

Does the Compass Rose pattern matter for testing normality?

18/06/03

Jan Annaert

Ghent University and University of Antwerp
W. Wilsonplein 5 D
9000 Ghent
Belgium
jan.annaert@rug.ac.be

Anouk G.P. Claes

University of Antwerp
Prinsstraat 13
2000 Antwerp
Belgium
anouk.claes@ua.ac.be

Marc J.K. De Ceuster*

University of Antwerp
Prinsstraat 13
2000 Antwerp
Belgium
marc.deceuster@ua.ac.be

&

Director Risk Management Deloitte & Touche
Lange Lozanastraat 270
2018 Antwerp
Belgium

Abstract

Some years ago, Crack and Ledoit (1996) discovered a strikingly geometric structure when plotting US stock returns against themselves. Since this pattern, in which lines radiating from the origin pop up, resembles the navigating tool it was named “Compass Rose”. Although authors differ in opinion when explaining the causes of the phenomenon, discreteness of price jumps is unanimously indicated as driver of the structure.

This paper first documents the presence of a Compass Rose Structure within the illiquid Belgian stock market, looking at both individual stocks and stock indices. We then examine whether the presence of a Compass Rose, i.e. the discreteness of prices, affects normality tests. Based on simulated Brownian Motions with rounded price increments, we notice that two commonly used normality tests react differently to discreteness in the underlying data. As the tick size increases, the popular Jarque-Bera test is not able to detect the deviations from normality. The Lilliefors test, however, clearly rejects the normality assumption when the data exhibit tick/volatility ratios in excess of 2.5.

*Corresponding author

1. Introduction

Some years ago, Crack and Ledoit (1996) discovered an intriguing and fairly robust structure within the graphical representation of daily US stock returns, plotted against themselves with a lag of one day. The structure of this phase diagram takes the form of evenly spaced rays emanating from the origin. The most prominent directions are the major directions of the compass. Hence, they labelled this phenomenon the "Compass Rose".

This Compass Rose can not only be observed in daily stock returns but evidence on the existence of a similar pattern in intra-day foreign exchange returns (Gleason, Lee and Mathur, 2000) and in some futures returns (Lee, Gleason and Mathur, 1999) was found. This empirical evidence gave rise to several interesting questions with respect to the explanation of the existence of this pattern. Can it be explained on purely statistical grounds, or are economic arguments needed? All authors agree that discreteness of observed prices (and hence market microstructure) on the one hand, and the observation frequency (a statistical measurement issue), on the other hand, play a prominent role in our understanding of this structure.

It has been shown that discreteness of traded prices can severely bias the estimation of second and higher moments of returns (Gottlieb and Kalay, 1985; Ball, 1988). Although these results have been around for more than a decade, it is remarkable how few people have used the suggested bias-adjustments. The Compass Rose literature, however, did spur a handful of authors to further examine the effects of discreteness on statistical tests. Whereas Krämer and Runde (1997) assessed the size of the Brock-Dechert-Scheinkmann (BDS) chaos test, Fang (2002) studied the influence of the Compass Rose on several random walk tests. Krämer and Runde (1997) show that under the null hypothesis of i.i.d. returns, the true rejection probability exceeds the nominal size severely. This result is important since it casts doubt on the evidence found in favour of chaos in financial markets. Fang (2002) concludes that random walk tests on transaction data with low levels of the tick/volatility ratio generate the correct size across different significance levels. However, the true rejection probability grows as the tick/volatility ratio rises. In fact it increasingly exceeds the nominal size of the test by increasingly wider margins.

In many theoretical applications financial asset prices are assumed to follow a geometric Brownian motion. This assumption implies that prices are lognormally distributed and hence that logarithmic returns are normally distributed. Of course, this theoretical continuous process cannot perfectly match real-world financial market data, as prices change only over discrete intervals, and the existence of minimal tick sizes prohibit a continuous distribution. It can therefore be expected that especially due to price discreteness tests for normality reject the null hypothesis, even despite the fact that the underlying (continuous) distribution is a normal one. In this paper, we investigate in a simulation setting to what extent this presents a problem for two commonly used tests for normality, the Lilliefors test and the Jarque-Bera test.

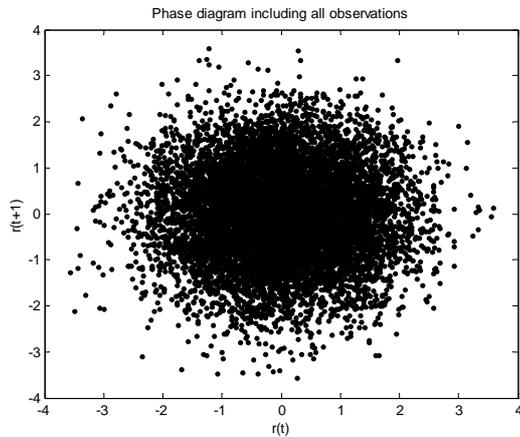
The remainder of the paper is structured as follows. In section 2, we will illustrate the existence of the Compass Rose pattern in daily Belgian stock and index returns. The controversy around the explanations of this phenomenon is discussed in section 3. It turns out that discreteness of price jumps, and hence the microstructure of financial returns is a necessary condition in order to generate a Compass Rose structure. This discreteness of prices may severely hamper statistical estimation (Ball, 1988) and testing (Krämer and Runde, 1997) procedures. Since normality tests are abundantly used in finance, section 4 will examine how price discreteness affects them. The test statistics we will use are the Jarque-Bera test and the Lilliefors test. Simulation experiments show that the Jarque-Bera test is unable to trace down any deviation from normality due to price discreteness. The Lilliefors-test, however, appears to be a more 'powerful' test.

2. Does a Compass Rose pattern exist for Belgian Stock Returns?

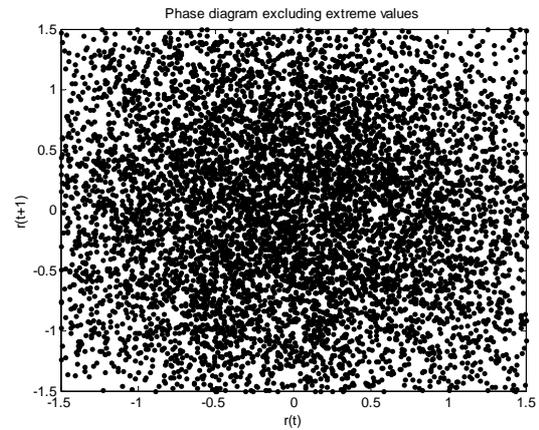
Financial theorists frequently model the dynamics of stock returns as a Geometric Brownian motion (see e.g. Black and Scholes, 1972). This implies that logarithmic stock returns are normally distributed. When generating a Compass Rose based on a simulated standard normal distribution, we obtain Figure 1. We clearly observe a completely random pattern (Panel A) even when zooming in on the centre of the distribution (Panel B). Within Compass Rose literature it is conventional to disregard 'extreme' returns. In the analysis of Crack and Ledoit (1996) the amount of 'extremes' ranged between 10 and 20%. For the purpose of our study we discarded approximately 15%.

Figure 1: Compass Rose based on normally distributed data

Panel A: Phase diagram including all observations



Panel B: Phase diagram excluding extreme values

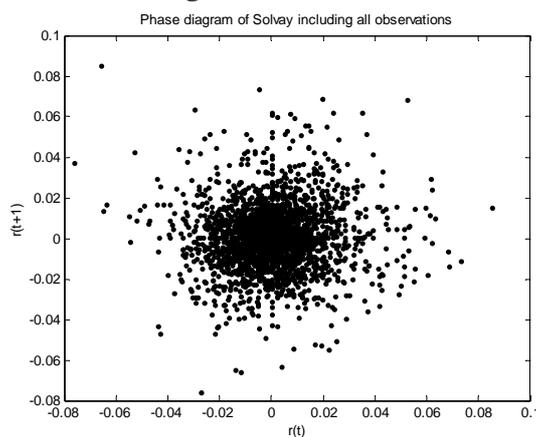


Note: The figures in Panel A and B are identical, except for a scale-adjustment. In Panel A we show the phase diagram of returns, generated from a standard normal distribution. We clearly observe a data point cloud. When zooming in on this diagram, we again observe a completely random pattern, no matter how deep we zoom into the figure.

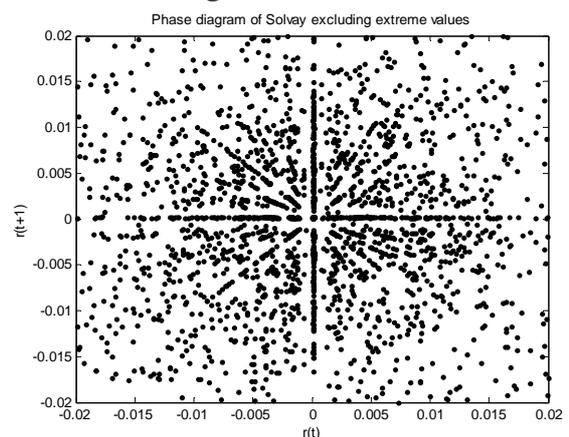
Figure 1 stands in sharp contrast with the Phase Portrait generated from real data. The scatter plot of Solvay's stock returns vis-à-vis its lagged daily returns is depicted in Figure 2. Just as in Figure 1, hardly any patterns can be recognized. The extreme returns, however, mask the Compass Rose.

Figure 2: Scatter plot of Solvay Stock returns (29/12/1989-17/10/2001)

Panel A: Phase diagram of Solvay returns including all observations



Panel B: Phase diagram of Solvay returns excluding extreme values



Note: The figure in Panel B is identical to the one in Panel A, except that we adjust the scale (in Panel B). When zooming in on the graph, the Compass Rose clearly arises. On the horizontal axis of Panel A and B, the return on the Solvay stock on a given day is shown, while on the vertical axis the return of the next day is plotted. Out of 3077 points in the plot of panel A, 440 lie outside the bounds of the plot of panel B.

By zooming in on the center of the distribution (Figure 2 Panel B), lines radiating from the origin turn up. The rays shooting in the major directions of the wind are the thickest. Even when the lag between the daily returns plotted on the axis is increased to 5 days (not shown), the Compass Rose structure remains visible. The diagonal rays, however, do tend to lose clarity. Crack and Ledoit (1996) assert that because the price change from one period to the next is no longer negligible with respect to the price level, the picture is a bit blurred. These findings do not stand-alone but can be found in any particular (Belgian) stock we studied.

3. How can one explain the existence of the Compass Rose structure?

Crack and Ledoit (1996) argue that the Compass Cross is present in all NYSE stocks, but note that the structure need not be revealed within portfolios. They claim that because the effective tick size¹ (with which the asset price changes) becomes negligibly small compared to the price changes, the Phase Portrait will not reveal the Compass Rose pattern. In addition to the requirement that stock price changes are small-integer multiples of some tick size (official or effective tick size), they identify three necessary and sufficient conditions for the appearance of the distinctive Phase Portrait. Hence the structure in the phase diagram arises whenever:

1. 'Daily price changes of Stock XYZ are small relative to the price level ;
2. Daily price changes of Stock XYZ are in discrete jumps of a small number of ticks ; and
3. The price of Stock XYZ varies over a relatively wide range.'(p. 754)

Contrasting the necessity of the three conditions by Crack and Ledoit (1996), Szpiro (1998) claims that all assets moving in discrete changes should have a Compass Rose, including portfolios of stocks. Unlike Crack and Ledoit (1996), the author derives the exact formulation of the return-ratio (R_{t+1}/R_t) that generates the Compass Rose structure. Simulating from this equation, he shows that the previous study was incorrect in claiming that the higher frequency of jumps of a small number of ticks causes the major directions of the Phase Portrait to be more distinctive than the minor directions. Szpiro (2000) indicates that their marking presence is merely an optical illusion. His analysis is based on the locus of the data points generating the Compass Rose. Using this approach, he shows that within a phase diagram, clusters of data-

¹ Effective tick size is defined by both Crack and Ledoit (1996) and Wang, et al. (2000) as 'the normal unit of change'.

points on the minor rays of the Compass Rose are much further apart from each other compared to data-points on the major rays. An increase in the stock price even intensifies this result, to the extent that the human eye will no longer see the clusters on the minor rays lying on one line. Based on these results he rejects the suggestion of Crack and Ledoit (1996) that small jumps in the stock price are more frequent than large jumps.

Yet Wang, et al. (2000) again shed a different light on these results. They reveal Szpiro's (1998) underlying assumption that all potential data points are realised, while in finite time this is of course not possible. They therefore support Crack and Ledoit (1996) and show that where the effective tick size is relatively small, such as in portfolios, the emerging of the Compass Rose phenomenon is not likely. They do, however, extend Szpiro's (1998) analysis by additionally taking the effect of time into account and conclude that the Compass Cross will appear when the following conditions are satisfied:

1. 'A Compass Rose is more likely if the effective tick size is large compared to the standard deviation of the returns;
2. A Compass Rose is more likely as the frequency of observation increases' (p. 121).

However, support for Szpiro's (1998) claim that the 'sole necessary and sufficient condition for the appearance of the micro and the nanostructure of the Compass Rose, is the requirement that price changes take place in discrete jumps' is also found in the literature. Fang (2002) confirms that 'the radically symmetric structure is solely attributable to price discreteness'. Based on simulation techniques he demonstrates the importance of the tick/volatility in the appearance of the Compass Rose phenomenon. Low volatility and large tick size induce more striking radically structures to appear. We now turn back to the data in order to see which theory can be confirmed by the Belgian stock market.

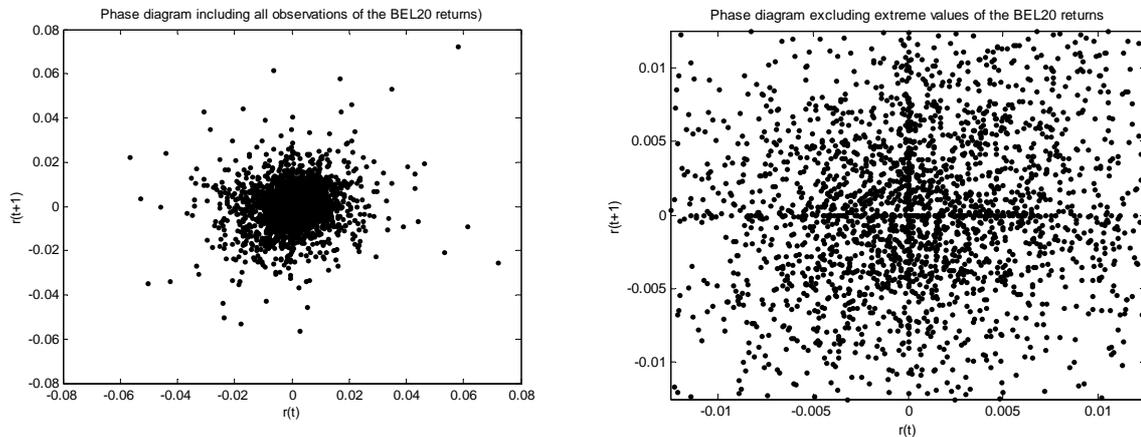
Since the most striking contrast between the two theories is related to the finding of the Compass Rose phenomenon within portfolios, we first concentrate on stock market indices. We both consider the BEL20-index computed by Euronext Brussels, and a self-constructed equally weighted index in which all sample stocks are included. In Figure 3 the phase diagram of the BEL20 is depicted. Panel A again gives the general picture,

including the 'outliers'. After zooming in on the index-returns (Panel B), we do see a pattern emerging, be it that only the main axes of the Compass Rose show up.

Figure 3: Scatter plot of the Belgian Price index BEL20 (29/12/1989-17/10/2001)

Panel A: Phase diagram of BEL20 returns including all observations

Panel B: Phase diagram of BEL20 returns excluding extreme values

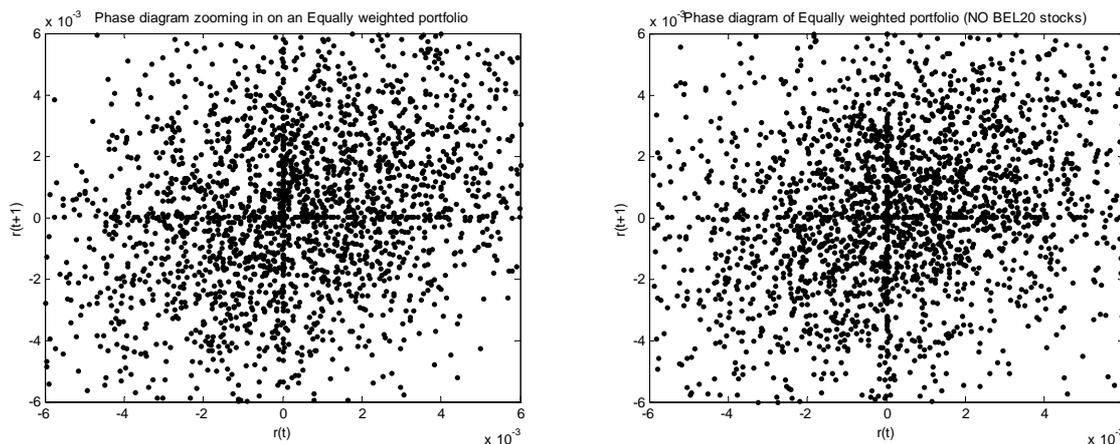


Note: The graph in Panel B is identical to the figure in panel A, except that we adjust the scale. Here a Compass Rose arises as well, be it only its main axis. On the horizontal axis the return on the BEL20 on a given day is shown, while on the vertical axis the return of the next day is plotted. Out of 3077 points in the plot of panel A, 416 lie outside the bounds of the plot of panel B. Not a single zero return was found.

Figure 4: Equally-weighted portfolio of Belgian Stocks

Panel A: Phase diagram zooming in on an Equally Weighted Portfolio (all stocks)

Panel B: Phase diagram zooming in on an Equally Weighted Portfolio (no BEL20 stocks)



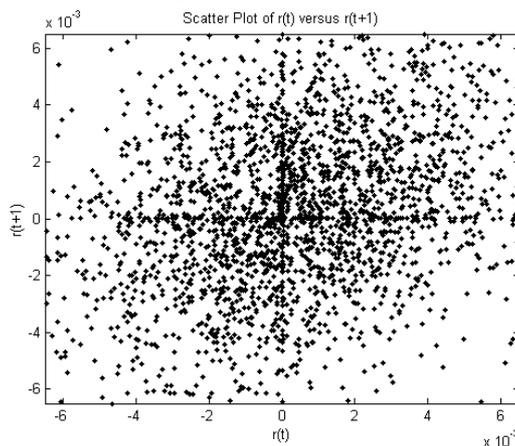
Note: This figure is again scale-adjusted. Here a (fuzzy) Compass Rose arises, be it only its main axis. In Panel A the return on an equally weighted portfolio of 78 Belgian Stocks on a given day is shown on the horizontal axis, while on the vertical axis the return of the next day is plotted. Panel B shows the same picture, except that here the 15 BEL20 stocks are excluded from the portfolio. Out of 3077 points, 427 lie outside the bounds of the plots of Panel A, 414 outside the scope of Panel B.

A similar structure is discovered when plotting the daily returns of the equally weighted portfolio against each other, as can be seen in Panel A of Figure 4. Even when considering higher lagged returns up to 5 days (not shown), the pattern remains. Within this portfolio, the 15 BEL20 stocks with continuous presence over our data sample were included. Since they may influence our results, we also formed a portfolio of the 63 remaining (less liquid) stocks (Figure 4 Panel B). Here, however, the (fuzzy) Compass Cross was found as well.

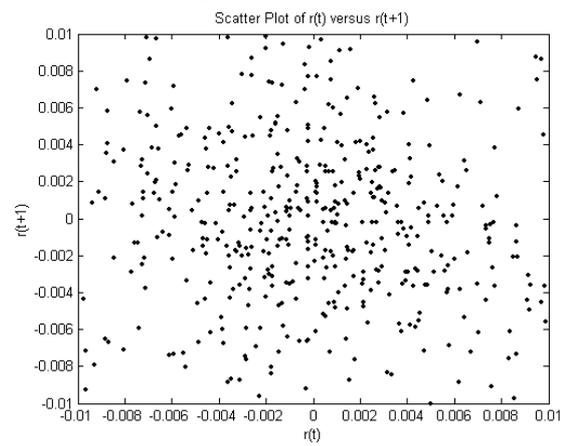
Interested in the time effect as suggested by Wang, et al. (2000), we also examine the frequency of the observations as a possible driver of the Compass Rose phenomenon. First, we therefore go back to the Solvay stock we used for illustrative purposes in section 2. Figure 5 represents the phase diagram for Solvay on a lower frequency rate. In Panel A, the 2 day-returns of the stock are plotted against each other. Compared to Figure 2, the Compass Rose structure has faded and increasingly does so when returns are calculated over more days. At the weekly level (see Panel B), the Compass Rose phenomenon has completely disappeared. This, however, can not be considered conclusive evidence since it may also be due to the lack of data (merely 614 observations).

Figure 5: Scatter plot of Solvay considering the time effect

*Panel A: Phase diagram of Solvay
zooming in on 2-day-returns*



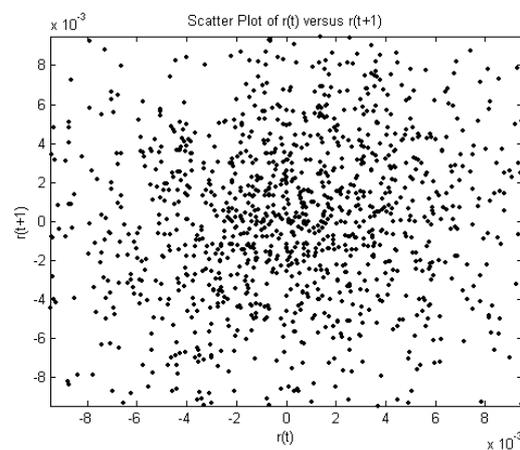
*Panel B: Phase diagram of Solvay
zooming in on week returns*



Note: Here the two-day returns (Panel A) and the week returns (Panel B) are plotted against each other. Compared to previous phase diagrams of Solvay shown, the Compass Rose structure is increasingly mystified and vanishes at the level of weekly returns. In Panel A, out of the 1538 data points originally obtained, 219 fall outside the bounds of the picture. For Panel B, out of 614, 84 extreme values are left out of the graph.

After looking at the time effect at the individual stock level, we now turn to the discussion of this effect on indices. Figure 6 shows the phase diagram of the BEL20 index, in which 2-day-returns are plotted against each other. Already at this (relatively high) frequency level, the Compass Rose structure has become invisible. A similar pattern is revealed by an equally weighted portfolio, both with and without containing the BEL20 stocks. Since the Compass Rose phenomenon in our data set is more likely for high frequency than for low frequency of observations, these Belgian data clearly confirm the second condition for detecting a Compass Rose structure as described by Wang, et al. (2000).

Figure 6: Scatter plot of BEL20 considering the time effect



Note: Here the two-day returns of the BEL20 are plotted against each other. Already at this level, the Compass Rose structure has completely disappeared. Out of the 1538 data points originally obtained, 215 fall outside the bounds of the picture

Based on this section, we now know that the Compass Rose structure is not a typically US phenomenon. Instead it is a revelation of a much more fundamental problem that affects many economic time series: i.e. price discreteness. In the remainder of the paper we will examine the effects of Compass Rose structure or price discreteness on normality tests. We have chosen normality tests since they are very often used in diagnostic checking. We will show that the Jarque-Bera test, which is without doubt the most popular one, has very low power vis-à-vis discretely measured prices. The less known Lilliefors-test, however, correctly rejects the null hypothesis of normality. First, we will describe the Jarque-Bera and Lilliefors test. Next we will consider their power.

4. Commonly used normality tests

Consider observations X_1, \dots, X_n from an unknown distribution $F(x)$. The two test statistics considered will all test the hypothesis H_0 that X_1, \dots, X_n are normally distributed. The Jarque-Bera test (JB) test is a 2-sided normality test with sample mean and sample variance used as estimates of the population mean and variance. Bera and Jarque (1981,1982) derived it as a Lagrange Multiplier test in the context of the Pearson distribution. This statistic is based on estimates of the sample skewness and kurtosis of the normalised data. Under the null hypothesis, the standardised 3rd and 4th moments are asymptotically normal and independent, and the test statistic has an asymptotic Chi-square distribution with two degrees of freedom. Explicitly, the

expression can be formulated as follows: $LM = \frac{n}{6} \frac{u_3^2}{u_2^3} + \left(\frac{u_4}{u_2^2} - 3 \right)^2 \frac{n}{24}$

n being the number of observations, u_2 , u_3 and u_4 the second third and fourth central moment of the observations.

The Lilliefors-test for normality (Lilliefors (1967)) tests whether the observations are drawn from a normal population with unknown mean and variance.

First the data are to be standardised, $Z_j = (X_j - \bar{X})/S$, where \bar{X} is the sample mean and S is the unbiased estimate of the sample standard deviation. The test statistic (D_n) is then the maximum absolute difference between the empirical distribution function of the standardised data and the standard normal distribution function:

$$D_n = \max_{j=1}^n |F_n(Z_j) - \phi(Z_j)|$$

where $F_n(\cdot)$ denotes the empirical distribution function and $\phi(\cdot)$ denotes the standard normal distribution function. We reject the null hypothesis of the sample having the normal distribution (with unspecified mean and variance) at the significance level of α if D_n exceeds the $1-\alpha$ quantile of the Lilliefors-test statistic for normality (Lilliefors (1967)). Using these test statistics, the size of the normality tests under the Compass Rose structure will be examined in section 5.

5. Size and power of normality tests under a Compass Rose structure

Evaluating the Compass Rose effect on the normality tests, we calculate the empirical size and power of both the Jarque-Bera and the Lilliefors test (Table 1). For this purpose, we set up simulation experiments under the Gaussian null hypothesis, using 10,000 replications. The randomly generated prices, with starting value equalling 100, are calculated based on 1,000 observations drawn from a Geometric Brownian motion for which the expected return and the volatility are given in the first and second column in Table 1. Following Fang (2002) daily σ was set to 0.01 (low volatility = 15.81% p.a.) and 0.05 (high volatility = 79.06% p.a.) in order to get representative annual numbers. To investigate the effect of the expected return, annual returns 6.25% and 12.5% were used (again derived from Fang (2002)). Both the Jarque-Bera test and the Lilliefors-test are calculated for 4 different sequentially generated price series: from original continuous prices, prices rounded to the nearest cent (0.01), by an eight of a dollar (0.125), to the nearest dollar (1).² We verify the normality assumption of both continuous and discrete prices. The empirical size and powers of the tests are given at the nominal significance level of 5%.

Table 1: Size and power of the normality tests

Panel A: Jarque-Bera test results

<i>Simulation set up</i>		Continuous Prices	Tick Size of Discrete Prices		
Expected Return	Daily Volatility		0.01	0.125	1
0.0625	0.01	0.0532	0.0533	0.0528	0.0482
0.1250	0.01	0.0481	0.0479	0.0465	0.0460
0.0625	0.05	0.0507	0.0500	0.0517	0.1019
0.1250	0.05	0.0497	0.0505	0.0506	0.0842

Panel B: Lilliefors test results

<i>Simulation set up</i>		Continuous Prices	Tick Size of Discrete Prices		
Expected Return	Volatility		0.01	0.125	1
0.0625	0.01	0.0610	0.0603	0.6065	1
0.1250	0.01	0.0579	0.0594	0.4600	1
0.0625	0.05	0.0659	0.0647	0.1451	0.7873
0.1250	0.05	0.0606	0.0612	0.1230	0.7545

Note: The size and power of the Jarque-Bera test and the Lilliefors test on normality was determined based on 10,000 simulation runs. The randomly sequentially generated prices (starting price equalling 100) were calculated based on 1,000 observations drawn from a Geometric Brownian motion for which the expected return and the volatility are given in the first and second column. Tick sizes are 0.01, 0.125 and 1 (forth, fifth and sixth column).

² All prices were obtained rounding the simulated price downward to the nearest tick, before generating a new observation. Non-sequentially generated price series (first generating the entire price series, afterwards rounding all prices down to the nearest tick) produced similar results for Table 1. Results are available upon request.

Panel A of table 1 reveals the size and power of the Jarque-Bera test. As predicted by both Crack and Ledoit (1996) and Fang (2002), the effect of different expected returns remains marginal. At the 5% significance level, the test statistic does not detect any deviation from normality for low volatility returns, independent of the specified tick size. A similar conclusion can be drawn for the higher volatility returns, except for the test computed on price series rounded to the nearest dollar. Here the rejection rates are 10.19% and 8.42%, for series with respectively low and high expected returns.

Table 1, Panel B displays the test results for the Lilliefors statistic. Again, the differences within the two volatility classes are very small. This normality test, however, appears to be much more sensitive to the rounding of prices. Here deviation from the expected rejection rate of 5% is not big for the tests performed on the continuous prices and prices with a tick size of 1 cent. The rejection rates of discrete prices, rounded to respectively one-eighth of a dollar and 1-dollar, however, clearly indicate non-normality. Since the Compass Rose is more likely to be distinctive for low volatility and price series with large tick sizes, we see from this panel that the Lilliefors-test is able to detect the Compass Cross.

Focussing on the tick/volatility ratio, the results shown in table 1 Panel B are now rearranged and shown in Table 2, in increasing tick/volatility ratio. This table informs us that using this normality test at the 5% level, deviations from normality are not detected, whenever the tick/volatility ratio is less than 1. In contrast to the Jarque-Bera test, the Lilliefors statistic does detect deflection from normality within price series containing an increasing tick/volatility ratio, and thus an intensifying Compass Rose pattern.

Table 2: Lilliefors power results versus tick/volatility

Tick/Volatility Ratio	Expected Return	
	0.0625	0.1250
0.2	0.0647	0.0612
1	0.0603	0.0594
2.5	0.1451	0.1230
12.5	0.6065	0.4600
20	0.7873	0.7545
100	1	1

Note: The power of the Lilliefors test results of Table 1 Tabled in order of increasing tick/volatility ratio.

6. Conclusions

Price changes of stocks typically occur as small-integer multiples of the tick size. This discreteness in stock price changes limits the possible values the ratio of successive returns can take. This phenomenon generates a striking geometrical pattern, commonly referred to as a Compass Rose, when daily stock returns are plotted against the returns of the day before. We reveal the Compass Rose structure within the Belgian illiquid market. The discovery of a Compass Rose for stocks in this less liquid market and moreover on less frequent data, generates evidence that this microstructure effect is not only associated with high frequency data. While the phenomenon is also present within portfolios, it is clearly less ostentatious.

Given this Compass Rose structure, we test the effects of price discreteness on the size of two normality tests, i.e. the Jarque-Bera test and the Lilliefors-test. While the Jarque-Bera test is unable trace to down any deviation from normality as the tick size of prices increases, the Lilliefors test appears to be a more sensitive measure since it does detect the increasing declination.

References

- Ball, C.A., 1988, Estimation Bias Induced by Discrete Security Prices, *Journal of Finance*, 43, 4, 841-865.
- Bera, A.K. and C.M. Jarque, 1981, Efficient Tests for Normality, Heteroscedasticity and Serial Independence of Regression Residuals: Monte Carlo Evidence, *Economics Letters*, 7, 313-318.
- Bera, A. and C. Jarque, 1982, Model Specification Tests: A Simultaneous Approach, *Journal of Econometrics*, 20,59-82.
- Black F. and M. Scholes, 1972, The Valuation of Option Contracts and a Test of Market Efficiency, *Journal of Finance*, 27, 399-418.
- Crack, T.F. and O. Ledoit, 1996, Robust Structure Without Predictability: The "Compass Rose" Pattern of the Stock Market, *Journal of Finance*, 51/2, 751-762.

Fang, Y., 2002, The Compass Rose and random walk tests, *Computational Statistics & Data Analysis*, 39, 299-310.

Gleason, K.C., C.L. Lee and I. Mathur, 2000, An explanation for the Compass Rose pattern, *Economics Letters*, 68, 127-133.

Gottlieb, G. and A. Kalay, 1985, Implications of the Discreteness of Observed Stock Prices, *Journal of Finance*, 40/1, 135-153.

Krämer, W. and R. Runde, 1997, Chaos and the Compass Rose, *Economics Letters*, 54, 113-118.

Lee C.L., K.C. Gleason and I. Mathur, 1999, A Comprehensive Examination of the Compass Rose Pattern in Futures Markets, *The Journal of Futures Markets*, 19/5, 541-564.

Lilliefors, 1967, H.W., On the Kolmogorov-Smirnov test for normality with mean and variance unknown, *Journal of the America Association*, 62, 399-402.

Szpiro, G.G., 1998, Tick Size, the Compass Rose and Market Nanostructure, *Journal of Banking and Finance*, 22, 1559-1569.

Wang, E., R. Hudson and K. Keasey, 2000, Tick size and the Compass Rose: further insights, *Economics Letters*, 68, 119-125.