

16th Annual Conference on Finance and Accounting, ACFA Prague 2015, 29th May 2015

Inference of Aggregational Gaussianity in Asset Returns Exhibiting a Paretian-distribution

Imran Amien^{a*}, Kanshukan Rajaratnam^b, Ryan Kruger^b

^aActuarial Science Section, University of Cape Town, Rondebosch, 7700, South Africa

^bDepartment of Finance and Tax, University of Cape Town, Rondebosch, 7700, South Africa

Abstract

Aggregational Gaussianity (AG) has long been considered a stylized fact of empirical asset return distributions. This research links existing work on the stable-Paretian Hypothesis with the Aggregational Gaussianity hypothesis and notes that the two are incompatible. We use simulation to show that under certain conditions, AG can be falsely inferred to hold in a data set exhibiting the stable-Paretian distribution.

© 2015 Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

Peer-review under responsibility of University of Economics, Prague, Faculty of Finance and Accounting

Keywords: Aggregational Gaussianity; Asset returns; Paretian distribution

1. Introduction

Amongst portfolio managers and academics, stochastic models are commonly used to forecast asset returns and prices. One of the most widely-used conventions in asset price modelling is to assume that that returns are log-normally distributed. This is equivalent to assuming that log-returns are normally distributed. It has been observed in many empirical studies that while log-returns of equity returns are non-normal over short timescales (see Kendall (1953), Fama (1965) and Cont (2001)), the distributions become closer to normal when measured over longer timescales (see Kendall (1953) and Eberlein and Keller (1995)). This convergence in distribution towards normality as the timescale increases is referred to as Aggregational Gaussianity (AG) and is well documented in the literature

* Corresponding author.

E-mail address: AMNIMR002@myuct.ac.za

and considered to be a stylized fact of equity returns (Cont 2001). However, empirical distributions tend to have taller peaks and fatter tails than the normal distribution would allow and are called leptokurtic. Mandelbrot (1963), in particular, has shown that the stable-Paretian distribution is a better fit of empirical asset distributions than the normal distribution. In this paper we show that when asset returns exhibit a Paretian distribution, normality tests, under small-sample size conditions, may result in failure to reject the presence of AG.

For the sake of convenience, the terms returns and log-returns will be used interchangeably in this paper although it is the distribution of log-returns that are primarily of interest. It should also be noted that while the most commonly modelled asset classes are stocks/equities of listed companies, this paper will refer to asset return distributions in general, being applicable to any asset class where normality is assumed.

The normal distribution has been studied in detail in statistical literature and its properties are well documented, thus assuming that log-returns are normally distributed simplifies the application of these models. The use of the normality assumption is seen in many widely-used models such as the Black-Scholes and Merton models and others that rely on Brownian Motion (Pascucci, 2011). It is also present indirectly in models that have normally distributed errors around some known mean, such as the Capital Asset Pricing Model (CAPM), multi-factor models, linear regression models and time series models (Cochrane, 2001).

However, the normality assumption is also one of the most widely-contradicted assumptions in the literature. Due to the observed non-normality of returns, many believe that using stochastic models which are reliant on the normality assumption misrepresents the risks in investment decision-making (Taleb 2001). Despite this, normally distributed models continue to see wide-spread use due to their desirable properties and the large volume of literature available on them.

In describing AG, it should be understood that the statistical distribution of asset returns depends implicitly on the timescale used to calculate returns. The relationship between returns and log>Returns is

$$R_t = \frac{S_{t+\Delta t} - S_t}{S_t}$$

$$\Leftrightarrow R_t = \ln S_{t+\Delta t} - \ln S_t \quad (1)$$

where R_t is the return on the asset at time t ,

S_t is the asset price at time t ,

Δt is the timescale over which returns are examined.

Thus, as can be seen from Equation 1, different log returns can be calculated depending on the timescale used (the size of Δt) and thus the statistical distribution of daily returns often differs from that of monthly returns and so on.

Measuring log-returns over longer timescales is equivalent to temporal aggregation (Rydberg 2000), i.e.,

$$\ln S_n - \ln S_0 = (\ln S_n - \ln S_{n-1}) + (\ln S_{n-1} - \ln S_{n-2}) + \dots + (\ln S_1 - \ln S_0), \quad (2)$$

where $n > 0$.

Thus, the terms 'increasing timescale' and 'aggregation' will be used interchangeably in this paper. Rydberg (2000) also argues that since longer timescales can be considered a sum of random variables then, provided its conditions hold, the Central Limit Theorem should come into effect and the distribution of log-returns should tend towards Gaussian as timescale increases.

Many previous studies on markets around the world examined stocks on the DAX (Eberlein and Keller 1995), NYSE (Rydberg 2000), CAC-40 (Herlemont 2003) and S\&\}P (Antypas, Koundouri and Kourogenis 2010) indices

respectively. Many studies used qualitative analysis such as Q-Q plots. There was evidence using Q-Q plots that while returns are severely non-Normal over shorter timescales, kurtosis decreases and tails become thinner as timescale increases. As a result, AG was inferred despite many of these studies using small samples. Aggregating data necessarily lowers sample size. For instance, in a market with 5 trading days a week, 4800 daily observations aggregates to only 20 yearly observations. Small samples can result in incorrect decisions especially since formal tests of normality become unreliable when samples are small. This is one possible explanation for the situation where normality is rejected at shorter timescales where data are plentiful but then is not rejected at longer timescales where data are scarce.

In summary, the recent work by Boavida (2011) suggests that AG cannot always be assumed as stylized fact in a returns data set. Consequently, it may not acceptable to use it as evidence to support the use of normally distributed models. In essence: in data sets that are non-normal at short timescales and do not aggregate to Gaussian, using normally-distributed models, even as approximations, may not be justifiable.

A parallel line of related research has been finding an appropriate distributional fit for asset returns. Mandelbrot (1963) suggested the stable-Paretian distribution to describe asset returns in asset price modelling. The stable-Paretian distribution, sometimes called the Paretian or α -stable distribution, is a 4-parameter multifractal distribution that can be used to describe the high-peaked fat-tailed distributions often found empirically¹.

Of particular importance to this work on the Paretian distribution is the α -parameter ($0 < \alpha \leq 2$) which dictates the heaviness of the tails. Eberlein and Keller (1995) note that for $\alpha < 2$, the Paretian distribution has infinite variance and for $\alpha \leq 1$, the first moment is undefined. For most cases, no closed form exists for the probability density function. These undesirable statistical properties may explain why the stable-Paretian distribution with $\alpha < 2$ is not as widely-used as the normal distribution despite being a more accurate description of asset returns (Fama 1963). For $\alpha = 2$, the Paretian becomes a normal distribution, indicating that the normal distribution belongs to the stable-Paretian family of distributions.

The Paretian distribution is also scale-invariant. This implies that the α -parameter does not change with aggregation. This property is also called summability or stability under addition. Thus, for $\alpha < 2$, the observations from a Paretian distribution cannot aggregate to normality since $\alpha \neq 2$. Thus, if the daily returns data have been sampled from a Paretian distribution with $\alpha < 2$, the data set cannot possess the AG property. Rydberg (2000) has noted that Paretian distributions and AG are incompatible and has suggested that by inferring AG, one can reject the stable-Paretian hypothesis. It follows that, when AG is spuriously inferred, it can lead to an incorrect rejection of the Paretian hypothesis.

In Section 2, we describe our simulation experiment and exhibit the results. This is followed by a discussion of the results in Section 3 and finally, we conclude and recommend next steps in Section 4.

2. Simulation and Results

AG is commonly inferred by noting that for empirical distributions, the peaks flatten and the tails become thinner with aggregation. Bearing in mind that many return distributions around the world can be considered Paretian at short timescales, and that Paretians cannot aggregate to Gaussian, this simulation will show that this criteria is insufficient to infer AG and can, in some cases, causes spurious inference of AG where it cannot be present.

¹ See Frain (2006) for additional literature on the stable-Paretian distribution and related processes.

Two identical simulations are carried out. A number of independent observations from a stable-Paretian distribution are generated. These are considered to be daily log-returns. The daily log-returns are then aggregated sequentially into weekly, monthly, quarterly and yearly log-returns. Q-Q plots and fitted histograms are then produced, and formal tests of normality are run, at each timescale. Where Simulation 1 and Simulation 2 differ is in the size of the data set. In Simulation 1: 120,000 daily returns were generated resulting in 24,000 weeks, 6000 months, 2000 quarters and 500 years. In Simulation 2: 4,800 Daily returns were generated resulting in 960 weeks, 240 months, 80 quarters and 20 years.

This simulation assumed 5 trading days in a week, 4 weeks in a month and 12 months in a year.² The parameters of the Paretian distribution used in this simulation, $\alpha=1.7$ (stability), $\beta=0$ (skewness), $\gamma=2.7$ (spread), $\delta=0.44$ (mean), are taken from Frain (2006) and previously used to test the robustness of tests of normality when the alternative is a Paretian.³

3. Analysis of Results

The results are shown in the figures and tables overleaf.

Figure 1a shows that the distribution of daily returns from Simulation 1 is heavily leptokurtic and many tail observations dot the horizontal axis. Some of these observations have a distance of as much as 25-sigma or greater. As this agrees with empirical observation, it should be apparent why the normal distribution inadequately explains the behaviour of returns. These tail observations have a low probability but a disproportionately high impact; Taleb (2001) has dubbed these the Black Swan events. As the data from Simulation 1 is aggregated, the peak of the distribution flattens and observations cluster around the middle as shown in Figures 1a - 1e. Note the continued prevalence of Black Swans at all timescales; the probability of this occurring for a sample of 500 observations from a normal distribution is negligible. Note that the distribution remains leptokurtic even at higher timescales. It would be expected that were there to be more observations, the yearly timescale would more closely resemble the daily timescale. The Q-Q plots in Figures 1a - 1e and the tests of normality in Table 1 also confirm these findings: they reject normality outright at all timescales.

² i.e. 240 trading days a year. Real stock markets have around 250 trading days a year.

³ The parameters have not been varied as this is merely an illustrative example of how AG can be falsely inferred under small samples.

Fig. 1: Results of Simulation 1 with equivalent of 500 years of data (a) Daily, (b) Weekly, (c) Monthly, (d)




Fig. 2: Results of Simulation 2 with equivalent of 20 years of data (a) Daily, (b) Weekly, (c) Monthly, (d) Quarterly and (e) Annually.




Table 1: Table of p-values from 3 Tests of Normality on Simulation 1 (note: AD – Andersen-Darling).

Test	Daily	Weekly	Monthly	Quarterly	Yearly
Jarque-Berra	0.000	0.000	0.000	0.000	0.000
Shapiro-Wilk	0.000	0.000	0.000	0.000	0.000
AD	0.000	0.000	0.000	0.000	0.000

Table 2: Table of p-values from 3 Tests of Normality on Simulation 2 (note: AD – Andersen-Darling).

Test	Daily	Weekly	Monthly	Quarterly	Yearly
Jarque-Berra	0.000	0.000	0.000	0.007	0.448
Shapiro-Wilk	0.000	0.000	0.000	0.016	0.859
AD	0.000	0.000	0.000	0.025	0.819

Simulation 2 is identical to Simulation 1 with but less data, i.e. 20 years as opposed to 500. It should be noted that 20 years is still generous for some markets and many empirical studies have used less. Figure 2a shows that the distribution of daily returns in Simulation 2 is also leptokurtic. Figures 2a - 2e show that as the data from Simulation 2 is aggregated, the peak flattens and the Black Swans disappear. This is in line with Simulation 1. Beginning at the monthly timescale (Figure 2c), however, changes are noted. As the peak shrinks and the tail observations disappear, the normal distribution appears to give a better fit. By the yearly timescale (Figure 2e), the peak is relatively flat and the Black Swans have disappeared under aggregation. The Q-Q plots appear to confirm the same. Most surprising are the normality Tests in Table 2: by the yearly timescale, all three tests fail to reject the null hypothesis that the data are normally-distributed. There are serious contradictions here. Firstly, Simulation 1 and Simulation 2 are sampled from the same distribution and should have identical statistical properties. Secondly, since the Paretian distribution is scale-invariant, it follows that the data can never aggregate to normal. However, Simulation 2 appears to aggregate to normality while Simulation 1 does not.

The work of Taleb (2001) provides a possible explanation for this. The shape of the Paretian distribution actually remains the same, but many of the distributional properties are not apparent when data are insufficient. Tail events are hidden from the distribution until they occur, while other events will cluster around the middle leading to excess kurtosis.⁴ Taleb (2001) calls this the principle of hidden evidence. This also explains why the daily returns from Simulation 1 have a taller peak and fatter tails than the daily returns from Simulation 2. This principle applies to AG as follows: returns data are plentiful at the daily scale but sparse at the yearly scale where it has been aggregated. The data thus appears to be leptokurtic at lower timescales, and so the Paretian fits well, but as the data are aggregated, the peak flattens and the tails become thin and so the data appears to become normal. However, it is known that a Paretian cannot aggregate to Gaussian. Hence, the yearly data spuriously appears normal because it is analysed over too short a period, such as 20 years. Indeed, when the data are analysed over a longer period, such as 500 years, it is almost certainly non-normal. Thus, this false aggregation to Gaussian under small samples can lead to

□ This is similar to the argument that when observing a man over a 10-year period it is unlikely that he will die (a tail event), but when he is observed over a 100-year period his death becomes almost a certainty. Over 10 years, there is spurious statistical evidence that he is immortal (Taleb , 2007).

the spurious rejection of the stable-Paretian Hypothesis and can appear to justify normally-distributed models which do not take Black Swans into account.

The above simulation has shown how AG can spuriously appear under small samples. This research terms this phenomenon Small Sample Pseudo-Aggregational Gaussianity, or Pseudo-Gaussianity for short. Empirically, returns which are leptokurtic at the daily scale but appear to aggregate to normal may be displaying this phenomenon. Data are most plentiful at the daily timescale and thus more reliable. Thus any returns data set which fails to reject the stable-Paretian Hypothesis at the daily timescale shows strong evidence to reject the AG property.

4. Conclusion

In summary, this paper has linked the existing work on the stable-Paretian hypothesis to the Aggregational Gaussianity hypothesis. The primary contribution of this work has been to show via simulation how a set of independent random variables sampled from a stable-Paretian distribution can spuriously appear to aggregate to Gaussian under small samples. This work suggests that since the stable-Paretian Hypothesis and the AG Hypothesis are incompatible, a failure to reject the stable-Paretian Hypothesis may be evidence to reject AG.

In light of the findings of this research and the recent work of others, it is strongly recommended that existing markets that have long been considered to possess the property the AG, be re-examined. Given that the stable-Paretian distribution is a good empirical fit of many markets around the world, and that this research has shown how Pseudo-Gaussianity can occur, it would be expected that many markets will find evidence against the AG hypothesis in the future.

References

- Boavida, J. (2011). Excess returns and normality. Technical University of Lisbon. Unpublished Masters Thesis.
- Cochrane, J.H. (2001). *Asset Pricing: Revised Edition*. Princeton University Press.
- Cont, R. (2001). Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1:223-226.
- Eberlein, E. & Keller, U. (1995). Hyperbolic distributions in finance. *Bernoulli*, 1(3): 281-299.
- Fama, E. (1963). Mandelbrot and the Stable Paretian Hypothesis. *The Journal of Business*, 36(4): 420-429.
- Fama, E. (1965). The behavior of share market prices. *Journal of Business*, 38: 34-105.
- Frain, J.C. (2006). Small sample power of tests of normality when the alternative is an α -stable distribution. Trinity College Dublin, working paper.
- Herlemont, D. (2003). Futures markets of stock indices: Stylized facts & automated trading. Yats Consulting USA, working paper.
- Kendall, M. (1953). The analysis of economic time series - Part I: prices. *Journal of the Royal Statistical Society*, 116(1): 11-34.
- Mandelbrot, B. (1963). The Variation of Certain Speculative Prices. *Journal of Business*, 36: 394-419.
- Pascucci, A. (2011). *PDE and Martingale Methods in Option Pricing*. Bocconi University Press.
- Rydberg, T. (2000). Realistic statistical modelling of financial data. *International Statistical Review*, 68(3): 233-258.
- Taleb, N.N. (2001). *Fooled by Randomness*. Random House.
- Taleb, N.N. (2007). *The Black Swan*. Random House.