure 11 — **How to look for fractals.** If fractality is present, the partition function should show an obvious scaling relationship. Multifractality means that the scale increases in a non-linear way for each increase in  $q$ .

Because of normalization, the partition function for each moment  $q$  should start at zero. Normalization makes it easy to see visually if the slope of the partition function increases as  $q$  increases.

You don't need to plot lines for all the  $q$ 's that you will actually use to estimate  $\tau(q)$  (which should be a lot of  $q$ 's). The convention is to plot about five, e.g. for  $q = 1, 2, 3, 4$  &  $5$ .

**Step 7:**

**Estimate the scaling function  $\tau(q)$  by running OLS linear regression on the different lines of the graph for different values of  $q$ .** Using the following formula<sup>26</sup>:

$$\ln \left[ \mathbb{E}[S_q(T, \Delta t)] \right] = \widehat{\tau}(q) \ln[\Delta t] + \widehat{c}(q) \ln[T]$$

we estimate the scaling function  $\hat{\tau}(q)$  and the prefactor  $\hat{c}(q)$ .

Under this definition, we run an Ordinary Least Squares (“OLS”) regression on the partition function curve for each  $q$ . This gives us a straight line for every time increment  $\Delta t$ . We treat these OLS regressions as the expected values of the partition function for each  $q$ . The gradient of the OLS line is the value of  $\tau(q)$  for that value of  $q$ .




Figure 12 — **How to calculate a scaling function.** For every partition function curve, estimate an OLS line going through it. The slope of the OLS line is the value of the scaling function for that value of  $q$ .

Through all the OLS's, you should have lots of different values for  $\tau(q)$  and the intercept  $c(q)$  for different values of  $q$ . The prefactor  $c(q)$  can mostly be ignored. Next, you have to use *non-linear regression* to estimate a *non-linear* curve  $\hat{\tau}(q)$  that goes through all of the estimated  $\tau(q)$  constants as closely as possible.




Figure 13 — **How to interpret a scaling function.** Non-linearity in  $\tau(q)$  is a common symptom of multifractality.

One technical issue arises here: usually, the scaling function is not very robust for high moments — after some point, the observed  $\tau(q)$  often starts becoming linear. This is a

<sup>26</sup> Adapted from Mandelbrot et al (1997) paper 3, page 4.

consequence of the real world not being perfectly multifractal, like an abstract mathematical shape. Thus, *it is better to run the non-linear regression only for lower moments*. It is not strictly necessary to get the function in continuous terms, but it makes the rest of the process easier, and ignoring the higher moments doesn't have much effect on the results. I estimate the non-linear curve  $\hat{\tau}(q)$  by only using observed  $\tau(q)$  values up to 4.0. This is also why the selection of  $q$ 's was concentrated around 2.

Also, from the definition of the partition function, we get that:  $\hat{\tau}(0) = -1$ . Finally,  $\hat{\tau}(q)$  should be concave.

### Step 8:

**Find Hurst exponent  $H$ .** This calculation is very easy. Use the estimated scaling function  $\hat{\tau}(q)$  and the following formula:

$$\tau_X\left(\frac{1}{H}\right) = 0$$

$$\therefore \hat{H} = \frac{1}{\hat{\tau}^{-1}(0)}$$

Simply put, find the point on the estimated scaling function where  $\hat{\tau}(q) = 0$ , then take its inverse. (The subscript  $X$  in  $\tau_X$  simply means that this is the scaling function for our stochastic price process.)

How do we know that  $\tau_X\left(\frac{1}{H}\right) = 0$ ? These relationships are all derived in the Mandelbrot et al (1997) papers.

### Step 9:

**Estimate the multifractal spectrum  $f(\alpha)$  by performing a Legendre transform on  $\hat{\tau}(q)$ .** A Legendre transformation is a fancy way of converting from some function with some variables.

For example, we may wish to convert some function with one set of variables (e.g.  $f(x)$  with  $x$ 's) to another function with some other variables (e.g.  $H(p)$  with  $p$ 's), where both

functions can be used to describe the same process.<sup>27</sup> The Legendre transform function (let's call it  $H(p)$ ) is defined as:

$$H(p) = \sup[xp - f(x)]$$

where  $p = \frac{df(x)}{dx}$  and  $\sup$  refers to the *supremum*. So the derivative of  $f(x)$  becomes our new dependent variable.

Defining the multifractal spectrum is very, very technical. According to the three MMAR papers, a *multifractal spectrum*  $f(\alpha)$  is a function that “describes the local growth rate in a multifractal process.” Thus, the spectrum “describes the fractal dimension of the set of instants having a given local [Hölder] exponent”. The fractal dimension measures the smoothness of a function at a given point. Conceptually, the multifractal spectrum allows the Hurst-exponent to vary over time, in which case we might call it the “Hölder exponent”. Confusingly, this is a different name for essentially the same idea, except the Hölder exponent can range from positive to negative infinity. We will mostly ignore the Hölder exponent. For our purposes, we only need the multifractal spectrum to estimate the second MMAR component — the “most probable” Hölder exponent,  $\alpha_0$ .

The scaling function  $\hat{t}(q)$  should be concave, which means we can define the multifractal spectrum of the price process as:

$$\hat{f}(\alpha) = \inf[q\alpha - \hat{t}(q)]$$

where  $\inf$  refers to the *infimum*.

The aim is to convert our quadratic scaling function  $\hat{t}(q)$  into a quadratic non-linear function in terms of  $\alpha$ . This involves differentiation. I outline my method for performing the Legendre transform in [appendix 1](#).

---

<sup>27</sup> For example: in physics (namely Classical Mechanics), there are two common approaches to describe motion — the Lagrangian and the Hamiltonian. The Legendre transform can be used to convert a Lagrangian function  $L\left(q, \frac{dq}{dt}, t\right)$  which considers a particle's velocity to a Hamiltonian function  $H\left(q, \frac{dL}{dq}, t\right)$  which considers its momentum.

**Step 10:**

**By using the estimated multifractal spectrum  $\hat{f}(\alpha)$ , find the most probable Hölder exponent  $\alpha_0$ .** Using the following definition:

$$\hat{\alpha}_0 = \hat{f}^{-1}(\max[\hat{f}(\alpha)])$$

Basically,  $\alpha_0$  is the value of the Hölder exponent  $\alpha$  for which  $\hat{f}(\alpha_0) = 1$ . This should also be the maximum value of the multifractal spectrum.

The interpretation for  $\alpha_0$  is this: if the data is multifractal — meaning that it can have different Hölder exponents at different timepoints — then  $\alpha_0$  is the most commonly occurring, most “dominant” Hölder exponent in the price data.

**Step 11:**

**Using the estimated values for  $H$  and  $\alpha_0$ , estimate the log-normal distribution parameters for the multiplicative cascade —  $\lambda$  and  $\sigma^2$ .** The MMAR involves generating a multiplicative cascade, whose masses are drawn from a continuous probability distribution.

The most common distribution used for these cascades<sup>28</sup>, and the one we will use, is the *log-normal* distribution, with a mean  $\lambda$  and a variance  $\sigma^2$ . Using the estimations that we found for  $\alpha_0$  and  $H$ , the second MMAR paper shows that these final two parameters are easily determined as follows<sup>29</sup>:

$$\hat{\lambda} = \frac{\hat{\alpha}_0}{\hat{H}}$$

$$\hat{\sigma}^2 = \frac{2(\hat{\lambda} - 1)}{\ln[b]}$$

We will use these two values to generate randomly distributed multipliers  $M_i$  when we construct multifractal cascades for the MMAR simulation.

---

<sup>28</sup> It is also possible to use Binomial, Poisson, and Gamma distributions (Calvet & Fisher 2002).

<sup>29</sup> For reference, all of these relations are derived in the three Mandelbrot et al (1997) papers.

### 4.3. Part 2: Constructing an MMAR simulation

#### Step 12:

**Choose how many data points you want to simulate.** Here's how you should make this decision:

The first step is to decide what data frequency you're thinking of (monthly, daily, hourly etc.), and then the total time period (30 minutes, 30 days, 30 years etc.).

For the next step, consider that you will simulate a total  $b^k$  data points (where  $b$  and  $k$  are both positive integers), so try to choose some values for  $b$  and  $k$  that will get you close to the number of data points you want to simulate. I recommend keeping  $b = 2$  (this is the standard value, used for the multiplicative cascade later). That way, for example, we could simulate 4,096 data points, or 8,192 points, or 16,384, or any other power of 2.

Since I'm working with around 30 years of daily data, which is around 7560 data points for each market (weekends and holidays not included), I will use  $b = 2$  and  $k = 13$  for 8192 simulated data points. With 250 trading days per year, that amounts to around 32 years of simulated daily prices. I do this 10,000 times for each of the three markets, which gives a total of about 82 million simulated days per country.

#### Step 13:

**Generate a log-normal multiplicative cascade  $\mu$ .** This is a very complicated process, so perhaps it is best to first explain how multiplicative cascades work.

First, choose the number of  $b$ -adic intervals  $b$  — this is the number of intervals into which you will divide your time series at every step  $k$ . Thus, at every step  $k$ , there should be  $b^{-k}$  intervals total.

An easy number of intervals is 2, and in fact, this is the most commonly used number for MMAR simulations. With  $b = 2$ , we have 1 interval at step zero; 2 intervals at step one; 4 at step two; 8 at step three; and so on.

Next, choose some masses, such as  $m_0$  and  $m_1$  if using a 2-adic cascade. At every step  $k$ , the "mass" of every interval (i.e. the height of the graph for that interval) will be multiplied by the allocated mass  $m_i$ . If we use  $m_0 = 0.6$ , and  $b = 2$ , then this gives us precisely the same cascade as the blue one that we saw in the [theory](#) section.

This is a very simple cascade, but for our purposes, we shall need to give it some extensions. With some thinking, we can see that if we were to continue this process indefinitely, we would get a predictable graph with a series of bars that look like scaled versions of each other, with one very long bar (tending to infinity, in fact) on the side where the higher  $m_i$  gets multiplied repeatedly. This is not what we want — what we want is some *randomness* in the system. So instead of having static mass allocations, give them some probabilities of occurrence. For example, if using 2 intervals, the probability of  $m_0$  occurring on the left could be  $\varphi_0 = 0.6$ , and thus 0.4 for it occurring on the right. This means that  $m_1$ 's probability of occurring on the left would be the leftover,  $\varphi_1 = 0.4$ . (These are the most common numbers used for demonstration.) Alternatively, we can make both masses ( $m_0$  and  $m_1$ ) have equal probabilities of appearing on either side ( $\varphi_0 = \varphi_1 = 0.5$ ). Similar rules can be set up for any number of intervals  $b$  higher than 2.

Now, so far our mass allocations have been mutually exclusive — if one side got  $m_0$ , the other had to get  $m_1$ . Since  $m_0 = 1 - m_1$ , this meant that the sum of the two masses added up to 1 each time. Since the mass is conserved, such a multiplicative cascade is called *conservative*.

But now imagine if we made the mass allocations independent of each other. This would mean that there would be no issue with getting two instances of  $m_0$  at any new interval — if we are using two intervals, and the first got mass  $m_0$ , the other interval does *not* have to get  $m_1$ . Thus, the sum of masses ( $m_0 + m_1 \dots + m_{b-1}$ ) doesn't have to add up to 1 each time. However, we keep it so that the *expected sum* of the distribution should be 1 (i.e.  $\mathbb{E}[\sum(\varphi_i m_i)] = 1$ ). In other words, the total weight of all masses should add up to one *on average*.

In fact, we don't even need to have discrete probabilities or multipliers (meaning a countable set of possible  $\varphi_i m_i$ ). Instead, we can draw the multipliers from a continuous probability distribution. In this final extension, we replace our discrete multipliers ( $\varphi_i m_i$ ) with some random multiplier  $M_i$ , drawn from a distribution of choice. An important point is to make sure that the *expected sum* of the multipliers equals one ( $\mathbb{E}[\sum(M_i)] = 1$ ).

This second type of multiplicative cascade, in which the masses only add up to 1 in expectation, is called a *canonical* cascade. This is the type most commonly used for simulating the MMAR-model, and it is the type which we will be using. We will also be using a *binomial cascade*, meaning that we will be splitting each interval into two each time (meaning  $b = 2$ ).

As mentioned before, the most common distribution used for these cascades is the *log-normal* distribution, with a mean  $\lambda$  and a standard deviation  $\sigma^2$ . These two parameters are easily determined as follows<sup>30</sup>:

$$\hat{\lambda} = \frac{\hat{\alpha}_0}{\hat{H}}$$

$$\hat{\sigma}^2 = \frac{2(\hat{\lambda} - 1)}{\ln[b]}$$

From this, we get that:

$$V = -\log_b M \sim \mathcal{N}(\hat{\lambda}, \hat{\sigma}^2)$$

and

$$\mathbb{E}[M] = \mathbb{E}[e^{-V \ln b}] = \frac{1}{b}$$

where  $V \sim \mathcal{N}(\hat{\lambda}, \hat{\sigma}^2)$ .

Below is an example of a binomial cascade which can be used for the MMAR-model:




Figure 14 — *The log-normal binomial cascade*. An example of a canonical cascade simulation for Norway, with log-normally distributed mass allocations. Since the parameters are  $k = 13$  and  $b = 2$ , we create  $2^{13} = 8192$  days.

#### **Step 14:**

**Convert the values of the generated cascade  $\mu$  into a cumulative distribution function (a “CDF”) to find the “trading time” function  $\theta(t)$ .**

The trading time CDF is the first main component of the MMAR, with the other being the fractional Brownian motion. This is the component that creates both high kurtosis and volatility clustering for our simulated returns.

Computing the CDF is simple. Basically, compute cumulative sum of the cascade from time 0 to the end of time (i.e.  $T$ ) for each value of  $t$ :

---

<sup>30</sup> For reference, all of these relations are derived in the three Mandelbrot et al (1997) papers.

$$\theta(t) = CDF_{\mu}(t) = \sum_{i=0}^t \mu(i)$$

The diagrams below offer an example. The first graph is a random binomial cascade using the data for Norway, and the second is its CDF. We can see that, when the cascade is tall, the CDF grows quickly, and the opposite is true when it is low.




Figure 15 — **How to construct the trading time component.** Take your simulated lognormal cascade, and calculate the cumulative sum of the “mass” as you move from the beginning (time 0) until the end (time T — here, day 8192). “Trading time”  $\theta(t)$  increases quickly when the cascade is tall.

It’s also useful to then divide everything by the sum on the last day, and multiply by the time period. This normalizes the trading time and makes it easier to interpret — we can see that the cascade has a truly massive range (the highest number is around  $2 \times 10^8$ ). By normalizing, instead of the last day being a random giant number, every “trading time” CDF finishes on 8192.

Since our cascade was made using a random distribution, we should generate a different CDF each time. Furthermore, when the cascade is made with a high variance parameter (a big  $\sigma^2$ ), we can get a wider range of trading time CDF’s. In the demonstration below, I simulate 1,000 trading time CDF’s for each country. Sweden has a lower kurtosis in the real data, so its band of probable CDF’s is narrower than the other two:




Figure 16 — **When time is non-linear.** 1,000 simulations of trading time for each market (normalized). “Trading time” is the result of turning the multiplicative cascade into a CDF. The CDF is normalized to always add up to 1.

The kurtosis of the original data determines the variety of our trading time simulations. Thus, the lines grow more distant from the center — Sweden has the lowest kurtosis and the narrowest band, while the opposite is true for the British LIBOR.

**Step 15:**

**Simulate a Fractional Brownian Motion  $B_H(\cdot)$ .**

Fractional Brownian motion (“fBm”) is the second stochastic component of the MMAR-model. This is the component that allows us to simulate non-random walks — i.e., random walks with Hurst exponents other than 0.5. Simulating the fBm is complicated, but many programming languages include a package for doing so. Here, we will briefly discuss its mathematics.

$B_H(\cdot)$  is a stochastic process characterized by the following equation:

$$B_H(t) = \frac{1}{\Gamma\left(H + \frac{1}{2}\right)} \left( \int_{-\infty}^0 \left[ (t-s)^{H-\frac{1}{2}} - (-s)^{H-\frac{1}{2}} \right] dB(s) + \int_0^t \left[ (t-s)^{H-\frac{1}{2}} \right] dB(s) \right)$$

where:

1.  $\Gamma(z)$  represents the Euler Gamma function  $\int_0^{\infty} [x^{z-1} e^{-x}] dx$
2.  $H$  represents the Hurst exponent,  $0 < H < 1$
3.  $B(s)$  represents a typical stochastic Brownian motion with  $H = 0.5$
4.  $t$  represents time

These equations are used to simulate a random process with a given Hurst exponent — in other words, a non-random walk.

Since I was simulating 8,192 trading days, I generated ten times that number (81,920) of data points for each fBm simulation. The reason for this is that later, when we combine the two processes, this would give plenty of room for the trading time to “jump” from one point on the fBm line to another. The way this works can be illustrated by example: if  $\theta(t) = 1.1$ , we go to  $B_H(\cdot) = 11$ . If  $\theta(t) = 21.3$ , we go to  $B_H(\cdot) = 213$ . If we had only generated 8,192 fBm days, fractional values of  $\theta(t)$  would get rounded — numbers like 1.1 or 1.2 would have to be interpreted as 1 — and this would cause too many days to appear as though there was zero change, because we would be stuck on the same point on the fBm line.




Figure 17 — **The non-random walk.** The fBm is similar to a random-walk except that the changes do not have to be independent, and can instead take on any  $H$  value.

Here, we see an example of a fractional Brownian motion simulation for Norway, with a Hurst exponent  $H = 0.432$ . Note that we purposefully generate 10 times more data-points than the length of our trading time CDF in days.

Some technical details: There are many methods for simulating an fBm process. The best and fastest known method is the one explained in Davies & Harte (1987), which lets us quickly simulate many fBm points for long time periods. To implement this method, I used the “fBm”

package for the Python language. One more important technical issue: simulating the fBm requires specifying the “length” of the fBm process, which determines the magnitude of movements. Consider our graph above: the range is approximately between 0 and -0.4. A simulation with greater “length” would have a larger range of values for  $fBm(t)$  — for example, a range of 10 units between 5 and -5. Consequently, the magnitude of price changes generated by the MMAR would also increase. The authors say nothing about this problem (mostly likely because they had their own way of simulating an fBm). To get around this, I had to manually adjust the fBm “length” for each country until the median standard deviation of returns was approximately the same as that of the real data. I report the lengths I used here:

	Norway (USD/NOK)	Sweden (OMXS30 index)	Britain (12-month LIBOR in £)
fBm length	0.10	3.24	4.0

*Table 4 — A technical detail for the fBm. This is not important conceptually. This is merely an input I had to manually adjust in the fBm simulation process. “Length” affects the magnitude of movements in the fBm path, and thus also the simulated  $X_t$ ’s. I adjusted it manually until the median standard deviation of the MMAR simulations was the same as the real standard deviation.*

### Step 16:

#### **Merge the trading time CDF with the fractional Brownian motion for a simulation of overall returns ( $X_t$ ).**

The final step of an MMAR simulation is to merge the two random components that we simulated in the previous two steps. This is meant to represent how much a price has grown since the beginning of the time period.

Formally, the MMAR-model is defined using the following formula:

$$X_t = \ln P(t) - \ln P(0) = B_H[\theta(t)]$$

Where  $B_H$  and  $\theta(t)$  are independent.

To explain this idea, I show a 27-day MMAR simulation in the diagram below, using a simplified version of the “trading time” CDF with only three speeds for time — normal, slow and fast. We take our data-points from the orange fBm line, and the ones we take are determined by our trading time CDF. Basically, the red CDF tells us how far to jump on our orange fBm line. When time is “slow”, i.e. trading time increases slowly, then we take small steps forward on our fBm simulation. Consequently, the box from which we take our data-points on the orange line is narrower when time is slow than when it is at normal speed. Similarly, when time is “fast”, meaning our CDF is increasing rapidly, then the steps between our fBm points increase as well.

This is why the third box is the widest. In the end, the data-points that we chose on the fBm are interpreted as price growth over time.

Because the jumps are small when time is slow, the resulting  $X_t$ -curve has low volatility when it's inside the second box. We can interpret this as the slow, boring days of the market when not much is happening. But when time is fast, then prices move quickly and drastically, and volatility goes up. These fast times are the days when fortunes are made and lost.




Figure 18 — **The MMAR explained.** The MMAR model simulates total growth for some time period (here, total growth at daily frequency for 27 days). In the MMAR, growth  $X_t$  is a compound process of two components: "trading time"  $\theta(t)$  and a fractional Brownian motion  $B_H[\cdot]$ .

In the trading time CDF (left graph), we imagine that time can move at different speeds. In this example, we simplify to only have three speeds: the usual increments of 1 day (blue area); slow increments of 0.5 days (orange area), and fast increments of 1.5 days (green area). Next, we simulate an fBm process (the  $B_H[\cdot]$  function). We take the  $B_H[\cdot]$  value for every  $\theta(t)$  value (the red crosses), and the result is an MMAR simulation of price growth:  $X_t = B_H[\theta(t)]$ , where  $X_t \equiv \ln \left[ \frac{P(t)}{P(0)} \right]$ .

Because of this jumpy nature of the model, we expect low volatility when time is slow, and high volatility when time is fast. Meanwhile, the fBm lets us generate non-independent price changes.

We've come a long way and seen that MMAR simulations are very complicated to construct. Unfortunately, the authors don't offer any code for making one. But there is good news: Wengert (2010) offers a robust version of MMAR code for users of MATLAB software. I modified this code for Python when constructing my multiplicative cascade, but otherwise wrote the full code on my own.

### Step 17:

**(Optional) Before simulating, add any desired restrictions or modifications to the outputs.**

To make the analysis a little easier, I also added some restrictions to the models. This was to combat some unrealistic issues arising just out of pure randomness. For instance, some simulations had the Swedish stock market increasing to a price of over 100,000 (this would've been equivalent to an annual growth rate around 25%), when the real end price was around 1,500. Others saw the LIBOR-rate shooting up to over 100%.

I assumed that any practitioner would just delete these wild simulations and run some more instead. So rather than doing so myself, I simply imposed some arbitrary restrictions on the outputs of both the MMAR and the Gaussian simulations. Specifically, I made it so the price of the Swedish index could never go over 5,000 and never below 100. A final day at 5,000 would represent approximately a 12.5% average annual growth rate over 32 years, which I felt was a reasonable upper bound for a stock market. I also made it so that the LIBOR could never go over 20% and never below 0.1%. In the past 30 years, the LIBOR has never gone over 16%. Finally, I also added the daily average price movement as a bias into the model, so that the model would move in that direction on average (we will later see that this made the mean daily price change of the simulations to usually be approximately the same as the real one.) The simulations for Norway were left untouched with no restrictions, because in general, for rich countries, we don't expect there to be a long-term trend in the exchange rate in either direction.

These restrictions are summarized below:

Restriction / Market	Norway (USD/NOK)	Sweden (OMXS30 index)	Britain (12-month LIBOR in £)
Price path restriction	None	$100 < P(t) < 5,000$	$0.1\% < P(t) < 20\%$
Daily price offset	None	+0.000325	-0.000322

*Table 5 — Restrictions added to the model. I purposefully restricted the model's simulations to never go out of the specified price ranges at any time  $t$ . I also made it so each day the price movement would be biased in some direction. The bias equals the average daily price change. Norway's simulations were left without restrictions.*

#### 4.4. Summary of Methodology

- Simulating markets with the MMAR-model is complicated business.
- The first step is to find four empirical parameters —  $H, \alpha, \lambda$  and  $\sigma^2$ . This is done by calculating some technical features the data — namely, its partition function ( $S_q(T, \Delta t)$ ), scaling function ( $\tau(q)$ ) and multifractal spectrum ( $f(\alpha)$ ).
- The next step is to simulate two random processes, which are used as the two components of the MMAR. These are: 1) the “trading time” function ( $\theta(t)$ ), made by converting a lognormal multiplicative cascade into a cumulative distribution; and 2) a fractional Brownian motion ( $B_H(\cdot)$ ), which is like a random-walk except that changes don't have to be independent.
- Trading time is plugged into the fBM, and we can interpret the compound process as a simulation of price growth over time ( $X_t = B_H[\theta(t)]$ ).

## V. Results

### 5.1. Empirical measurements

Unfortunately, the MMAR's authors do not provide any clear way to calculate an error range for these figures. But I suspect it doesn't really matter. The values only have to be close enough to generate reasonably realistic simulations.

The following table shows the values I found for the four MMAR parameters.

	$H$	$\alpha_0$	$\lambda$	$\sigma^2$
Norway (USD/NOK)	0.432	0.483	1.118	0.341
Sweden (OMXS30 index)	0.542	0.590	1.090	0.258
Britain (12-month LIBOR in £)	0.645	0.726	1.126	0.364

Table 6 — *The most important table.* The four MMAR parameters for each market — all we need to run 10,000 simulations.

The  $H$ -exponent values for Sweden and the LIBOR are consistent with theory — namely that many financial series show persistent behaviour in price returns. Their Hurst values are greater than the random-walk value of 0.5, which suggests that their time series are more likely to show trends. Up movements are more likely to be followed by up movements, and the converse is true for down movements.

Interestingly, the Norwegian krone seems to have an anti-persistent Hurst exponent (recall that this suggests that an up-movement is more likely to be followed by a down-movement). Some small anti-persistence is even given by the most common Hölder exponent ( $\alpha_0$ ) of the multifractal spectrum. Anti-persistence seems to be a common feature of currency markets, as I found the same result for several other countries<sup>31</sup>. We can interpret this as suggesting that market forces tend to keep the exchange rate at the same level as it was before — if the price moves too far in any direction on a given day, then it tends to go in the opposite direction on the next day, and any trends get arbitrated away. In some sense, then, the exchange rate tends to fluctuate around an equilibrium price. In one way, this makes sense. We typically don't expect rich countries to show any trend in the foreign exchange rate, and in fact we might

---

<sup>31</sup> See [appendix 7](#). Interestingly, this conflicts with Peters (1994), who found that currencies showed no long-term cycles and were thus completely describable by fractional noise processes (meaning  $H$  should be close to 0.5), whereas stocks and bonds were akin to fractional noise in the short term and chaotic in the long term.

expect exchange rates to be quite stable. (For example, 100 years ago, the British-Pound / US-dollar exchange rate, was almost the same at about 3.5£/\$.) For stock markets, on the other hand, we expect there to be growth trends as the country's economy grows, and perhaps even some bubble trends. This is particularly true for stocks, because their fair value is perhaps harder to define than that of exchange rates, which are often explicitly defined by central bank targets.

We arrived at these four parameters by the process outlined below. First, we found the following partition functions<sup>32</sup>:




Figure 19 — *Is the data fractal?* The partition functions for all three markets seem to show scaling behaviour. This suggests that the price movements are fractal. These graphs have also been normalized to start at zero to make the scaling behaviour easier to see.

By running linear regressions on the partition function for each raw moment  $q$ , we can determine the scaling behaviour of each market. For every market, we observe a nonlinear scaling function  $\tau(q)$ , which we can interpret as evidence of multifractality. Multifractal scaling is consistent with the MMAR's assumptions.




Figure 20 — *Is the data MULTI-fractal?* Here we see that the scaling function  $\tau(q)$  is non-linear, which suggests that, yes, it is. The scaling function describes the relationship between the partition function ( $S_q(T, \Delta t)$ ) and its time increments ( $\Delta t$ ) with respect to the raw moment ( $q$ ) This nonlinearity is the most common symptom of multifractality.

<sup>32</sup> Note that here, I only show the partition functions for select moments. Having used 121  $q$ 's, the full partition function graphs have 121 lines each, and thus would look like a mess. Nevertheless, I include the full graphs in [appendix 6](#).

Using these empirical scaling functions, we can get a generally good idea of what we should expect for our Hurst coefficients. But before we find those, we estimate the multifractal spectra — the  $f(\alpha)$  functions:




Figure 21 — **What is  $\alpha_0$ ?** The multifractal spectrum determines  $\alpha_0$ , one of the four parameters of the MMAR.

The above multifractal spectra were used to find our  $\alpha_0$  coefficients, which is just the point where the  $f(\alpha) = 1$ . They were estimated by performing nonlinear regression on the empirical  $\tau(q)$  scaling function to estimate some parameters for a continuous function. and then performing a Legendre transform on the result.

Below are the estimated scaling functions, which we can use to find our Hurst coefficients.




Figure 22 — **What is H?** Here, we re-estimate the scaling function to have it in its continuous form and use it to determine H, another key parameter in the MMAR. Estimating the function is not a problem, as it passes very closely through each real value.

Below is a table summarizing the parameters for our estimated  $\tau(q)$  and our  $f(\alpha)$  functions:

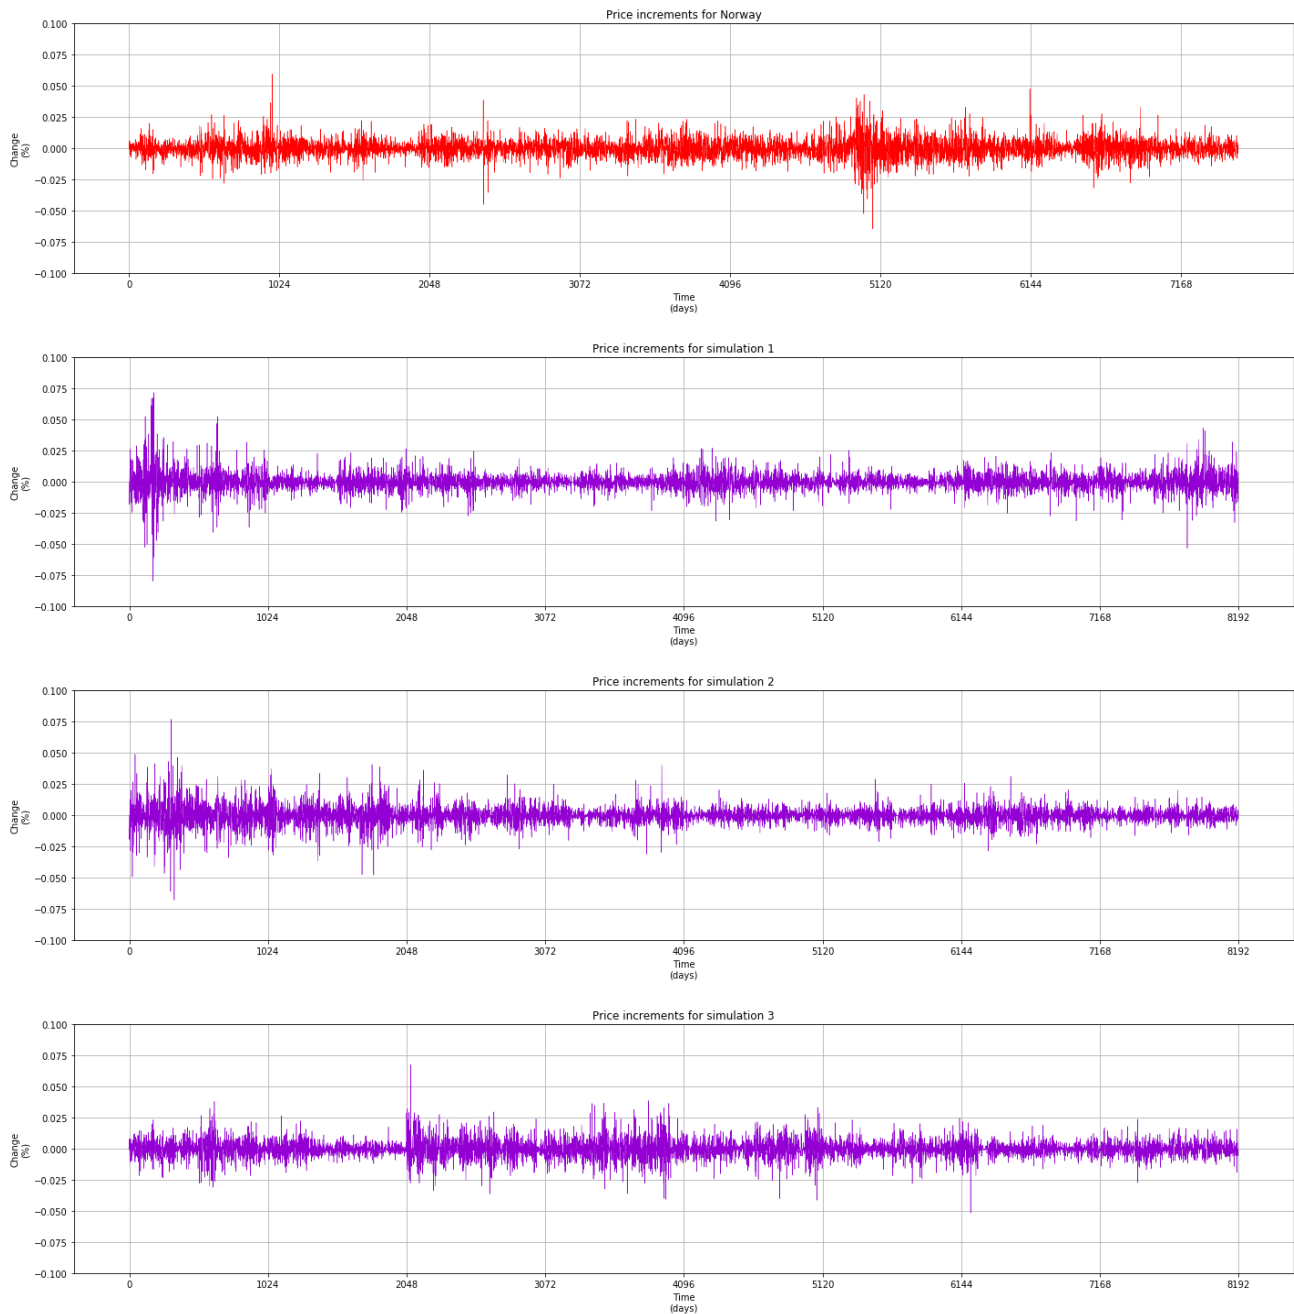
	Estimated $\tau(q)$	$f(\alpha)$
Norway (USD/NOK)	$-0.0212q^2 + 0.484q - 1.01$	$-11.79\alpha^2 + 11.41\alpha - 1.75$
Sweden (OMXS30 index)	$-0.0269q^2 + 0.591q - 0.999$	$-9.30\alpha^2 + 10.99\alpha - 2.25$
Britain (12-month LIBOR in £)	$-0.0544q^2 + 0.728q - 0.997$	$-4.59\alpha^2 + 6.69\alpha - 1.44$

Table 7 — **Continuous functions are easier to work with.** Here I show the formulae for the multifractal spectra and the estimated scaling functions, which are shown above. These were estimated using moments up to  $q = 4.0$ .

We are now ready to consider what type of simulations these parameters give us. It is helpful to consider them *visually* before we move on to studying them *statistically*. In the next

few sections, for each country, I show the real daily price changes followed by the first five simulations. I also show the 10,000 simulated price paths that the MMAR gave me, as compared to the real price path over thirty years.

## 5.2. Results from Norwegian currency market






Figure 23 — *How do prices move?* The actual price changes for Norway (red) and five MMAR simulations (purple).






Figure 24 — **10,000 universes.** 10,000 MMAR price paths for the USD/NOK exchange rate. The black line shows how the price actually moved in real life.

### 5.3. Results from Swedish stock markets






Figure 25 — *How do prices move?* The actual price changes for Sweden (gold) and five MMAR simulations (orange).





Figure 26 — **10,000 universes.** 10,000 MMAR price paths for the Swedish stock market index. Note that these price paths have been restricted never to go over 5000 nor under 100. The black line shows how the price actually moved in real life.

### 5.4. Results from British LIBOR market



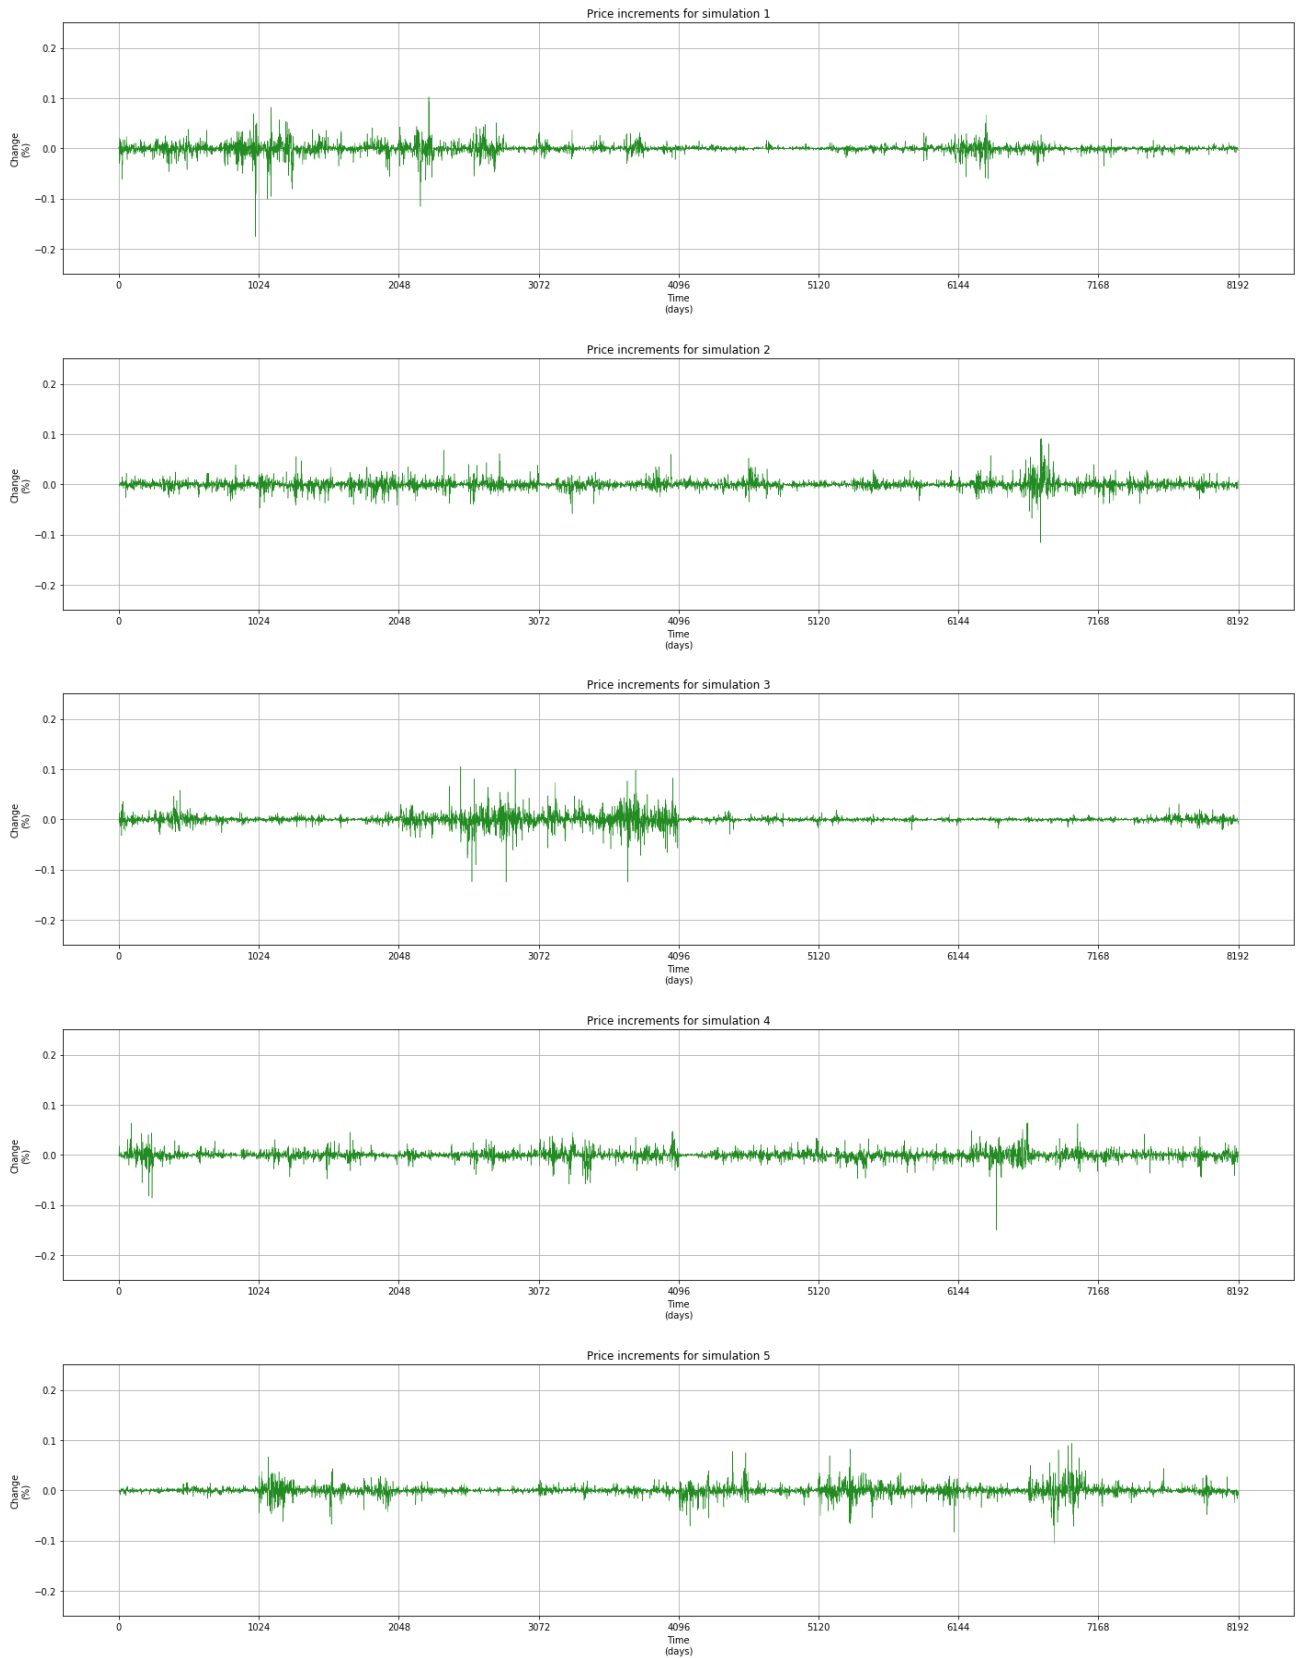


Figure 27 — *How do prices move?* The actual price changes for Britain (blue) and five MMAR simulations (green).




Figure 28 — **10,000 universes**. 10,000 MMAR price paths for the 12-month LIBOR in British pounds. As with the Swedish simulation, these price paths have been restricted never to exceed a yield of 20% nor to drop under 0.1%. The black line shows how the LIBOR-rate actually moved in real life.

### 5.5. Analysis of simulated moments (final price, mean, standard deviation and kurtosis)

From the graphs above, we can see that the MMAR achieves its three main goals. The graphs show that it can simulate both high kurtosis and volatility clustering. We also know that it simulates non-independence, because it is made using an fBm with an  $H$  exponent that we choose in advance.

We can begin our analysis of the MMAR's performance by considering the histogram graphs of the simulations. Loosely speaking, for each simulation, here I show the results for:

- The zeroth moment — where the price ended up after 32 years.
- The first moment — the mean of price changes.
- The second moment — the standard deviation of price changes.
- The fourth moment — the kurtosis of price changes.

Note the *smoothness* of the distributions. This smoothness suggests that the MMAR can be used to make *probabilistic forecasts* about the future.

The graphs are also reasonably symmetrical. Because prices can grow more than they can contract, we would expect the final price curve to be skewed to the right — there should be a long right tail. Excess kurtosis (which can't go below -3) is also skewed. But we would also expect the distribution of its higher moments to be a symmetrical bell curve — if we had an infinite

number of simulations, the distribution of their means and that of their standard deviations should be a smooth curve, with a peak in the middle and smoothly decreasing tails. In fact, this is precisely what we see for Norway. In fact, we would have probably seen the same thing for the means of the two other countries, had I not added restrictions to the model to prevent prices from flying off to extremely high amounts. Nevertheless, the histogram of the standard deviations appears to be converging to a clear, smooth bell curve.




Figure 29 — *What can we expect in thirty years?* Histograms of last-day prices for each MMAR simulation.

I can report that the  $H$ -parameter can greatly affect price paths in the long run. With a high  $H$ , *trends get amplified*, so many simulations can fly off to astronomically high amounts (this is one of the main reasons why I decided to add restrictions to the price range of simulations). On the other hand, with a low  $H$ , *trends get dampened*.

For example, though I do not show the graphs here, I can report one basic comparison to the random walk. Since Norway had a low  $H$ -exponent (meaning that up and down movements tended to reverse) it actually had a much lower range of final prices than the Gaussian random walk — the MMAR's highest price was about 30, while the random walk's was around 140.<sup>33</sup>

For the mean and standard deviation, I show the Gaussian random walk comparison in an overlaid histogram in light blue:

<sup>33</sup> Some Swedish stock market simulations (without restrictions) flew off into the millions — compare that to its real final-day value which was around 1,500. Clearly, a high  $H$  exponent can strongly amplify any trend.




Figure 30 — **The average price change for every simulation.** These histograms show the mean price change over 8,192 simulated days for each MMAR simulation, with Gaussian comparisons in light blue. We can see that a low  $H$  contracts the range of means, while a high  $H$  widens it. For Norway, the most common mean was zero, which was done on purpose. Similarly, I purposefully sought to make the expected means of the other two markets equal their real mean returns.

We can see that the  $H$ -exponent also affects the *mean returns*. For Norway, because of the low  $H$ , the range of expected mean returns was much lower for the MMAR than for the random walk — the MMAR’s mean returns had a range of about -0.00015 to +0.00015, while the random walk had one about twice as wide from -0.0003 to +0.0003. The range became twice as narrow when going from the random walk  $H$  of 0.5 to the anti-persistent 0.432. The other two markets had a high  $H$ , so their range of mean returns would have been much wider — had I not added restrictions. Because of the restrictions, the histogram of mean returns for Sweden is visually nearly the same. However, for the LIBOR, the wider range is clearly visible even with restrictions.

So the  $H$  exponent tells us something important about what we might expect in the future. According to the MMAR, we might conclude that the  $H$ -value affects our *range* of expected future returns — a high  $H$  makes both very high returns and very low returns more probable, whereas a low  $H$  suggests that the mean return will probably change very little. And as for the price itself, a high  $H$  makes the long run more unpredictable *despite* having more trends in the short run. Meanwhile, a low  $H$  suggests that prices will probably not change very much in the long run, because short run trends will tend to cancel out quickly. And furthermore, this is consistent with the assumption that returns are random.


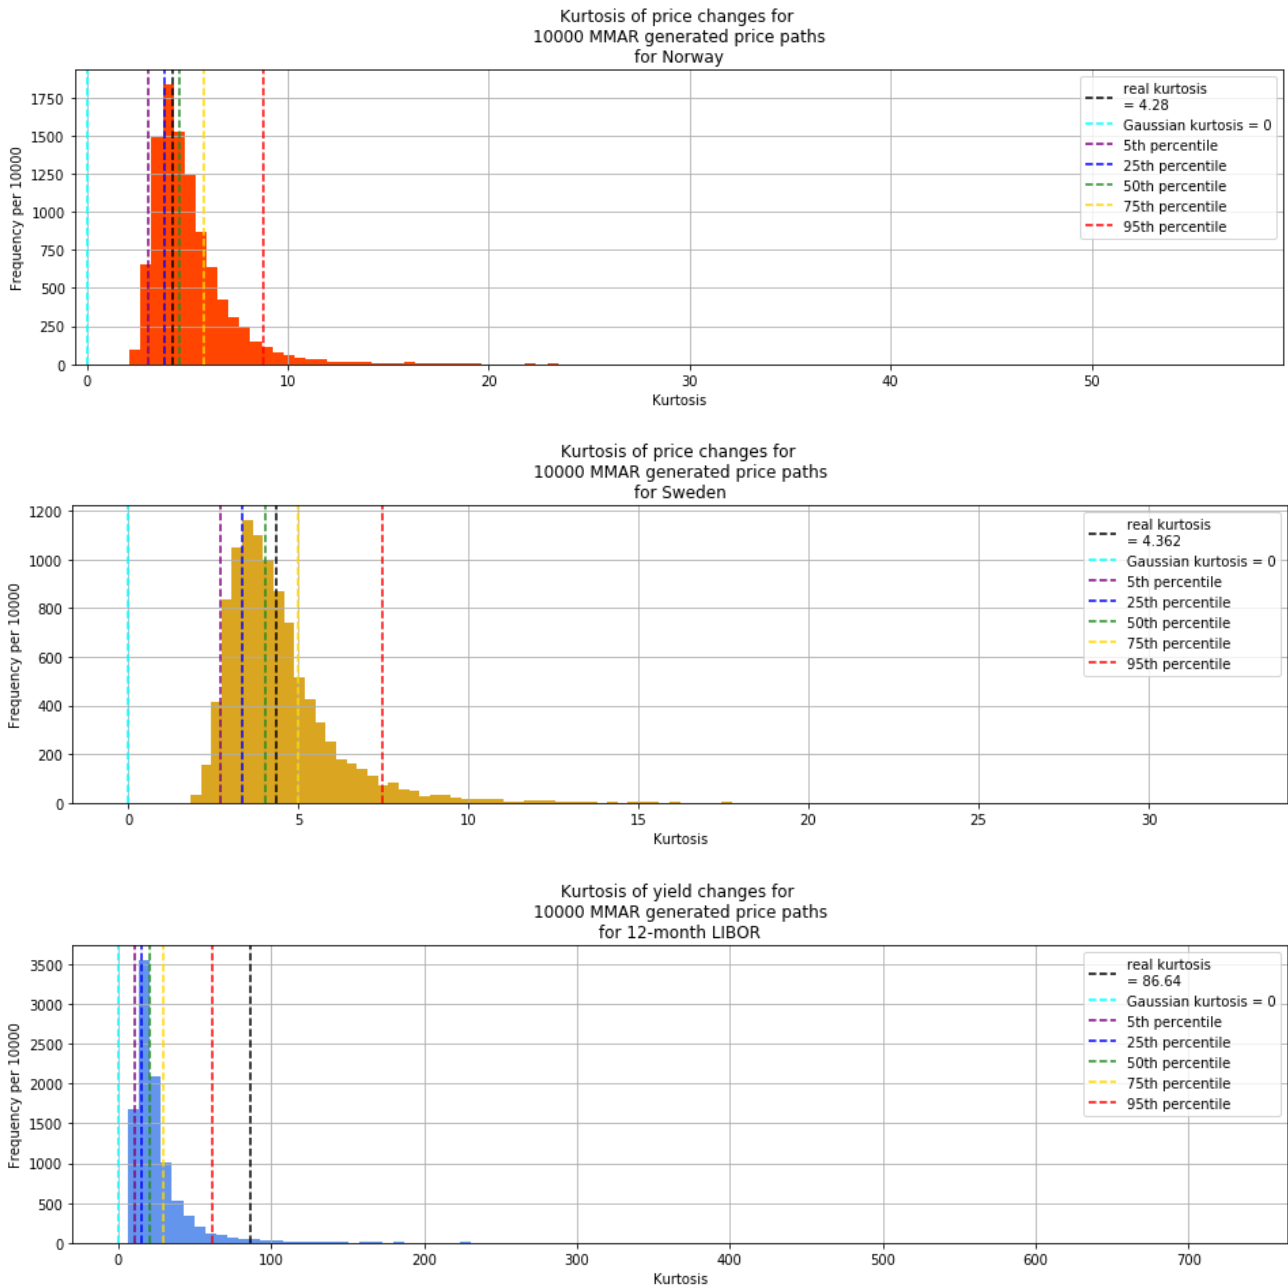


Figure 31 — **The mean is nothing without the variance.** The standard deviations of our simulated returns. We can see that, on average, the variance of the MMAR simulation is more variable than for the Gaussian simulations.

We can also see that the range of probable variances increases when using the MMAR. Interestingly, the range of standard deviations seems to increase with the kurtosis of the data — Sweden had the lowest kurtosis, and here the difference with the Gaussian is the smallest, whereas the LIBOR had the highest kurtosis for both the data and the simulations, so the difference between the MMAR and the Gaussian simulations is much clearer.

What about kurtosis? What can we expect from the MMAR?



*Figure 32 — Consider the odds of catastrophe — the kurtosis for every MMAR simulation. We can see that, in general, the MMAR tends to produce kurtosis levels that are smaller than that of the real data (the black dashed lines). However, it practically never gets close to the Gaussian kurtosis of zero. Data with higher kurtosis produces a wider variety of simulations. The British LIBOR, for example, produced one monstrous price path with a kurtosis over 700!*

In these graphs, the black dashed line represents the kurtosis of the real data, while the light-blue line represents the theoretical Gaussian kurtosis of zero. All of our 30,000 MMAR simulations are leptokurtic — they have excess kurtosis compared to the Normal distribution (none of them showed a kurtosis of zero). In other words, all the distributions of returns have a tall peak and fat tails. This is similar to what we find in the real data. (It also seems to imply that, for these markets, the odds of getting a price path with Gaussian kurtosis are practically zero).

We can see that the MMAR can produce simulations with quite a large variety of kurtosis values, and these appear to show a smooth progression. This range can be very large, and increases with the kurtosis of the original data — witness the LIBOR simulation with kurtosis over 700. However — for data sets with very high kurtosis to begin with — it appears that the simulated kurtosis is usually lower than that of the real data. For the LIBOR, the real kurtosis (black dashed line) was higher than the 95<sup>th</sup> percentile simulation (red line). The implication is that price paths with high kurtoses are improbable, but possible. This “under-simulation” problem doesn’t seem to happen for datasets with smaller kurtosis — for Sweden and Norway, the 50<sup>th</sup> percentile of simulations (green line) was very close to the real data.

What about the Gaussian random walk? What type of kurtosis did it show?




Figure 33 — *Gaussian excess kurtosis equals zero.* We can see that using the Normal distribution yields simulations that show practically no excess kurtosis.

For all three countries, the kurtosis of the random-walk simulation never went out of the range of -0.2 and 0.2. Thus, normally distributed simulations tend to stay quite close to the theoretical excess kurtosis of zero. This is far closer than the real data. So the basic Gaussian simulations suggest that high-kurtosis events are, for practical purposes, impossible.

Based on kurtosis, then, the MMAR fits the data much better than a basic random walk. This is not much of a surprise. The next question is — how realistic are the MMAR’s simulations?


## 5.6. [Are the simulations realistic? Kolmogorov-Smirnov testing](#)

Interestingly, the three authors spend a great deal of time discussing the mathematical derivations of the MMAR and its advantages over typical models, but they never provide any good way to test its predictions statistically. To check whether the simulations are realistic, nothing is offered beyond the eye-ball test — the only argument seems to be that if you plot the returns and multifractal spectrum of the real data, then generate some simulations, then the MMAR's graphs should look more like that of the real data than do graphs of random-walk models or GARCH. (To be fair, I admit that the simulations do at least pass my eye-ball test).

Is it a problem that the model is not really statistically testable? Physicists have a designated phrase for models that are not testable — such models are often known as “not even wrong”. But Mandelbrot's (2004) writing seems to suggest that he merely wished the model to be “good enough” to make financial decisions, rather than fitting the data perfectly.

The only method I could find for comparing two statistical distributions was the Kolmogorov-Smirnov test, which can be used to test whether two arrays of data could have been drawn from the same statistical distribution. The KS-test works by comparing the cumulative distributions of two datasets. It is non-parametric, and does not assume any type of distribution beforehand. The null hypothesis is that the two sets of data are generated in the same way, so if the P-value is low, we reject the null hypothesis and conclude that the two arrays are drawn from different distributions. Using the KS-test, I compared the returns of each MMAR simulation to the returns of the real data for that market.

How well does the MMAR perform on the KS-test? Well, here there's bad news — it seems that most simulations fail. At the 5% significance level, over half of the simulations for both Norway and Sweden failed the test. Worse, for the LIBOR, not a even single simulation could manage to get a P-value that wasn't a flat-out zero. The sorted P-values for every KS-test are shown in the diagrams below:



*Figure 34 — Failing the test for realism. These graphs show the P-Values of the Kolmogorov-Smirnov test for every MMAR simulation, where we compared the distribution of returns to those of the real data. We can see that most simulations fail the KS-test — with 95% confidence, we reject the null hypothesis that the simulated returns are drawn from the same distribution as the real returns. In particular, not a single LIBOR simulation passed the test.*

*However, I suspect the MMAR's creators probably wouldn't be bothered by this result. The model is probably designed on purpose in a way that would make it fail most KS-tests. The feature is built-in — like this, we can simulate more possible price histories. The failures most likely arise from the many different kurtosis values. Unfortunately, however, the MMAR's creators never specified a way to test the model's realism.*

It also appears that the model fails the KS-test mainly because of different kurtosis values. I tested the sensitivity of the KS-test by comparing Normally distributed arrays with different values for the mean and variance, resembling the ranges found in the simulations, and found little change in the number of failing P-values.

So kurtosis is the most likely culprit for the testing failures — but I suspect that the MMAR's creators wanted the model to produce a variety of kurtoses. Massive one-day price shifts that led to one-day bankruptcies seemed to be the main thing they were warning about. So failing the KS-test appears to be a built-in feature of the model — since it's not constrained by the test, it can suggest a greater variety of possible price histories. Practically, this greater variety of simulations should let us prepare better for many different eventualities.

Is there a way to make the MMAR perform better on the KS-test? In [section 3.3](#) of the literature review, I mentioned that Lux (2001) was able to get a perfect success rate for the KS-test and his MMAR simulations. He did this by manipulating the parameters of his multiplicative cascade — specifically, the second-last MMAR parameter,  $\lambda$  — until every simulation had a P-value that was higher 0.05. His paper argued that it was “remarkable” that a stochastic model could get such a success rate by only manipulating one parameter. One might consider this to be an advanced form of P-hacking, and I find it dubious — what is the point of doing empirical work, if in the end we just optimize our parameters anyway? Regardless, I was unable to replicate this result.

Meanwhile, every single Gaussian simulation failed the KS-test, with a flat P-Value of 0.0. Thus, the random walk performed even worse than the MMAR. So although the MMAR performs badly on the test, we can see that it could be worse.

### **5.7. Results from a basic random walk model**

Before we look at some “monster” simulations in the next section, let us take a look at what we will be comparing them to. Every random walk simulation looks basically like this:




Figure 35 — **Gaussian returns**. Simulated Normally distributed returns for Norway (purple), Sweden (orange) and Britain (green). These look obviously different from the real data.

Although they all look basically the same, these graphs have the great virtue of being very quick and easy to generate. Where the simulation time for the MMAR took me somewhere between 8 to 20 hours (depending on restrictions), these basic ones never took longer than 20 minutes. And of course, they are also easier to understand, and much easier to code.

### 5.8. Monsters — the simulations with the highest kurtosis

The MMAR-model is supposed to be able to suggest a possibility for very high-sigma events. Here I show the simulations that the MMAR gave me that reached the highest kurtosis values. For convenience, I shall refer to them as “monsters”.

These monster graphs look perhaps too extreme. The monster kurtosis values were 57 for Norway, 32 for Sweden, and an enormous 729 for the British LIBOR. This is approximately 10 times higher than the kurtosis of the real data.

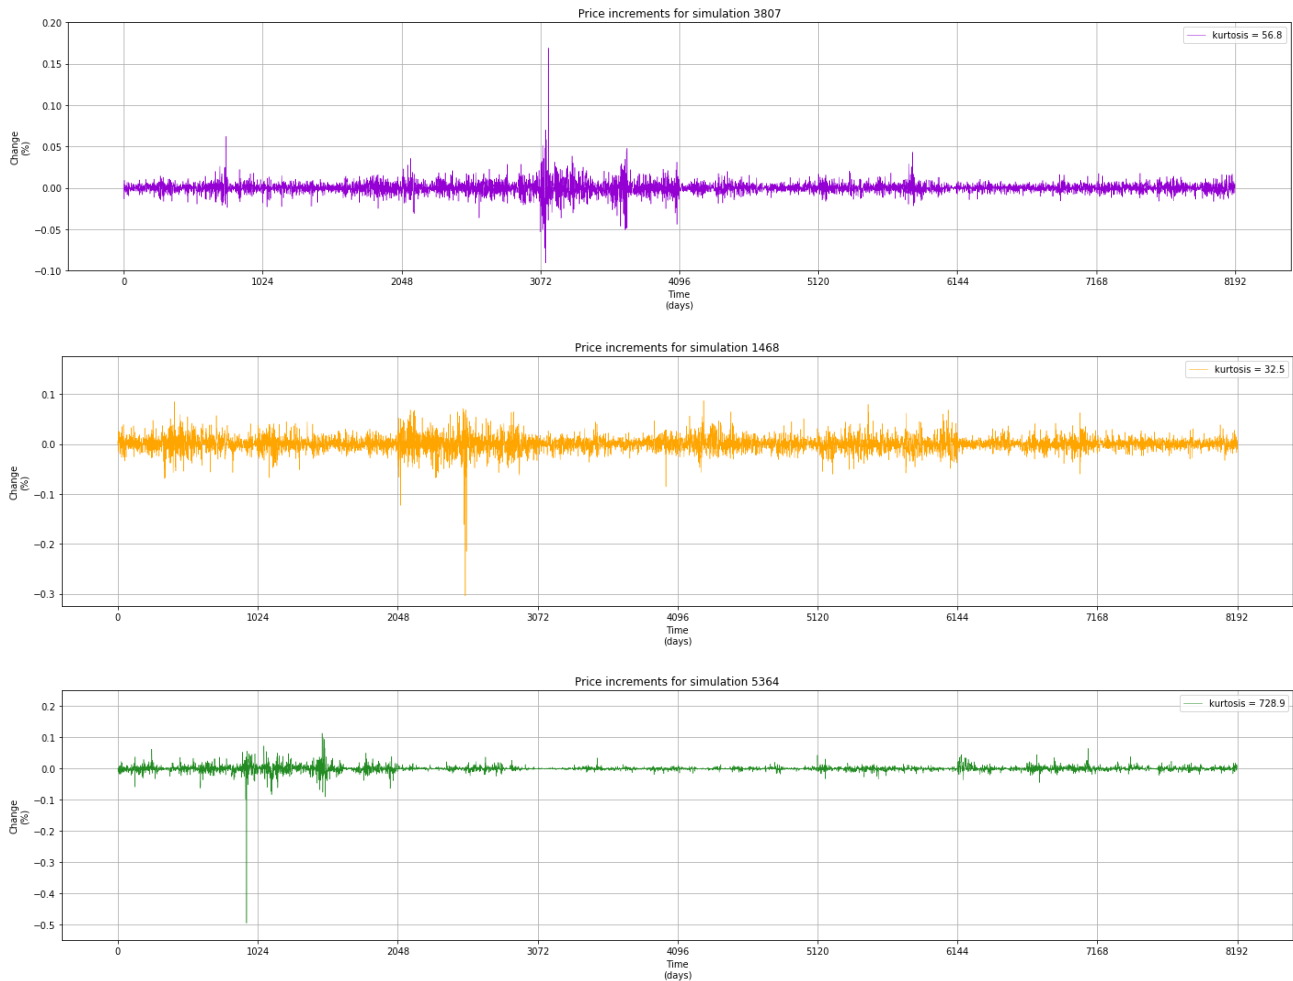


Figure 36 — *Now that's a day for the history books.* The MMAR simulations with the highest kurtosis for each market. As before, Norway is purple, Sweden is orange and the British LIBOR is green.

How wild is a kurtosis of 729? The 50% move downwards is approximately 60 standard deviations from the average<sup>34</sup>. To put it into perspective, the average height for men in the United States is about 176 centimetres, with a standard deviation around 7.42 centimetres. So at 60 standard deviations from the average, a man would be about 621cm tall, which is about the height of a typical two-story house.

These simulations portray a very unlikely event. One that — according to the MMAR — is so unlikely that it should only occur once in 10,000 thirty-two-year periods. This is so unlikely that it is clearly not possible to test empirically.

But I suspect that getting the probability correct is not the point — the real virtue of the MMAR is that it can generate these extreme simulations in the first place. Presumably, if we consider the model to be realistic enough, then we can use these simulations to stress-test our

---

<sup>34</sup> Calculated using the original LIBOR data, with a mean of -0.0003 and a standard deviation of 0.008319.

portfolios for the worst possible outcome. We can test what it would be like to invest in the unluckiest of 10,000 universes, and how that would affect us, and whether we would go bankrupt. In the next (and final) section of the results, I show a very simple way in which we might use the MMAR to stress-test a financial portfolio.

### **5.9. Using the MMAR, what can we say about how to avoid bankruptcy? — A very basic model of payoffs from trading**

All of this modelling means nothing unless we can also say something about how these wild times would affect our portfolios. In this section, I will construct two payoff models with some very simple rules, in order to try to compare the payoffs of investing in these imaginary markets to the real market's price path. These are not at all meant to be rigorous, so we can refer to them as our two little "games".

Here I must first point out a very happy coincidence. If we consider the general shape of the price path for our high kurtosis simulation (see below), then we can find something that resembles a financial bubble. We can see that it showed rapid growth in the first few years, before experiencing a rapid collapse. On one day, the price dropped by about 30.4%, which was closely followed by another 20% drop. After that crash, the imaginary market barely managed to get to the same level over two decades later.

In fact, it just so happens that this is quite similar to what happened in the real Swedish stock market! The OMXS30 had a peak at around 1600 SEK in the early 2000's, then dropped rapidly, and only recovered to the same level around 15 years later (I presume this was Sweden's experience of the "dot-com" stock bubble). The real Swedish market dropped much more slowly than our monster simulation, and it never went as high — but nevertheless, the similarity of the overall shape is helpful for us to compare the "monster" to the real market.




Figure 37 — ***This too shall pass***. The price path of the MMAR simulation for Sweden with the highest kurtosis. We can call it the “monster”. Notice the “bubble” and its massive drop occurring around the 2500th day. We shall use this price chart to study the odds of bankruptcy for different levels of leverage.

Since I also ran 10,000 random-walk simulations, I also happen to have a Gaussian simulation that looks very similar to both of these. This Gaussian simulation also had something resembling a bubble — but this bubble didn’t burst rapidly. Instead, it grew slowly and then deflated slowly over many years, before finally starting to grow again.




Figure 38 — ***When bubbles don't exist***. We will compare the above MMAR simulation with this Gaussian simulation. This price path was chosen mainly because it looks similar to the MMAR one, and ends up in nearly the same place.

Both the “monster” price path and the Gaussian price path above end up in the same place, so we can use these two simulations to compare their payoffs under MMAR assumptions versus random-walk assumptions.

For the first game, let us imagine a group of traders, who started trading in the Swedish stock market around 30 years ago. Each one had a different risk tolerance borrowed a different amount of money for their investments. How would their portfolio look today?


We will be studying the cumulative payoff of each trader, while varying their “leverage multipliers”. The idea is similar to the models found in the book by Embrechts et al (1997), who study the influence of extreme events on insurance companies. The question we are asking is: “how badly might leverage affect us under the MMAR’s assumptions?” Presumably, high leverage

makes us gain fast, but also lose fast, so we imagine that leverage multiplies any return for a given day. For simplicity, let us also imagine that each trader reinvests all his money every day. Thus, for each trader, the cumulative payoff for any given day will be determined by their cumulative payoff from the days before, their leverage, and the market's return for that day:

$$U_t = U_{t-1} \times (1 + r_t \times L)$$

Where  $U_t$  is the accumulated payoff by day  $t$  for the particular investor,  $r_t$  is the market return for that day, and  $L$  is the leverage multiplier. All traders start with their assets at  $U_0 = 1$ .

Let us define a profit of zero or less as “bankruptcy”. In the graphs below, I used the payoff equation above to compare how quickly the various traders would go “bankrupt”, if every day their returns were multiplied depending on their leverage. The range of “players” is from 1 (no leverage) to 40 (extremely high leverage — this one usually gets wiped out immediately). On the left I show the formula applied to the real Swedish stock market over thirty years. In the middle I show the milder Gaussian results, while on the right we can see that dramatic results based on our “monster” simulation.



*Figure 39 — Losing, fast and slow.* Here, we compare what how leverage multiplies both wins and losses, for the real Swedish stock market (left), our Gaussian simulation (middle) and our MMAR simulation (right). The results suggest that high leverage eventually ruins all debtors. But high kurtosis, dependence and volatility clustering make it possible to lose more in one day than we ever earned. The massive move in the MMAR is 20 standard deviations, which is still less than the 25-sigma events reported by Goldman Sachs in 2008.

The graphs here are in “symlog” form, meaning they are logarithmically scaled both for positive and negative values. (Normally, logarithms for values less than 0 are undefined.)

It seems that one can do quite well with leverage, but only in the short run. (Which is more or less what we expect to see for most real investors). In each simulation, the cumulative losses in market downturns eventually bankrupted every investor with leverage greater than 3.

What is interesting is the speed of bankruptcy. In all three graphs, we can see that many high-leverage “traders” in our game were quite successful until the downturn. In the hypothetical world of the Gaussian simulation, one could presumably trade for a long time, and then get out of the market when the downturn starts — “taking their winnings” and going home,

to put it in gambling terms. This is not the case with the monster simulation (third diagram). There, the red high-leverage trader lost over 90% of his earnings within a few days (notice that these are logarithmically-scaled diagrams). Worse, a few days later, the “monster” move of the market threw him into negative-payoff territory, where he owed 10 times more money than he had started with. In other words, within a few days, he lost more money than he ever made. (The imaginary traders with even higher leverage were almost all wiped out before the our simulated crash, so we can mostly ignore them.)

With these graphs, we are starting to see what Mandelbrot meant when he argued that many financial models underestimate risks, and offered fractals as an alternative. Volatility clustering coupled with dependent movements can lead to prices falling downwards very fast. This is true even if we ignore extreme-kurtosis events — in the left graph, where we apply the simple payoff model to the real Swedish stock market, high-leverage traders saw very rapid reversals of fortune as well, when prices appeared to only go downward.

From the MMAR’s three features, volatility clustering and dependent moves are the difficult phenomena to study — but fortunately we can still try to contemplate the impact of high-kurtosis events on their own. To do this, we can construct another simple game, where we imagine that traders employ some form of *hedging* — but where the hedge can only protect them *up to a point*. To give the game some basis in reality, we might imagine the break-point representing some form of “counter-party risk” — an extreme event makes it so the counterparty can’t pay us back, and so we are dragged down into bankruptcy along with them.


Let us imagine that our traders are perfectly hedged, and earn the market return every day, plus some small variation (we ignore their leverage levels for now). However, the hedge “breaks” if we experience a high-sigma event. At that point, the leverage multiplies the return of that day. Since the “monster”-event in our simulation was about 21 standard deviations (i.e.  $21\sigma$ ), let us choose  $20\sigma$  as the breakpoint. As a formula, we might write down the rules of the game as follows:

$$U_t = \begin{cases} U_{t-1} \times (1 + r_t + \varepsilon_t) & \text{for } r_t < 20\sigma \\ U_{t-1} \times (1 + r_t \times L) & \text{for } r_t \geq 20\sigma \end{cases}$$

Where, as before,  $U_t$  is the accumulated payoff by day  $t$  for the particular investor,  $r_t$  is the market return for that day, and  $L$  is the leverage multiplier. Here we also add two more terms:  $\sigma$

is the standard deviation of daily returns in the Swedish stock market (in our case,  $\sigma \approx 0.0146$ ), and  $\varepsilon_t$  is some randomly distributed error term (arbitrarily, I used  $\varepsilon_t \sim N\left(\mu, \frac{\sigma}{2}\right)$ , where  $\mu$  is the mean daily return  $\mu \approx 0.000325$ )<sup>35</sup>. All traders start with their assets at  $U_0 = 1$ .

The results of this game are shown in the graph below. Unsurprisingly, it shows the high-leverage traders getting wiped out on the high- $\sigma$  day. What was perhaps surprising is that the no-leverage trader (the blue line) could do better even in the very long run, just because of this one market crash.



*Figure 40 — When the levee breaks — Hedging and high-kurtosis events. In this graph, each trader earns the market return with some variation, no matter their leverage. The hedge breaks if the market incurs a 20-sigma event, at which point leverage becomes relevant. With high leverage-multipliers, traders can lose everything in one day. We might imagine the break-point as some sort of “counter-party risk”.*

These are, of course, extremely simplified models. They should not be taken as a practical lesson in market dynamics. Not only does real trading not work this way, but there is also more to risk than leverage.

But the reasoning may be useful to some traders. Specifically, it may be useful to watch out for high-kurtosis events, periods of high-volatility and long-downward price spirals caused by non-independence. In venture capital jargon, we might call this a “proof of concept”. The MMAR’s virtue is that it suggests that although these events are improbable, they are at least not astronomically improbable. And it allows the risks to be somewhat quantifiable. Investors who are exposed to short run risks may find the model helpful for thinking about their exposure. I shall discuss the MMAR’s usefulness more in [section 6.4](#).

Finally, I note two factors which might suggest that things can be worse than these games suggest. Firstly, the “monster” event in the MMAR-simulation — a 30% price drop — was only around 21 standard deviations, meaning that, in some sense, it was still less wild than the several

---

<sup>35</sup> When choosing the variance of the error term, the aim was merely visual. The goal was to keep the variance of the error term low enough so as not to be relevant, but just high enough so we could see the different lines in the beginning the graph.

days of 25-sigma moves reported by Goldman Sachs in 2008. Secondly, we also ignored the fact that traders usually compete not on absolute returns, but on returns relative to the market and their peers.

### 5.10. Summary of results

- The data appears to fit the MMAR's assumptions — namely, non-independent price changes and multifractal scaling behaviour.
- The results of the MMAR's simulations form nice smooth histograms, which suggests that the model might be useable for probabilistic forecasts.
- The  $H$ -exponent can affect price simulations in the long-run — an anti-persistent  $H$  keeps our prices more or less in the same place, while a high  $H$  can make a price simulation fly off to astronomical highs, or shrink to miniscule lows.  $H$ 's. In some sense, then, a high  $H$  suggests that the long run is more unpredictable. The krone has a low  $H$ ; Swedish stocks and the LIBOR have high  $H$ 's.
- All MMAR simulations are leptokurtic (they have excess kurtosis). The range of kurtoses increases if the real data had higher kurtosis. Gaussian kurtosis almost never goes out of the range of  $\pm 0.02$ .
- Visually, the MMAR's simulations look much more like the real data than a random walk. However, their accuracy is untestable. Most simulations fail the Kolmogorov-Smirnov test for identical distributions. This is most likely because of kurtosis differences — but this wide range of probable kurtoses is probably a built-in feature of the model.
- The MMAR can occasionally simulate some markets with truly monstrous kurtosis.
- A basic simulation of payoffs suggests that high leverage in the Gaussian world should lead to slow bankruptcy over many years, but in the MMAR world suggests bankruptcy can come much faster — perhaps even in one day.

## **VI. Analysis — Limitations and motives for future study**

*“Catastrophe mathematics, dealing with such events as falling off a height, is a new branch of the discipline [of mathematical economics], I am told, which has yet to demonstrate its rigor or usefulness. I had better wait.”*

Charles P. Kindleberger, *Manias, Panics, and Crashes*, 1978

Is Kindleberger’s wait finally over? In this classic work, the author laments that typical econometric regression techniques are unable to deal with market panics, and instead researchers have to introduce them exogenously as ‘dummy variables’<sup>36</sup>. But the MMAR and other fractal models have panics built in to them in the form of high kurtosis and volatility clustering. We have seen in [part V](#) that the MMAR can occasionally predict very high kurtosis indeed, even surpassing the 25-sigma events experienced by Goldman Sachs during the financial crisis.

But fractal models have their own share of problems. In this section, I try to argue that there are still many limitations to fractal modelling, which may help explain why Kindleberger’s ghost has not come back to congratulate us. One by one, we shall look at the MMAR’s methodological, practical and philosophical considerations, before finally considering the model’s usefulness in general.

### **6.1. Methodological considerations**

One issue that jumps off the page is the MMAR’s flimsy method for calculating the  $H$ -exponent. The authors give no way to calculate error bars — how do we know if we have enough evidence to reject the hypothesis that  $H = 0.5$  and that price moves are independent? Furthermore, how robust is this method for changes in the data? Peters (1994) gives one way to check the robustness of a method for checking the robustness of a methodology — if we

---

<sup>36</sup> A dummy variable is a special type of variable in a model, which can take either the value of 1 (meaning it is present) or zero (not present) at different time-points. Dummy variables are “exogenous” — the researcher has to specify manually whether their effects are present or not. Thus, the researcher must choose to turn it on or off, so to speak. It is used to represent factors which cannot normally be modelled or predicted but whose effects are nevertheless substantial. For example, an unforeseen financial crash.

shuffled the data-points around so that price moves had a random order, would  $H$  get closer to 0.5? Twenty years later, it looks like the authors have never published any robustness tests for their methodology.

Another issue is that there's also a lot of artistic license to these simulations — one has to make a lot of methodological judgement calls. For instance, how should one choose:

- the set of  $\Delta t$ 's (time increments) for the partition function?
- or the number and size of  $q$ 's (the moments) for the partition functions, which are used to perform OLS regressions to estimate  $\tau(q)$ ?
- or the number of splits  $b$  for simulating the multiplicative cascade? (as mentioned, the convention seems to be 2)
- or at which point to stop considering higher moments for the  $\tau(q)$  function? (I had to arbitrarily choose  $q = 4.0$ )
- or if the error range for the four parameters is important?

These may be pedantic details, but it may be fruitful to develop a standardised methodology that would make results more interpretable. A topic for future research would be to check the robustness of this methodology — do these judgements make no difference in the eventual simulations, or is the MMAR sensitive to initial conditions? In that case, should a methodology be standardized? Is a sensitive model even useful?

Another issue is the importance of *reverse causality*. Bianchi & Pianese (2005) point out that, while multifractality implies both a linear partition function on a log-log plot and a concave  $\tau(q)$ , the reverse does not imply multifractality. In fact, the authors show that a large number of Gaussian, non-fractal processes can be made to show the same two characteristics (linear partition functions and concave scaling functions). So the model is prone to *false positives*. (Although these are quite unlikely.)

As a criticism of this paper, one may also point out I have used an oversimplified strawman to represent “standard” finance — real financiers use more sophisticated models that don't just input variance and the mean. This is true, and I don't expect the Gaussian comparisons to be taken too seriously. Comparing the MMAR to other models can be a topic for future research. In the [literature review](#), I cited one such comparison to a sophisticated model. There, I pointed out that the main criticism of fractal researchers against GARCH models, namely that they are scale invariant — the simulations show high-kurtosis on a short time-scale, but using the same short-run parameters to simulate longer time-scales makes the simulation look almost

like a random walk. One thing to point out however, is that due to the competitive nature of the business, we don't really know what any of these financiers are doing, or what assumptions underlie their models. And if it helps the thesis at all, I can offer anecdotal evidence that the random-walk model is being used to guide the general public: when I made my first bank account in Swedbank, the risks of every savings portfolio offered by the bank were explained to me using a basic random walk model, with its corresponding confidence intervals for returns.

## **6.2. Practical considerations**

I can accept all the methodological quirks of the previous section with the excuse that it the MMAR is already clunky enough and doesn't need any more complications, as they would make it impractical to use. (Though presumably those problems could all be solved once-and-for-all with pre-built code, if only one of the authors bothered to make any code openly available.) So methodology aside, here we will consider the practicality of the model.

First, the obvious: we cannot draw a final practical conclusion from one study of simulations. The real test if fractal analysis is valuable is to select a large group of investors, tell half of them to use fractal analysis and the other to avoid it, and then to monitor them over a long time (e.g. 30 years). But seeing as this type of study is practically impossible, simulations are the next best thing.

If MMAR simulations are so realistic, then why haven't multifractal models caught on? One thing to point out is that using the model involves estimating quite a few things. As a list, it would be a large partition function, as well as the two functions  $\tau(q)$  and  $f(\alpha)$ , before finally getting to the four parameters  $H$ ,  $\alpha$ ,  $\lambda$  and  $\sigma^2$ . On top of that, it requires conceptual knowledge of statistical moments, Hurst exponents, Legendre transformations, multiplicative cascades, fractional Brownian Motion, multifractals, and other such complicated concepts. This mathematical background is pretty demanding for your average curious reader, and given that the theory doesn't purport to help anyone get rich, it's no surprise that it few people jot it down on their priority lists. Mandelbot's (2004) popular book is helpful for conceptual understanding, but it gives almost no practical guidance for how to use any of these ideas.

There are some other details which may be of historical significance. In his book on financial history, MacKenzie (2006) recounts how Mandelbrot's earlier models in the 1960's were branded useless for implying infinite variance: an implication of low- $\alpha$ , fat-tailed Stable distributions is that — due to infinite variance — the *expected value* of an asset price becomes

infinite. Understandably, most readers at the time deemed this to be a very impractical limitation. But the limitation didn't seem to bother Mandelbrot at all. In an amusing story, MacKenzie pointed this out to Mandelbrot in an interview, who reportedly looked back sternly and answered, "Why not?". Apparently, finite variance is just a "convention" that can be abandoned. That said, the MMAR does *not* imply infinite variance, and Mandelbrot and his co-authors saw it fit to mention this as an advantage.

Another problem is that the model also takes a while to calculate. Where the 10,000 random-walks took me perhaps 30 minutes to simulate for each market, it took over 20 hours to simulate the MMAR for Sweden (although, granted, the restrictions I added were prohibitive. Other countries were faster, but still on the scale of hours.) Needing to wait hours might be impractical when clients expect prompt advice. But perhaps this problem might be solved by using better code and simulating shorter timescales.

Requiring long timespans of data is certainly a practical problem. Fractal processes take a long time to develop, meaning that a distribution may not appear fractal in the short run. This is why I used 30 years of data for the empirical section. But many important markets *don't have* 30 years of data. For example, the euro versus US dollar (EUR/USD) is the most traded currency market in the world, but its data only goes back to around 1999, when the euro was first introduced. The MMAR's authors don't seem to offer a guide on how much data is enough.

Furthermore, the MMAR only works on medium timescales. As Mandelbrot (2004) points out: on very small-scale, high-frequency data, we encounter the effects of "market microstructure", defined by economists as the special conditions that occur on very short scales. Over the very long run, meanwhile, the "long memory" effect eventually disappears. But how wide is the range of these viable timescales? Surprisingly, quite wide. The original empirical MMAR study by Mandelbrot et al (1997), which looked at high-frequency data for the Deutschemark, saw multifractal scaling begin at as little as  $10^{2.5}$  seconds, which is about 5 minutes (see [appendix 8](#)). As for the upper limit, it seems that no one has looked into it. But if it already works reasonably well on timescales of three decades, then do they need to be any bigger?

Another problem with prescribing the MMAR, or indeed any model, for industry-wide standard use is that humans have the ability to react to other humans in unpredictable ways. We can consider a version of the Lucas Critique — what would happen if *everyone* started using multifractal models and became risk-savvy about wild, high-kurtosis randomness and periods of high-volatility? Would the effects of more conservative investment policies be cancelled out?

Would bankruptcy risks stay the same? I am inclined to say no. Presumably, wild markets are made up of wild investors — presumably, if the number of investors exposed to short-term, one-day risks went down, then the high volatility caused by market panics ought to go down as well. On the other hand, complacency could cause exacerbated bubbles which burst suddenly and unpredictably. To my knowledge, Mandelbrot never said anything about this issue.

An interesting point for further research would be some agent-based simulations, which could model how individual investors change their behaviour based on their “expectation” of market risks, and how these micro-level decisions affect the macro-level of the simulated market. LeBaron (2006) offers a clear and comprehensive literature review on agent-based computational finance. He notes that even very simple agent-based models lead to issues encountered in nonlinear dynamics and chaos theory — namely, very sensitive dependence on initial conditions. A small amount of statistical noise can get amplified by the simulated market system, which makes the outcome difficult to forecast. To my knowledge, no one has combined agent-based modelling with multifractal simulation of markets<sup>37</sup>.

### **6.3. Philosophical considerations**

More fundamentally, what does it mean to be predictive? In his 2015 book, *Superforecasting*, the political scientist Philip E. Tetlock points out a key problem in measuring forecasting success — probabilistic forecasts must not only be clearly defined, but also plentifully numerous. The outcome of a single probabilistic forecast is meaningless. Suppose one said there was a 70% chance of rain tomorrow. The next day, it didn't rain. Was the forecast wrong? After all, there was a 30% chance of *no rain*. Only with numerous forecasts of different probabilities can we assess predictive power — with a perfect forecaster, rain would happen 70% of the time when she said there was 70% chance of rain, and 30% of time when she gave a 30% chance, and so on. This is called *calibration*, and it can be used to measure if forecasts are well-calibrated versus over- or under-confident. Furthermore, for good forecasting, calibration must be supplemented by *resolution* — a forecaster (or model) cannot stay in the “cowardly” area of 40%–60% probabilities, but also make “brave” predictions about rare and near-certain events, i.e. those with very low or high probabilities. These two criteria are essential for assessing the predictiveness of both people and models. The implication is that, to check the

---

<sup>37</sup> Originally, I aspired to do this myself, but unfortunately I ran out of time (and page-count space!)

MMAR's predictiveness, we should compare it to 10,000 thirty-year periods in real financial markets(!). Since there is no chance of that happening, the model is in some sense untestable, despite the smooth probabilistic histograms that we saw in the [results section](#). However, smaller scale comparisons could be a viable alternative — we could test it over a few years or months, rather than decades. Modern financial markets are perhaps the most data-abundant structures in history — thus, they are very well suited for this type of testing, for both fractal and standard models.

But there is a bigger point to be made here — *why* are markets fractal, and what is *the point* of all this modelling? It cannot be to predict where or how prices will move— after all, we have randomness built into the simulation (in the form of the random fBm). So far, this appears to have mainly been an interesting mathematical exercise, whose results let us create charts that can fool the human eye. Theoretically, we have learned that market returns look like fractals — but *why* do they look like fractals? What is the underlying economic reasoning that causes these fractal relationships? And — more importantly — *so what?* What can we do with this model? And in what situations does it fail us? For an attempt to answer these questions, the usefulness of the MMAR will be the subject of the next section.

#### **6.4. So is any of this useful?**

After this long essay, one may be forgiven for wondering why anyone should care about fractal market analysis. Do we really need so many simulations to show that markets can be wild? Mervyn King, the Governor of the Bank of England from 2003-2013, summed it up nicely in his recent book: “The lesson is that no amount of sophisticated statistical analysis is a match for the historical experience that ‘stuff happens’” (King 2016), and I would argue that for most people this is good advice.

But the question is “how much stuff?”. When making practical decisions, there comes a point at which it becomes absurd to imagine worst-case scenarios, like markets crashing because of a meteor strike. The question is how much absurdity are we willing to consider, because preparing for all of it is too costly.

Fractal analysis is meant to help answer this question. In the introduction, I argued that most models leave out high-impact events which are actually quite common, and thus they underestimate risks. Fractal analysis is meant to give a more realistic portrayal of market risks, in a way that's not far from intuition. Its other benefit is that the risks can be at least somewhat

quantified, rather than told as anecdotal stories of market crashes. I believe that fractal analysis can be a useful tool for thinking about what to expect when investing, both in the short- and the long-term, without paralyzing the investor with inaction caused by excessive risk-avoidance.

That said, these tools don't appear to be useful for pricing assets or forecasting market movements. Two co-creators of the MMAR seem to disagree — the MMAR's co-authors, Calvet & Fisher (2008), have written an extensive mathematical treatise arguing that multifractal analysis can in fact be used for both forecasting volatility and pricing assets (particularly derivatives and options). The book is renowned by acquainted economists<sup>38</sup> — including some esteemed ones, such as John Y. Campbell of Harvard, who wrote the foreword.


But I am sceptical, for the following reasons: first, if price movements are discontinuous and fat-tailed, one can never really know when the next market jump will come. Second, the infinite variance aspect of fat-tailed, low- $\alpha$  distributions implies that the expected value of the asset — and thus its price — must also be infinite. Yes, the MMAR does not imply infinite variance, but it implies a wide variety of possible tail events — and in principle, that should imply a very wide range of reasonable prices. Third, and most importantly, measuring  $H$ , or  $\alpha_0$ , or the expected kurtosis accurately requires decades of data — and even then, the measurement can be very imprecise. Nassim Taleb (2007) has pointed out that any measurement of market kurtosis is likely to be an underestimate, meaning larger deviations will be underrepresented. This is because of the “Black Swan” problem — any measurement of past data is unlikely to include extreme events that are nevertheless possible in principle and also unpredictable in practice. Thus, at any given time, “even if you had a million data points,” it is more likely that you would compute that markets are less wild than in reality. Crucially, the MMAR makes tail events seem more probable *after they already happened*, not before — so it might be underestimating their probability, and without living in 10,000 universes, we wouldn't even know it. If the probability of rapid downturns is how we judge the price, and the probability number that our model gives us can itself change rapidly, then this is a problem — if that is the case, then how can any of this really be used to compute precise asset prices? And what is even a reasonable range? Are we forever destined to just shrug and say “stuff happens”?

The model also doesn't seem to be robust towards certain fat-tail events, which can make it behave inconsistently with its own predictions. Originally, this thesis featured an analysis of the Swiss exchange rate (USD/CHF) instead of the Norwegian. But unfortunately, the kurtosis of

---

<sup>38</sup> Although to me the book seemed esoteric and practically unreadable.

that market was so heavily dominated by one event that the MMAR generated silly-looking simulations. The event in question was the 15<sup>th</sup> January 2015, when the Swiss central bank spontaneously unpegged the franc from the euro, without giving any prior warning. The Swiss franc soared in value and the USD/CHF rate dropped by over 17% in one day, after which trading was uneventful. Many hedge funds made big losses.<sup>39</sup> This one event shoots the kurtosis up to 44.4; without it, kurtosis is only a modest 5.96<sup>40</sup>. And without the second biggest event (the peak in 2011), it falls to 2.3.



*Figure 41 — **Black Swans can ruin a perfectly good model.** This graph shows the price changes of the Swiss exchange rate (USD/CHF) and its MMAR simulation. Because the Swiss National Bank unannouncedly unpegged its exchange rate in January 2015, the price chart is very much dominated by one event, which ruins all subsequent simulations.*

The problem is that the exchange rate unpegging causes all subsequent simulations to make the market look wilder than it really is. The real price graph for the franc looked almost Gaussian in comparison to its simulation, which had multiple clustered periods of high volatility. Yet this simulation also had a lower kurtosis than the real data because it didn't have any humongous price moves. Traders of the franc probably had a much more comfortable life than the MMAR would suggest — except for that one day in 2015 which bankrupted a few of them. Interestingly, if we had used the data without that huge move (and perhaps also without the big jump upwards

<sup>39</sup> See *the Economist* magazine for a quick overview of the story: <https://www.economist.com/the-economist-explains/2015/01/18/why-the-swiss-unpegged-the-franc>

<sup>40</sup> And without the second biggest event (the tall move in 2011, which was a depreciation of about 9%), the kurtosis falls even more to a very modest 2.3. Still higher than Gaussian, but not as bad as the simulations would suggest.

in 2011), then the MMAR would probably have generated graphs that fit the data quite well, with a small probability of a massive move like the one in 2015. It is *only after the fact* that all subsequent simulations break.

So here's the quandary: the MMAR-model is supposed to suggest that high-kurtosis events are consistent with the data, but then breaks after those events actually happen. When a model's predictions are confirmed, shouldn't it start to fit *better*, rather than *worse*? To my knowledge, no one has ever pointed out this problem, nor suggested any solutions.<sup>41</sup>

So what's the final verdict on the usefulness of the MMAR? Mandelbrot (1999a) recommends using multifractals to "stress-test" financial portfolios. For my part, I can imagine using fractal simulations to create a trading simulator game for training employees at the trading desk of a financial firm. Players could learn to minimize trading losses on both slow, boring days, as well as fast, highly volatile days. For those who use computerized, algorithmic trading, fractal models could be used as a stress-test for the algorithm. Finally, it could also be used to model the payoffs (though perhaps not the "correct" pricing) of various financial options.

Beyond that, I'm out of ideas. I shall close with the truism that further research is needed.

### 6.5. Summary of limitations

- Like any model, the MMAR also has limitations.
- The methodology issues include: a flimsy method for calculating  $H$ ; the fact that there is no standardized methodology, which means the model requires a lot of personal judgements; and that the model doesn't take into account reverse causality (non-fractal data can sometimes still look as though it was fractal, giving a false positive).
- The practical issues include: the need to calculate several parameters; too much mathematical complexity for an average user; long calculation times; needing too many years of data; working only on medium timescales (though this range of these is quite wide); and that humans can react in unpredictable ways, so we don't know how the market would change if everyone adopted fractal models.

---

<sup>41</sup> Our LIBOR data had an even higher kurtosis of nearly 87. Perhaps this 'after-the-fact'-breaking is another reason why the LIBOR's MMAR simulations performed so poorly in the KS-test.

- The philosophical issues include: that if we want to test if the model is good at probabilistic forecasts, we would need to check it in 10,000 thirty-year periods, which is certainly impractical (although the model's predictability may be testable on shorter timescales); and that we don't seem to have gained any insight into *why* markets often look like fractals.
- Finally, is the model useful? In many ways, it has potential. For one thing, it gives a quantitative way to assess the impact of extreme events, which may be better than simply saying "stuff happens".
- However, the model does not appear useful for pricing assets (though two of its creators disagree); there is no way of knowing if it is underestimating potential risks until after they happen; and worst of all, the model can sometimes break when its own predictions are confirmed! A dataset with one very unusual event can make every subsequent simulation look clearly different from most of the data.
- As a final verdict, the model may be useful for stress-testing portfolios and trading algorithms, or for simulating "trading games" for aspiring traders to practice on.

## VII. Conclusions

*“To say much with little: Such is the goal of good science.”*

Benoit B. Mandelbrot (2006)

*“A framework that claims that the future is inherently uncertain cannot be accused of perfection. Yet it can provide important insights into reality; it can even anticipate the future within bounds, although the bounds themselves are uncertain and variable.”*

George Soros, Hungarian-American investor and philanthropist (2009), quoted from *The Soros Lectures* (2010)

In this thesis, I tried to study the Multifractal Model of Asset Returns, including its theoretical reasoning and its outputs. I studied that daily price data for three markets — the Norwegian USD/NOK exchange rate, the Swedish OMXS30 stock market index, and the 12-month LIBOR-rate in British pounds.

I found that the data fits well with many of the model’s assumptions — namely, non-independent price movements, many high-kurtosis events, clustered periods of high volatility and visible scaling behaviour in raw moments. I also found that the model does a reasonably good job of simulating all these different types of markets — but this is a tentative conclusion, because the model is not easily testable. The model suggests that the Hurst exponent can greatly affect probable price paths in the long run — a high  $H$  amplifies trends (suggesting that prices have a greater chance of ending up very high or very low), while a low  $H$  cancels them (suggesting that the price will probably not change much in the long run). Meanwhile, the model can simulate a very wide range of kurtoses, and this range increases if the kurtosis of the data is higher. The distributions of the simulations’ moments (the means, standard deviations, kurtoses etc.) also appear quite smooth, suggesting that the MMAR might be useable for probabilistic forecasts. The “monster” simulations that the MMAR occasionally gives may be useful for stress-testing portfolios and trading algorithms.

Mandelbrot was a revolutionary mathematician and scientist who happened to have an on-and-off relationship with financial modelling for five decades of his life. The MMAR was his last contribution to this field before his death in 2010. His memoirs<sup>42</sup> portray a man who loved

---

<sup>42</sup> In the afterword of *The Fractalist* (Mandelbrot 2012), Michael Frame sums up Mandelbrot’s approach quite nicely: “Look at the pictures, remember the simple formula, think, and be amazed.”

to describe complex phenomena with profoundly simple equations. The MMAR is a model like no other before, both in the way it is constructed and in what it predicts about market dynamics.

But who was Mandelbrot kidding when he said the MMAR was “saying much with little”? The model involves very complicated mathematical reasoning that can be very difficult to follow. Its unveiling in 1997 took three full-length academic papers. It takes a while to simulate. And it requires years of data to be realistic. On top of that, I’m not sure to what extent we can call it “science”. The model only seems to make mostly qualitative predictions, and sometimes it even breaks when its own predictions actually occur. And its creators don’t seem to offer any definitive ways of testing the model. It seems that it remains untestable even 20 years later.

It is my conclusion that, unfortunately, Mandelbrot and his creators have still left many questions unanswered. As it stands, the MMAR and its offshoots seem to only be useful in the sense described above by George Soros — the MMAR can provide important insights about scaling behaviour, high-kurtosis events, volatility clustering and non-random walks. Qualitatively, all these features are clearly present in the model’s simulations, as well as in the real data. To an extent, the model might even be able to anticipate the level of market volatility and the likelihood of huge market moves — but only within bounds, and those bounds are uncertain and variable. According to the MMAR’s surviving co-authors, Calvet & Fisher (2008), the offshoot models can predict some of these behaviours ahead of time, sometimes. But none of the models seem to give any fundamental reason for *why* markets behave the way they do. It seems that, overall, many market phenomena remain a mystery.

With that in mind, I can still praise the model as an excellent first step. The MMAR has the benefit that its maps of the sea make the reader very aware that whirlpools exist, even if it can’t tell you where they are. This is an important improvement on a basic random-walk model, or other models with similar limitations. In a very basic simulation, I tried to show that ignoring the possibility of volatile, high-kurtosis events can lead to rapid bankruptcy, equivalent to losing all the money we ever made in one day. With its random but qualitatively accurate simulations, the MMAR could be a very useful way of stress-testing portfolios so that investors could withstand these events. Its main benefit is that it offers some *quantitative* way of thinking about these problems — rather than just saying “stuff happens”, we are one step closer to saying how much “stuff”. Better stress-testing should lead to less bankruptcy in our financial system, and hopefully fewer financial crises.

I believe that — despite having been mostly ignored by mainstream financial theorists — research into fractal financial modelling may teach some interesting lessons about market

dynamics. It might also offer important insights for how traders can avoid bankruptcy. But it is a shame that most of this research is written in a very esoteric style that is not readily accessible for newcomers.

I close this thesis with a short summary of what I've learned by dipping my toes in the water of this arcane discipline.

### **7.1. Summary — What can fractals do and not do?**

**Does empirical market data fit the MMAR model?** This exploration has found that yes, it does. (Albeit according to statistical tests made by an indignant mathematician to check his own theories.) For all three markets, I found high kurtosis, non-independent price changes and scaling behaviour in the partition function.

**What do fractal models show?** The MMAR model and its offshoots simulate a market with three key features — high kurtosis in price changes, non-independent price movements and clustered periods of high volatility.

**How is the MMAR better than standard models?** It shows the three characteristics mentioned above, and it can simulate different timescales with the same parameters (e.g. days, months, or decades). The basic random walk shows no kurtosis and no price independence, but at least it calculates quickly. GARCH calculates slowly but looks Gaussian if you use the same parameters for longer timescales.

**What kind of data do fractal models need?** Price data from any financial market. The more the better, although millisecond-data is probably not necessary; daily frequency will suffice. What is necessary is a long time-scale: years rather than days.

**What do we need for an MMAR simulation?** All we need is to do is calculate the four parameters:  $H$ ,  $\alpha$ ,  $\lambda$  and  $\sigma^2$ .

**What kind of markets can the MMAR analyse?** Apparently, any kind, as long as they show multifractal scaling. In this thesis, I found that the model works well for stocks, interest rates and exchange rates. In principle I could have used it to study oil or banana prices as well.

**How are MMAR simulations made?** By combining two random processes: (1) a stochastic fractional Brownian motion; and (2) a “trading time” cumulative probability distribution, constructed using a multifractal binomial cascade.

**Are the simulations realistic?** To the naked eye, yes. To statistical testing, maybe not. The graphs of price returns look like they could have been from real data, but the most simulations seem to fail the Kolmogorov-Smirnov test for identical distributions. But this may be a fault of the test. To my knowledge, there is no clear way to test the model’s validity.

**Can fractals predict where the markets will go, or when volatility will change?** Probably not. But I suspect that that’s not the point; the point is to simulate what *could* in principle happen, and then to consider how that would affect us.

**Can fractals price assets?** Some say yes, but that literature is so far perhaps too technical to be of common, practical use. I suspect that this will be difficult, because it’s not easy to get 30 years of relevant data for newly-made securities. If the technical hurdles are overcome, I suspect that fractal models might become useful for pricing options.

**Can fractals predict who will eventually go bankrupt?** In a simple simulation, yes. In real life, perhaps. I invite other, more mathematically creative researchers to check the model.

**Where’s the catch? What are the MMAR’s limitations?** The model might be too mathematically complicated to become mainstream; many details require a lot of human judgement when being used; the data might not be multifractal on *all* timescales, such as the very short or very long; there isn’t really a good way to check the model’s accuracy beyond its qualitative features; the model needs years of data and can’t be used for new securities; finally, in a strange turn of scientific logic, the model appears to break after one of its high-kurtosis events actually happens.

With all that in mind, I still think the model is a good first step in the right direction.

**Who is the model useful for?** The model may be useful for stress-testing financial portfolios, and so is in principle useful for any investor, but especially those who are exposed to short-term risks. It might also be useful for a trading-simulation game, or for testing the trading algorithms of computers.

**What should be the aim of future research?** Some good starting points would be to devise a way to test if the model is accurate, including making a better definition for what is a realistic simulation. We should also standardize a methodology for how to perform fractal simulations and fractal financial analysis. Another interesting idea would be to combine fractal simulation with agent-based modelling, to consider how investors behave under MMAR-assumptions, and to see if there is anything they can do to avoid bankruptcy caused by persistent down-periods, periods with high volatility, or simply massive one-day price moves. One more idea could be to see why the model breaks if the data has a very high kurtosis caused mostly by one event, if the newer fractal models suffer from the same problem, and whether there's any easy way to solve this issue. On top of all of that, I think the whole field needs some more philosophical consideration — what are we even trying to predict with our modelling, and how do we know what we know? And also, *why* do markets show multifractal characteristics at all? These fundamental questions could be made clearer.

## Glossary

**Canonical** — referring to the multiplicative cascade. In a canonical measure, the total mass at each step is random. Meanwhile, in a conservative measure, the total mass is fixed at each step (typically as  $\sum[M_\beta] = 1$ ).

**Cascade** — a cascade is a type of fractal process. It is created by generating a flat line, dividing it into sections, and then making each of those sections either taller or shorter according to some pre-defined rules. See sections [2.3](#) or [4.3](#) for some examples.

**Compound process** — a function within a function. For example, we can combine  $f(x)$  and  $g(x)$  to form  $f(g(x))$ . The MMAR-model is a compound process, combining a simulated fractional Brownian motion ( $B_H(t)$ ) and a simulated “trading time” ( $\theta(t)$ ) to form a compound function of time:  $X_t \equiv B_H[\theta(t)]$ .

**Cross-over point** — the bound in fractal data where a new mathematical relation takes hold; common in real fractal, rather than in theoretical. An example is that market patterns may be different at intervals below two hours and above 180 days. Cross-over points are why multifractal models can be more useful than simple fractal models.

**Dependence** — the opposite of statistical independence, i.e. when yesterday’s price *does* in fact influence today’s. Measured by the parameter  $H$ , where  $0 \leq H \leq 1$ . For the normal distribution, random-walk case:  $H = 0.5$ , i.e. independent. At higher  $H$ ’s, trends tend to go in the same direction; at lower  $H$ ’s, effects tend to cancel themselves out quickly.

**Fractional Brownian Motion (fBm)** — this is a generalization of Brownian motion, which is a random process. For example, gas particles are often modelled using Brownian motion, the assumption being that a particle’s next move can be in any random direction. In the fBm, however, increments do not have to be independent of each other — an up-movement can either make the next movement more likely to also be an up-movement, or more likely to be a down-movement, depending on the  $H$ -exponent. The fBm is usually simulated with computers. See [section 4.3](#).

**Fractal** — an abstract object used to describe and simulate naturally occurring objects. Artificially created fractals commonly exhibit similar patterns at increasingly small scales. Fractals can be in the form of a curve (such as the [Mandelbrot set](#)) or a geometrical figure (such as the [Sierpinski triangle](#)).

**Generator** — the repeated “process” which is applied to the initial shape at each step — for the Sierpinski triangle, the generator is cutting out a triangle in the middle. For a cascade, the generator is the process of splitting the cascade into parts, and then applying the masses.

**Heterogeneous (fractal)** — a fractal model with various different fractal properties embedded within it. A heterogeneous fractal is a *multifractal*.

**Hurst exponent ( $H$ )** — this is a parameter ranging from 0 to 1. It tells us whether changes are independent or not. Specifically,  $H = 0.5$  is the random-walk case, in which past events have no influence on future ones, such as in a coin flip. A big  $H$  (meaning  $H > 0.5$ ) suggests that changes are dependent — an up-movement is more likely to be followed by another up-movement. The opposite is true for a small  $H$  under 0.5 — up-movements are more likely to be followed by a down-movement, and vice-versa.

**Initiator** — the starting shape that begins the fractal process — in the Sierpinski triangle, it is the initial blue triangle. In a fractal cascade, it is the initial flat line, which gets divided.

**Kolmogorov-Smirnov test** — this is a statistical test, designed to check if two sets of data could have come from the same probability distribution. See [section 5.6](#).

**Kurtosis** — a measure of the “tailedness” of a probability distribution. A distribution with high kurtosis has somewhat non-intuitive scaling behaviour, in which extreme events are more relatively common. The Normal distribution has a kurtosis of 3. Excess kurtosis is thus defined as kurtosis minus 3.

**Leptokurtic** — Meaning the distribution has excess kurtosis. This statistical property of a probability distribution means that the distribution has “fat tails” and a tall peak in the middle. Leptokurticity implies that extreme events are more likely than under a simple normal distribution.

**Long-term Dependence** — If two variables are long-term dependent, then they are correlated over long time scales — days, weeks, years etc. Events even in the distant past can show a correlational effect with today. For instance, the break-up of Rockefeller’s Standard Oil company in 1911 continues to have an effect on the newly formed companies (ExxonMobil, ChevronTexaco etc.) For water reservoirs, past years have a long-term effect: past rainfall fills the reservoir, which drains slowly, so the volume of water in the reservoir is affected by rainfall even from many years before; one could have many dry years and still see a correlation.

**Moment** — in the statistical sense, a moment is a measure of the shape of a function, such as a probability distribution. For example, the average is the 1<sup>st</sup> moment. Other moments are

not easy to interpret in their “raw” form (sometimes called “crude”), but their centralized or normalized versions do. For instance, the variance is the 2<sup>nd</sup> moment centralized, skewness is the 3<sup>rd</sup> moment normalized, and non-excess kurtosis is the 4<sup>th</sup> moment normalized.

**Multifractal** — a heterogeneous fractal — one that has different properties at different locations, but which are all nevertheless fractal. A mountain can be very sharply jagged at its higher points and peaks, while being relatively smooth in the low foothills. Making a multifractal involves ‘function-within-a-function’ modelling.

**Power-law distribution** — a distribution between variables where the relationship is not linear, but rather involves exponents. Mathematically, it can be expressed as:

$$Y = kX^\alpha$$

Where  $Y$  and  $X$  are the variables of interest,  $k$  is a constant and  $\alpha$  is the (measured) exponent. A simple example of a power-law relationship involves squares: if you double the length of the sides, the area will quadruple. More generally, increase side length by  $x$ , and area will increase by  $x^2$ .

**Scaling** — this is a rule that tells us what the shape looks like at different scales. What happens if we zoom in or zoom out? A linear scaling rule suggests self-similarity, meaning that the shape is a simple fractal and that it looks the same at every scale. A complication occurs if the scaling rule is non-linear. This means that the process is merely self-affine, and that the shape is a multifractal. Financial data is usually shows non-linear scaling.

**Self-affinity** — Related to self-similarity, but a more general case. Like with self-similarity, the object has a similar structure on different scales. However, when scaling, one must scale by different amounts for the  $x$ - and  $y$ -directions. (E.g., instead of a pure 2x zoom, scale down  $y$ -axis by 2 and scale down the  $x$ -axis by 1.5 to see the same shape.) In financial markets, self-affinity is far more common than the “perfect” self-similarity.

**Self-similarity** — The property of having a substructure analogous or identical to an overall structure. Basically, an object is “self-similar” if it looks very similar at any scale. (Examples from real-life and nature include: broccoli; snowflakes; bronchi inside lungs; coastlines; earthquake seismic activity patterns; galaxy distributions). This is the “perfect” case of self-affinity.

**Stable distribution** — a technical term; a distribution is said to be *stable* if a combination of two independent random variables with this particular distribution yields the same distribution, after rescaling for size and location

**Trading time** — Formally, this is a lognormal cascade converted into a cumulative distribution function (CDF). Conceptually, “trading time” is meant to represent the real-life human experience that time doesn’t feel like it moves in a linear way. Sometimes, times are “slow” and uneventful. Meanwhile, at other times, time seems to move fast with lots of big events happening in a short time-frame. In markets, some days are more volatile than others. “Trading time” is the component of the MMAR that helps to capture the clustering of high volatility periods.

**Volatility** — this is a measure of how big the changes in a process tend to be. In finance, volatility is usually measured using standard deviation. Some time-periods are more volatile than others — for example, the onset of financial crises usually causes high volatility.

**Zipf’s law** — an empirical proposition, named after George Kingsley Zipf, dealing with word frequencies. The law states that the frequency of any word is inversely proportional to its rank on the frequency table. Thus, the  $N^{\text{th}}$  most frequent word appears with a frequency proportional to  $1/N$ . Zipf’s Law holds for most languages, including English.

## **Bibliography**

- Addison, P.S., 2002. *The illustrated wavelet transform handbook: introductory theory and applications in science, engineering, medicine and finance*. CRC press.
- Alford, M. Legendre Transforms. Available from: <http://alford.fastmail.us/physics/legendre.pdf>
- Bachelier, L., 1900. *Théorie de la spéculation*. Gauthier-Villars.
- Backhouse, R.E., 2010. *The Puzzle of Modern Economics: Science or Ideology?*. Cambridge University Press.
- Barunik, J. and Kristoufek, L., 2010. On Hurst exponent estimation under heavy-tailed distributions. *Physica A: Statistical Mechanics and its Applications*, 389(18), pp.3844-3855.
- Batten, J.A., Kinatader, H. and Wagner, N., 2014. Multifractality and value-at-risk forecasting of exchange rates. *Physica A: Statistical Mechanics and its Applications*, 401, pp.71-81.
- Bell, P., 2012. Goodness of fit test for the multifractal model of asset returns.
- Bianchi, S. and Pianese, A., 2007. Modelling stock price movements: multifractality or multifractionality?. *Quantitative Finance*, 7(3), pp.301-319.
- Black, F. and Scholes, M., 1973. The pricing of options and corporate liabilities. *Journal of political economy*, 81(3), pp.637-654.
- Brock, W.A. and De Lima, P.J., 1996. 11 Nonlinear time series, complexity theory, and finance. *Handbook of Statistics*, 14, pp.317-361.
- Bundt, T. and Murphy, R.P., 2008. Are residual economic relationships normally distributed? Testing an assumption of neoclassical economics. *The Review of Austrian Economics*, 21(4), p.329.
- Calvet, L., Fisher, A. and Mandelbrot, B., 1997. Large deviations and the distribution of price changes. *Cowles Foundation Discussion Paper No. 1165; Sauder School of Business Working Paper*.
- Calvet, L. and Fisher, A., 2002. Multifractality in asset returns: theory and evidence. *Review of Economics and Statistics*, 84(3), pp.381-406.
- Calvet, L.E. and Fisher, A.J., 2008. *Multifractal volatility: theory, forecasting, and pricing*. Academic Press.

- Cantor, G., 1884. De la puissance des ensembles parfaits de points. *Acta Mathematica*, 4(1), pp.381-392.
- Carfi, D., Lapidus, M.L., Pearse, E.P. and Van Frankenhuisen, M. eds., 2013. *Fractal Geometry and Dynamical Systems in Pure and Applied Mathematics: Fractals in pure mathematics*(Vol. 600). American Mathematical Society.
- Cheung, Y.W. and Lai, K.S., 1995. A search for long memory in international stock market returns. *Journal of International Money and Finance*, 14(4), pp.597-615.
- Cont, R., 2001. Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance* 1(2), pp.223-236
- Cont, R., 2007. Volatility clustering in financial markets: empirical facts and agent-based models. In *Long memory in economics* (pp. 289-309). Springer, Berlin, Heidelberg.
- Davies, R.B. and Harte, D.S., 1987. Tests for Hurst effect. *Biometrika*, 74(1), pp.95-101.
- Davis, M.S., 1971. That's interesting! Towards a phenomenology of sociology and a sociology of phenomenology. *Philosophy of the social sciences*, 1(2), pp.309-344.
- Dieker, T., 2004. *Simulation of fractional Brownian motion* (Doctoral dissertation, Masters Thesis, Department of Mathematical Sciences, University of Twente, The Netherlands).
- Dowd, K., Cotter, J., Humphrey, C. and Woods, M., 2011. How unlucky is 25-sigma?. *arXiv preprint arXiv:1103.5672*. <https://arxiv.org/pdf/1103.5672.pdf>
- Ekins, E., 2017. 'Wall Street vs. The Regulators: Public Attitudes on Banks, Financial Regulation, Consumer Finance, and the Federal Reserve: Results from the Cato Institute 2017 Financial Regulation Survey'. *Cato Institute* [Accessed on 2 April 2019]: <https://www.cato.org/survey-reports/wall-street-vs-regulators-public-attitudes-banks-financial-regulation-consumer>
- Embrechts, P., Klüppelberg, C. and Mikosch, T., 1997/2013. *Modelling extremal events: for insurance and finance* (Vol. 33). Springer Science & Business Media.
- Fama, E.F., 1965. The behavior of stock-market prices. *The journal of Business*, 38(1), pp.34-105.
- Fama, E.F., 1970. Efficient capital markets: A review of theory and empirical work. *The journal of Finance*, 25(2), pp.383-417.
- Feder, J., 2013. *Fractals*. Springer Science & Business Media.
- Fisher, A.J., Calvet, L.E. and Mandelbrot, B.B., 1997. Multifractality of Deutschemark/US dollar exchange rates. *Cowles Foundation Discussion Paper No. 1166; Sauder School of Business Working Paper*.

- Fillol, J., 2003. Multifractality: theory and evidence, an application to the french stock market. *Economics Bulletin*, 3(31), pp.1-12
- Gleick, J., 1987/2011. *Chaos: Making a new science*. Open Road Media.
- Granero, M.S., Segovia, J.T. and Pérez, J.G., 2008. Some comments on Hurst exponent and the long memory processes on capital markets. *Physica A: Statistical Mechanics and its applications*, 387(22), pp.5543-5551.
- Gneiting, T. and Schlather, M., 2004. Stochastic models that separate fractal dimension and the Hurst effect. *SIAM review*, 46(2), pp.269-282.
- Günay, S., 2016. Performance of the Multifractal Model of Asset Returns (MMAR): evidence from emerging stock markets. *International Journal of Financial studies*, 4(2), p.11.
- Harford, T., 2014. An astonishing record-of complete failure. *FT Magazine-Undercover Economist*, 30, p.2014.
- Hurst, H.E., 1951. Long term storage capacity of reservoirs. *Transactions of the American Society of Civil Engineers*, 116(776), pp.770-808.
- Jamdee, S., 2005. *Multifractal models and simulations of the US term structure* (Doctoral dissertation, Kent State University).
- Kay, J., 2015. *Other People's Money*. Profile Books Limited.
- Keppler, M., 1990. Risk is not the same as volatility. *Die Bank*, 11, pp.610-614.  
[http://www.kepplerfunds.com/load/publications/p03\\_eng.pdf](http://www.kepplerfunds.com/load/publications/p03_eng.pdf)
- Kindleberger, C.P., 1978. *Manias, Panics and Crashes*. The Macmillan Press Ltd.
- King, M., 2016. *The End of Alchemy: Money, banking, and the future of the global economy*. WW Norton & Company.
- Kobeissi, Y.H., 2013. *Multifractal Financial Markets: An Alternative Approach to Asset and Risk Management*. New York: Springer.
- Kristoufek, L., 2011. Multifractal height cross-correlation analysis: A new method for analyzing long-range cross-correlations. *EPL (Europhysics Letters)*, 95(6), p.68001.
- Kristoufek, L., 2012. Fractal markets hypothesis and the global financial crisis: Scaling, investment horizons and liquidity. *Advances in Complex Systems*, 15(06), p.1250065.
- Kristoufek, L., 2013. Fractal markets hypothesis and the global financial crisis: wavelet power evidence. *Scientific reports*, 3, p.2857.
- Kristoufek, L. and Vosvrda, M., 2013. Measuring capital market efficiency: Global and local correlations structure. *Physica A: Statistical Mechanics and its Applications*, 392(1), pp.184-193.

- Kristoufek, L. and Vosvrda, M., 2014. Measuring capital market efficiency: long-term memory, fractal dimension and approximate entropy. *The European Physical Journal B*, 87(7), p.162.
- Larsen, P.T., 2007. Goldman pays the price of being big. *Financial Times*, August 13, p.3. Available from: <https://www.ft.com/content/d2121cb6-49cb-11dc-9ffe-0000779fd2ac>
- LeBaron, B., 2006. Agent-based computational finance. *Handbook of computational economics*, 2, pp.1187-1233.
- Liu, R., 2008. Multivariate multifractal models: Estimation of parameters and applications to risk management (Doctoral dissertation, Christian-Albrechts Universität Kiel).
- Lo, A.W. and MacKinlay, A.C., 2002. *A non-random walk down Wall Street*. Princeton University Press.
- Lux, T., 2001. Turbulence in financial markets: the surprising explanatory power of simple cascade models. *Quantitative finance*, 1(6), pp.632-640.
- MacKenzie, D., 2006. *An engine, not a camera: How financial models shape markets*. Mit Press.
- Makridakis, S. and Hibon, M., 2000. The M3-Competition: results, conclusions and implications. *International journal of forecasting*, 16(4), pp.451-476.
- Makridakis, S., Hogarth, R.M. and Gaba, A., 2009. Forecasting and uncertainty in the economic and business world. *International Journal of Forecasting*, 25(4), pp.794-812.
- Makridakis, S., Spiliotis, E. and Assimakopoulos, V., 2018. The M4 Competition: Results, findings, conclusion and way forward. *International Journal of Forecasting*, 34(4), pp.802-808.
- Malkiel, B.G., 2003. The efficient market hypothesis and its critics. *Journal of economic perspectives*, 17(1), pp.59-82.
- Malkiel, B.G. and McCue, K., 1985. *A random walk down Wall Street* (Vol. 332). New York: Norton.
- Mandelbrot, B.B., 1963. The variation of certain speculative prices. In *Fractals and scaling in finance* (pp. 371-418). Springer, New York, NY.
- Mandelbrot, B., 1967. The variation of some other speculative prices. *The Journal of Business*, 40(4), pp.393-413.
- Mandelbrot, B.B., 1972. Possible refinement of the lognormal hypothesis concerning the distribution of energy dissipation in intermittent turbulence. In *Statistical models and turbulence* (pp. 333-351). Springer, Berlin, Heidelberg.
- Mandelbrot, B.B., 1975. Les objets fractals: forme, hasard et dimension.
- Mandelbrot, B.B., 1989. Multifractal measures, especially for the geophysicist. In *Fractals in geophysics* (pp. 5-42). Birkhäuser, Basel.

- Mandelbrot, B.B., 1997. *Fractals and Scaling in Finance: Discontinuity, concentration, risk*. Selecta volume E. Springer Science & Business Media.
- Mandelbrot, B.B., 1999. A multifractal walk down Wall Street. *Scientific American*, 280(2), pp.70-73.
- Mandelbrot, B.B., 1999/2013. *Multifractals and 1/f Noise: Wild Self-Affinity in Physics (1963-1976)*. Springer.
- Mandelbrot, B.B., 1999/2008. How Fractals Can Explain What's Wrong with Wall Street. *Scientific American* 15.  
<https://www.scientificamerican.com/article/multifractals-explain-wall-street/>
- Mandelbrot, B.B., 2005. Parallel cartoons of fractal models of finance. *Annals of Finance*, 1(2), pp.179-192.
- Mandelbrot, B.B., 2005. The inescapable need for fractal tools in finance. *Annals of Finance*, 1(2), pp.193-195.
- Mandelbrot, B.B., 2012. *The Fractalist: Memoir of a Scientific Maverick*. Vintage.
- Mandelbrot, B., Fisher, A. and Calvet, L., 1997. A multifractal model of asset returns. *Cowles Foundation Discussion Paper No. 1164; Sauder School of Business Working Paper*.
- Mandelbrot, B.B. and Hudson, R.L., 2004/2010. *The (Mis)behaviour of Markets: a fractal view of risk, ruin and reward*. Basic Books.
- Mandelbrot, B.B. and Wallis, J.R., 1968. Noah, Joseph, and operational hydrology. *Water resources research*, 4(5), pp.909-918.
- Mandelbrot, B.B. and Wallis, J.R., 1969. Computer experiments with fractional Gaussian noises: Part 2, rescaled ranges and spectra. *Water resources research*, 5(1), pp.242-259.
- Markowitz, H., 1952. Portfolio selection. *The journal of finance*, 7(1), pp.77-91.
- Martino, W. and Frame, M., 2010. Multifractal Measures of Time Series:  $f(a)$  Surfaces. *International Journal of Bifurcation and Chaos*, 20(08), pp.2453-2470.
- Munger, C.T., Kaufman, P.D. and Wexler, W.E.B., 2005/2011. *Poor Charlie's Almanack: The Wit and Wisdom of Charles T. Munger*. PCA Publications L.L.C.
- Muzy, J.F., Bacry, E. and Arneodo, A., 1994. The multifractal formalism revisited with wavelets. *International Journal of Bifurcation and Chaos*, 4(02), pp.245-302.
- Peitgen, H.O., Jürgens, H. and Saupe, D., 2006. *Chaos and fractals: new frontiers of science*. Springer Science & Business Media.
- Peng, C.K., Buldyrev, S.V., Havlin, S., Simons, M., Stanley, H.E. and Goldberger, A.L., 1994. Mosaic organization of DNA nucleotides. *Physical review e*, 49(2), p.1685.

- Penrose, R., 2004/2006. *The Road to Reality: A complete guide to the laws of the universe*. Vintage.
- Peters, E.E., 1991. *Chaos and order in the capital markets: a new view of cycles, prices, and market volatility*. John Wiley & Sons.
- Peters, E.E., 1994. *Fractal market analysis: applying chaos theory to investment and economics*. John Wiley & Sons.
- Peters, E.E., 1999. *Patterns in the dark: understanding risk and financial crisis with complexity theory* (Vol. 71). John Wiley & Sons Inc.
- Qian, B. and Rasheed, K., 2004, November. Hurst exponent and financial market predictability. In *IASTED conference on Financial Engineering and Applications* (pp. 203-209).
- Reinhart, C.M. and Rogoff, K.S., 2009. *This Time is Different: Eight centuries of financial folly*. Princeton University Press.
- Reinhart, C.M. and Rogoff, K.S., 2010. Growth in a Time of Debt. *American Economic Review*, 100(2), pp.573-78.
- Segnon, M. and Lux, T., 2013. Multifractal models in finance: Their origin, properties, and applications (No. 1860). *Kiel working paper*.
- Sharpe, W.F., 1964. Capital asset prices: A theory of market equilibrium under conditions of risk. *The journal of finance*, 19(3), pp.425-442.
- Shiller, R.J., 2000/2015. *Irrational exuberance: Revised and expanded third edition*. Princeton university press.
- Shiller, R.J., Kon-Ya, F. and Tsutsui, Y., 1996. Why did the Nikkei crash? Expanding the scope of expectations data collection. *The Review of Economics and Statistics*, pp.156-164.
- Sobolev, D. 2014. Financial Applications of Human Perceptions of Fractal Time Series, Doctoral thesis, UCL (University College London). Available from: [http://discovery.ucl.ac.uk/1461731/5/Thesis%2011\\_02\\_2015%20Final%20version%20for%20submission.pdf](http://discovery.ucl.ac.uk/1461731/5/Thesis%2011_02_2015%20Final%20version%20for%20submission.pdf)
- Sornette, D., 2003/2017. Why stock markets crash: critical events in complex financial systems (Vol. 49). Princeton University Press.
- Soros, G., 2010. *The Soros Lectures: At the Central European University*. PublicAffairs.
- Stanley, H.E., 2003. Statistical physics and economic fluctuations: do outliers exist?. *Physica A: Statistical Mechanics and its Applications*, 318(1-2), pp.279-292.
- Taleb, N.N., 2007. *The Black Swan: The impact of the highly improbable*. Random house.
- Taylor, S.J., 1986/2008. *Modelling financial time series*. New York: Wiley.

- Tetlock, P.E. and Gardner, D., 2015. *Superforecasting: The art and science of prediction*. Crown Publishing Group.
- Thaler, R.H. and Sunstein, C.R., 2009. *Nudge: Improving decisions about health, wealth, and happiness*. Penguin.
- Tolstoy, L. N., 1877. *Anna Karenina*. Translated by Richard Pevear and Larissa Volokhonsky. Penguin Classics.
- Webster, M. and Sell, J. eds., 2007/2014. *Laboratory experiments in the social sciences*. Elsevier.
- Weierstrass, K., 1872. Über continuirliche functionen eines reellen arguments, die für keinen wirth des letzteren einen bestimmten differentailquotienten besitzen, Akademievortrag. *Math. Werke*, pp.71-74.
- Wengert, C., 2010. Multifractal Model of Asset Returns (MMAR). *Last modified December, 12*. Available from: <https://www.mathworks.com/matlabcentral/fileexchange/29686-multifractal-model-of-asset-returns-mmara>
- Weron, R. and Przybyłowicz, B., 2000. Hurst analysis of electricity price dynamics. *Physica A: Statistical Mechanics and its Applications*, 283(3-4), pp.462-468.
- Weron, R., 2002. Estimating long-range dependence: finite sample properties and confidence intervals. *Physica A: Statistical Mechanics and its Applications*, 312(1-2), pp.285-299.
- Yamasaki, K. and Mackin, K.J., 2004. Market simulation displaying multifractality. In *The Application of Econophysics* (pp. 223-228). Springer, Tokyo.
- Zhong, J. and Zhao, X., 2012. Modeling complicated behavior of stock prices using discrete self-excited multifractal process. *Systems Engineering Procedia*, 3, pp.110-118.
- Zweig, J., 1998. Five investing lessons from America's top pension fund. *Money*, 27(1), pp.115-18.

## Appendix

### 1. The Legendre Transformation for continuous quadratic functions

We are going to perform a Legendre Transformation by doing some basic calculus. The equation for a Legendre transform  $L(p)$  for a quadratic is:

$$L(p) = \max[xp - f(x)]$$

Where  $p$  is our new dependent variable. If we take the first derivative of the inside of those brackets, we get that  $dL(p)/dx = p - f'(x)$ . To maximize the difference between the two terms  $xp$  and  $f(x)$ , the first derivative must equal zero (i.e.  $p - f'(x) = 0$ ). We can use this to find  $x$  in terms of  $p$ .

We know that  $f(x)$  is quadratic, because it is our scaling function. In general, for any quadratic equation  $y = ax^2 + bx + c$ , where  $a$ ,  $b$  and  $c$  are constants:

$$dy/dx = 2ax + b$$

From this, we know that our derivative  $dL(p)/dx$  should be in the form:

$$dL(p)/dx = p - (2ax + b)$$

We want to maximise  $L(p)$ , which we know is concave (like an upside-down U). The maximum is at the point where the derivative equals zero, so:

$$dL(p)/dx = 0$$

$$\therefore p - (2ax + b) = 0$$

$$\therefore x = (p - b) / (2a)$$

We can plug this into our Legendre Transform equation to get the final result. For every quadratic equation  $ax^2 + bx + c$ , the Legendre Transform equation will be:

$$L(p) = \max[xp - f(x)]$$

$$L(p) = \max[xp - (ax^2 + bx + c)]$$


So finally, by replacing all the  $x$ 's for  $(p - b) / (2a)$ , we find the Legendre Transform to be:

$$L(p) = ((p - b) / (2a)) * p - (a * ((p - b) / (2a))^2 + b * ((p - b) / (2a)) + c)$$


This last line can more or less be copy-pasted into code!

## 2. Examples of fractals in mathematics and nature

1. The Mandelbrot set is perhaps the most famous fractal shape. Here we see a repetition of the same shape when zoomed in to as scale a googol times smaller — i.e. a zoom scale of  $10^0$  vs.  $10^{-100}$ . (Source: <https://www.youtube.com/watch?v=PD2XgQOyCck&t=120s> )




$10^0$




$10^{-100}$

But the main feature of the Mandelbrot Set is probably its “infinite complexity”, as shown in the diagrams below (adapted from Roger Penrose’s 2004 book):




(Clockwise from top left) The magnification factors are: picture *a* is full size, picture *b* is magnified 11.6 times; picture *c* is magnified 168.9 times; and picture *d* is magnified 1042 times.

2. The Sierpinski Triangle and Sierpinski Pyramid — another very famous fractal model. (Source: Wikipedia — [https://en.wikipedia.org/wiki/Sierpinski\\_triangle](https://en.wikipedia.org/wiki/Sierpinski_triangle) )



Sierpinski Triangle



Sierpinski pyramid and its inverse

3. Romanesco cauliflower — an edible vegetable demonstrating a fractal growth pattern. Sometimes called “Romanesco broccoli”. (Source: Wikipedia — [https://en.wikipedia.org/wiki/Romanesco\\_broccoli](https://en.wikipedia.org/wiki/Romanesco_broccoli) )




### 3. Sequential standard deviation for Norway, Sweden and the LIBOR — Variance does not converge

The mathematics of Levy-Stable distributions imply that distributions with a small  $\alpha$ , and therefore high kurtosis, should have infinite variance. This means that any calculation of variance will never converge as we add more and more observations. Of course, in real life, variance cannot be infinite. But we can see that the variance for our three markets shows very slow convergence — much slower than a Normally distributed simulation!

In the graphs below, I show how the variance of the price returns changes as we add more and more observations. The black line is a Normally distributed simulation — the variance for this was taken by measuring the real data's variance on the last day. The second row of graphs shows the same data, but zoomed in to only include the last 1,250 days, which at approximately 250 trading days per year, is almost the same as the past five years.


We see that the graphs appear to be converging — but very slowly. Furthermore, if we zoom in, we again see a very different structure for the real data. Its sequential standard deviation is much “jumper” — it takes it much longer to converge than for the Gaussian, and the value can change rapidly. The implication is this: if we are using standard deviation to calculate price, then our price may vary a lot if simply calculate the standard deviation using different dates.



**4. Mandelbrot's (1963) cotton price graph — an example of power-laws in economic price series**

When data comes from a stable distribution, the empirical distribution function should show a straight line in a log-log plot of price changes plotted against frequency.

Mandelbrot (1963) studied cotton prices and found that price changes followed a power-law distribution. For reference, the log-log graph is included here. The different lines (from  $a^+$  to  $c^-$ ) represent three different time-periods, for which the plus-label shows the price increases and minus-label shows the price falls. Seeing the graph gives a general idea of the power-law distributions that we are looking for in the data.



5. American confidence in banks — the trend in Gallup polls

According to Gallup, a renowned polling organization, the proportion of Americans who express high confidence in banks has been falling steadily since 1979 and the Reagan period of the 80's. It dropped drastically after 2008.

*Americans' Confidence in Banks, 1979-2016 Trend*

Now I am going to read you a list of institutions in American society. Please tell me how much confidence you, yourself, have in each one -- a great deal, quite a lot, some or very little.




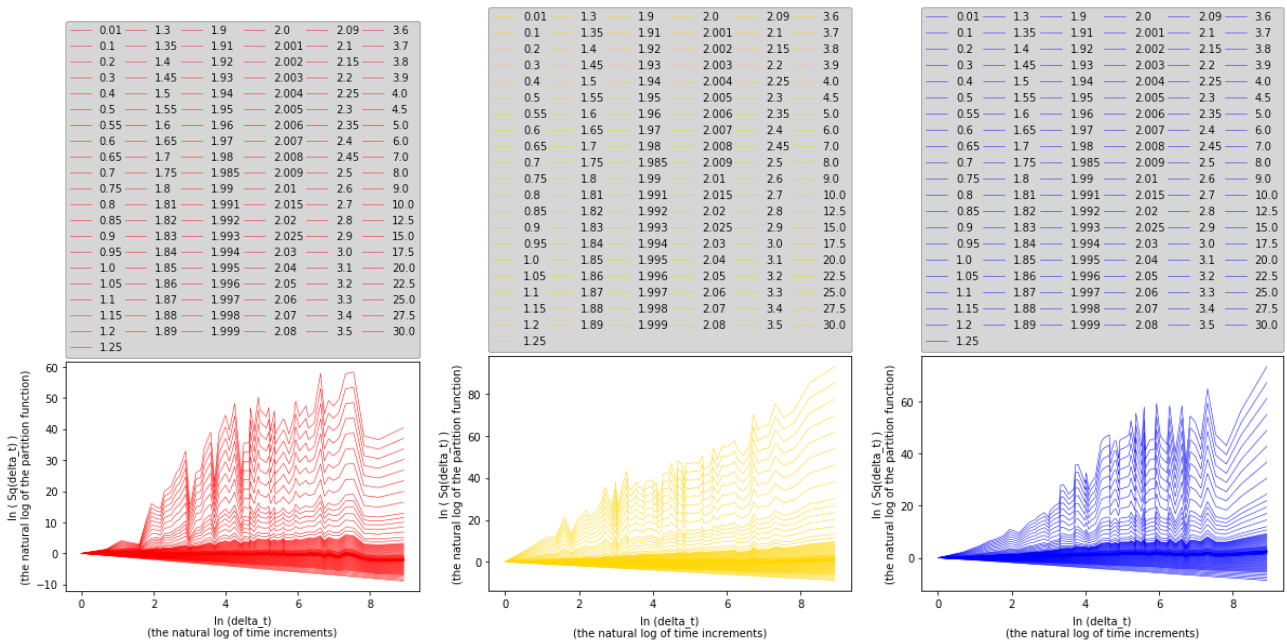
Figure 42 — Source: <https://news.gallup.com/poll/192719/americans-confidence-banks-languishing-below.aspx>

The drop is more pronounced among low-income individuals. Conservatives, liberals and independents show similar levels of low confidence.

**6. Full partition functions for all three markets**

These partition function graphs have 121 lines, because I chose 121 values of  $q$  (121 “raw” moments) with which to measure the scaling function. I include them here for the sake of completeness.


But despite the mess, the scaling relationship is still quite visible. As always, red is for the Norwegian krone, yellow is for the Swedish stock market, and blue is for the 12-month LIBOR in British pounds. The 121 moments are listed in the upper boxes.



**7. Scaling functions for various exchange rate markets — Lots of anti-persistent behaviour**

It would seem that most exchange rates seem to have anti-persistent behaviour — or at least they do against the US dollar. Recall that anti-persistence means a low  $H$  exponent (under the random-walk value of 0.5). This means that price movements are not independent — instead, an up-movement is more likely to be followed by a down-movement.

In the graph below, I show the scaling functions for several currency markets, all trading the local currency against the US dollar. I also added the Swedish stock market’s scaling function for comparison, as it has a persistent  $H$  of 0.542. Every currency shows an anti-persistent  $H$  — all of them are under 0.5.



We might interpret this as evidence that exchange rate markets tend to have a lot of arbitrage behavior. When a price trend occurs, FX traders begin trading in the opposite direction, and the price tends to go back to where it was before.

Exchange rate data was taken from the FRED database. All data spanned approximately thirty years (specifically, 7560 days) between 1989 and 2019.

**8. Scaling after 5 minutes in the Deutschemark (Fisher et al 1997)**

One of the original MMAR papers, which studied the Deutschemark with high-frequency data, suggests that fractal scaling can begin at as short a time interval as 5 minutes (i.e. here,  $10^{2.5}$  seconds). More serious scaling begins at  $10^4$  seconds, which is a little under 3 hours.

