

Amilon, Henrik; Byström, Hans

Working Paper

The Compass Rose Pattern of the Stock Market: How Does it Affect Parameter Estimates, Forecasts, and Statistical Tests?

Working Paper, No. 2000:18

Provided in Cooperation with:

Department of Economics, School of Economics and Management, Lund University

Suggested Citation: Amilon, Henrik; Byström, Hans (2000) : The Compass Rose Pattern of the Stock Market: How Does it Affect Parameter Estimates, Forecasts, and Statistical Tests?, Working Paper, No. 2000:18, Lund University, School of Economics and Management, Department of Economics, Lund

This Version is available at:

<https://hdl.handle.net/10419/259840>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

The Compass Rose Pattern of the Stock Market: How Does it Affect Parameter Estimates, Forecasts, and Statistical Tests?

Henrik Amilon and Hans NE Byström*

Department of Economics

Lund University

P.O. Box 7082

S-220 07 Lund

Sweden

E-mail: henrik.amilon@nek.lu.se, hans.bystrom@nek.lu.se

Phone: +46-46-222 79 09

Abstract

A "compass rose" pattern sometimes appears when stock returns are plotted against themselves with a one-day lag, since stock prices move in discrete steps. In this paper, we perform a Monte Carlo study on simulated stock price series rounded in different ways to mirror the behavior of stocks on the Stockholm Stock Exchange. We find AR-GARCH parameter estimates to be affected by the discreteness imposed by rounding. Based on the compass rose and the discreteness, we investigate, theoretically and empirically, different possibilities of improving predictions of stock returns. The distributions of the BDS test as well as Savit and Green's dependability index are also influenced by the compass rose pattern. However, throughout the paper, we must impose unrealistically heavy rounding of the stock prices to find significant effects on our estimates, forecasts, and statistical tests.

JEL codes: C15, C22, G19

Keywords: GARCH, discrete prices, forecasts, correlation integral statistics

*The authors are grateful to Professor Tobias Rydén for most helpful comments on how dependences occur in rounded series.

1 Introduction

Typically stock prices, as well as numerous other financial time series, move in small-integer multiples of a minimum "tick size". This discrete nature of stock prices restricts stock returns to take on a limited number of values only, a restriction which is one of the necessary conditions for creating the so-called "compass rose" pattern, a geometrical pattern in a scatter plot of returns versus lagged returns. Crack and Ledoit (1996) were the first to discover this pattern in daily return plots, finding compass rose patterns in all stocks they investigated on NYSE. Crack and Ledoit describe the compass rose pattern as "subjective" and impossible to use for predictive purposes. Contrary to the beliefs of Crack and Ledoit, Chen (1997) demonstrates some evidence of how the compass rose of a stock can be used to improve stock return forecasts. Crack and Ledoit also suggest that the discrete nature of stock prices might affect GARCH estimates as well as statistical tests. So far, no studies of discrete prices and GARCH estimates have been made but Krämer and Runde (1997) show how the compass rose seriously distorts the null distribution of the BDS test (Brock et al. (1996)). The BDS test is a widely used statistical portmanteau test detecting deviations from IID in general and the existence of chaos in particular. Based on their own findings, Krämer and Runde suggest that this test should be used with caution whenever discrete data is used. However, we believe their return series to be highly unrealistic and less suitable for drawing conclusions.

There is more to the compass rose than meets the eye and in this chapter, we use Monte Carlo simulations to test whether the existence of a compass rose and the associated discreteness affect estimates, forecasts, and correlation integral based statistical tests. Possible effects on parameter estimates are investigated as is the use of the compass rose to enhance forecasts in a GARCH framework (as suggested by Chen) which is tested and further developed in a more theoretical framework. We also test how the BDS test and the associated Savit and Green (1991) dependability index are influenced by the discreteness of stock prices. When researchers look for evidence of chaos in the stock market, they often study stock indexes and not individual stocks. However, as shown by Atchinson and White (1996), an aggregation of chaotic processes may very well be non-chaotic. This turns the focus to the study of individual stocks and therefore, it is important to clarify whether correlation integral based tests can be used when examining stock returns with a discrete nature and compass rose patterns.

Section 2 describes the compass rose. In section 3, we investigate how rounding affects AR-GARCH parameter estimates. Section 4 deals with our compass rose enhanced forecasts of stock returns. Section 5 shows how correlation integral based statistical tests can be affected by discrete prices and the compass rose. Concluding remarks are found in section 6.

2 The Compass Rose

A compass rose pattern sometimes appears when returns from a financial time series are plotted against lagged returns. The pattern is characterized by several evenly spaced lines radiating from the origin of the graph; the thickest lines pointing in the major directions of a compass. Crack and Ledoit (1996) were the first to recognize and explain this phenomenon, induced by discreteness in stock price data. They found that the compass rose appears clearly if the stock in question satisfies three conditions:

1. Daily stock price changes are small relative to the price level.
2. Daily stock price changes occur in discrete jumps of a small number of ticks.
3. The stock price varies over a relatively wide range.

The derivation of these three conditions is straightforward and the details can be found in Crack and Ledoit (1996)¹. The patterns in Figures 1a-d show when the compass rose appears. In Figure 1a, we plot the daily log-returns (1977-1984) from a stock, "Atlas Copco A Fria", listed on the Stockholm Stock Exchange.² The compass rose pattern can be seen clearly. In Figures 1b-d, we plot simulated returns based on simulated prices from an AR-GARCH model (see section 3) fitted to the same Atlas Copco stock. In Figure 1b, it can be seen how the original simulated returns show no pattern. This is radically changed in Figure 1c, where we have rounded the simulated prices to mirror the behavior on the Stockholm Stock Exchange³. This pattern is very similar to the one in Figure 1a. It is important to notice, however, that the discreteness induced by the official tick size alone is not the cause of the striking patterns in Figure 1a and Figure 1c. In Figure 1d, we plot the same rounded returns as in Figure 1c but randomly permuted (scrambled). The compass rose pattern disappears, which indicates that something must be added to the discreteness to create the pattern. The remaining cross-shaped pattern is merely a consequence of the large number of zero returns that remain in the scrambled series.

¹Szpiro (1998) shows how a more rigorous treatment of the compass rose makes the assumption that asset prices change by only small amounts relative to the price level superfluous.

²Crack and Ledoit derive their three conditions on basis of the assumption of percentage returns. In this paper, we use the more widely used log-returns and we see how the compass rose is not a consequence of the use of percentage returns.

³On the Stockholm Stock Exchange, the prices are allowed to take only certain discrete values. The level of discreteness depends on the price level; below 5 SEK the smallest price jump is 0.01 SEK, between 5 SEK and 10 SEK the jump is 0.05 SEK, between 10 and 50 SEK it is 0.1 SEK, between 50 SEK and 500 SEK it is 0.5 SEK, and above 500 SEK it is 1 SEK.




Figure 1: The compass rose pattern in: (a) Atlas Copco, (b) simulated returns, (c) stock exchange rounded simulated returns, and (d) scrambled stock exchange rounded simulated returns. All time series are of the length 2000.

3 Parameter Estimation

Crack and Ledoit (1996) hypothesize that Autoregressive Conditional Heteroscedasticity (ARCH) models (Engle (1982)) might be influenced by the compass rose. Chen (1997) uses information contained in the compass rose to improve forecasts from an ARMA-GARCH model. He does not study the effect of the discreteness on his parameter estimates, though, which is exactly what we try to investigate with Monte Carlo simulations in this section.

In order to study the effects of the discreteness, return series differing in nothing but their level of discreteness must be found. Since this is not easily achieved with empirical data, we have chosen to simulate series which we then round in order to get discrete time series⁴. The

⁴When we use the word discrete, we mean discrete in state space, not in time space. All statistical tools used are based on common discrete-time models.

simulated series used in this section all come from the same AR-GARCH model:

$$\begin{aligned} r_t &= \alpha_0 + \alpha_1 r_{t-1} + \alpha_2 r_{t-2} + \varepsilon_t \\ \sigma_t^2 &= \phi_0 + \phi_1 \varepsilon_{t-1}^2 + \phi_2 \sigma_{t-1}^2, \end{aligned} \tag{1}$$

where r_t is the stock log-return, σ_t^2 is the conditional variance of ε_t , and where the parameter vector $\theta = \{\alpha_0, \alpha_1, \alpha_2, \phi_0, \phi_1, \phi_2\} = \{5.2 \cdot 10^{-4}, 0.12, -0.053, 2.2 \cdot 10^{-5}, 0.11, 0.83\}$ comes from an AR-GARCH estimation of the "Atlas Copco A fria" stock on the Stockholm Stock Exchange over the period 1977-1990⁵. Atlas Copco is one of the largest companies on the exchange and the stock is liquid and representative for the time period. Moreover, it also shows a distinct compass rose pattern. Our choice of simulated series is different from the treatment in Krämer and Runde (1997). Their simulated stock price series come from IID normally distributed returns with means and variances chosen without resemblance to stocks in the real world. Therefore, any relationships between the level of rounding and the potential effects detected are not useful.

Using the model in (1), we simulate 1000 return series, each being 2000 observations long. The return series are exponentiated and 1000 price series are computed, all with the starting value of 50. These "original" price series are not discrete and do not show any compass rose pattern. To get price series with varying degrees of discreteness, we round the original price series; either *Integer Rounding* to the nearest integer or *Stock Exchange Rounding* to mirror the stocks on the Stockholm Stock Exchange, as described in section 2. After calculating log-returns from these series we have *three* \times 1000 return series originating from non-rounded prices, realistically discretized prices, and heavily rounded prices.

Our purpose is to study how the parameter estimates from the *three* \times 1000 return series change with the *three* different degrees of discreteness. The model we apply to these series is exactly the AR-GARCH model in (1) and therefore, we expect the estimated parameter vectors, $\hat{\theta}$ for the original series, $\hat{\theta}_S$ for the stock exchange rounded series, and $\hat{\theta}_I$ for the integer rounded series, to come rather close to θ . However, it is important to note that the AR-GARCH model we use is misspecified when applied to the rounded series, since ε_t cannot be continuous if r_t is discretely distributed around the conditional mean.

In Table 1, we present some statistics such as the mean, the 95% confidence interval, the minimum, and the maximum, of the *three* \times 1000 AR-GARCH estimates. The AR-GARCH parameters were jointly estimated by Maximum Likelihood methods (the BHHH algorithm)⁶. From Table 1, we see that the small sample distributions of some of the parameters in $\hat{\theta}$, $\hat{\phi}_0$ and

⁵Totally, 3489 daily observations.

⁶Estimating the AR parameters prior to the GARCH parameters, did not considerably affect the estimates and these results are therefore not presented.

Table 1: Parameter estimates and likelihood values. $\hat{\theta}$, $\hat{\theta}_S$, and $\hat{\theta}_I$ denote the parameter vector estimates of the original, the stock exchange rounded, and the integer rounded series, respectively. The empirical α -percentile is denoted I_α . Small numbers are 95% confidence intervals.

	$\hat{\phi}_0 \cdot 10^5$	$\hat{\phi}_1$	$\hat{\phi}_2$	$\hat{\alpha}_0 \cdot 10^4$	$\hat{\alpha}_1$	$\hat{\alpha}_2$	$\ln L$	
$\hat{\theta}$	mean	2.314 (2.273,2.355)	0.110 (0.109,0.111)	0.826 (0.824,0.828)	5.342 (5.103,5.582)	0.118 (0.117,0.120)	-0.054 (-0.056,-0.053)	3.502 (3.500,3.505)
	$I_{0.025}$	1.258	0.075	0.757	-2.395	0.073	-0.103	3.419
	$I_{0.975}$	3.815	0.147	0.882	12.819	0.165	-0.010	3.579
	min.	0.910	0.056	0.709	-8.080	0.037	-0.135	3.367
	max.	5.471	0.175	0.908	16.914	0.191	0.014	3.624
$\hat{\theta}_S$	mean	2.370 (2.328,2.412)	0.107 (0.106,0.108)	0.828 (0.826,0.830)	5.372 (5.131,5.612)	0.108 (0.107,0.110)	-0.052 (-0.053,-0.050)	3.493 (3.491,3.496)
	$I_{0.025}$	1.267	0.073	0.757	-2.385	0.062	-0.098	3.414
	$I_{0.975}$	3.928	0.144	0.885	12.791	0.155	-0.008	3.568
	min.	0.908	0.053	0.689	-8.246	0.033	-0.130	3.364
	max.	5.495	0.179	0.908	16.881	0.178	0.023	3.607
$\hat{\theta}_I$	mean	2.539 (2.484,2.595)	0.091 (0.090,0.093)	0.846 (0.844,0.848)	5.546 (5.289,5.802)	0.048 (0.044,0.051)	-0.042 (-0.044,-0.042)	3.433 (3.429,3.437)
	$I_{0.025}$	1.181	0.042	0.769	-3.261	-0.092	-0.092	3.267
	$I_{0.975}$	4.589	0.135	0.918	13.296	0.128	0.001	3.537
	min.	0.621	0.047	0.681	-9.071	-0.215	-0.122	3.040
	max.	10.022	0.162	0.959	17.885	0.152	0.027	3.571

$\hat{\phi}_2$, are biased, since the true parameters, ϕ_0 and ϕ_2 , are outside the 95% confidence intervals of the means of $\hat{\phi}_0$ and $\hat{\phi}_2$. The distribution of the parameter estimates in $\hat{\theta}_S$ and $\hat{\theta}_I$ change with the level of discreteness. The confidence intervals are wider and the means of the estimates differ compared to $\hat{\theta}$, which is most obvious for $\hat{\alpha}_1$. It is important to emphasize that a combination of two effects is observed. First, in infinite samples, we know that $\hat{\theta}$ will converge to θ , but since we have rounded the series, we do not know to what $\hat{\theta}_S$ and $\hat{\theta}_I$ will converge asymptotically. Second, as mentioned above, the continuous-state AR-GARCH framework is not the proper one to use in modeling discrete return series. Hence, we cannot separate the effect of the rounding on the asymptotics from the effect of the misspecified AR-GARCH model; we can only investigate the compound effect of discreteness on parameter estimation⁷.

To say something about how the distributions of $\hat{\theta}$, $\hat{\theta}_S$, and $\hat{\theta}_I$ differ, we analyze the deviation vectors, $\tilde{\theta}_S = \hat{\theta} - \hat{\theta}_S$ and $\tilde{\theta}_I = \hat{\theta} - \hat{\theta}_I$. Comparing individual estimates from the stock exchange

⁷An extension of the GARCH framework to handle cases where the return series are discrete is given in Amilon (1999).

Table 2: Parameter estimate and likelihood value differences. $\tilde{\theta}_S = \hat{\theta} - \hat{\theta}_S$ and $\tilde{\theta}_I = \hat{\theta} - \hat{\theta}_I$, where $\hat{\theta}$, $\hat{\theta}_S$, and $\hat{\theta}_I$ denote the parameter vector estimates of the original, the stock exchange rounded, and the integer rounded series, respectively. The empirical α -percentile is denoted I_α . Small numbers are 95% confidence intervals.

		$\tilde{\phi}_0 \cdot 10^5$	$\tilde{\phi}_1$	$\tilde{\phi}_2$	$\tilde{\alpha}_0 \cdot 10^4$	$\tilde{\alpha}_1$	$\tilde{\alpha}_2$	$\widetilde{\ln L}$
$\tilde{\theta}_S$	mean	-0.056 (-0.065,-0.048)	0.003 (0.002,0.003)	-0.002 (-0.003,-0.002)	-0.029 (-0.042,-0.016)	0.010 (0.010,0.010)	-0.003 (-0.003,-0.002)	0.009 (0.009,0.009)
	$I_{0.025}$	-0.360	-0.005	-0.017	-0.460	-0.000	-0.012	0.001
	$I_{0.975}$	0.223	0.011	0.012	0.411	0.023	0.005	0.020
$\tilde{\theta}_I$	mean	-0.225 (-0.261,-0.188)	0.018 (0.017,0.019)	-0.019 (-0.021,-0.018)	-0.203 (-0.249,-0.157)	0.070 (0.068,0.074)	-0.012 (-0.013,-0.011)	0.069 (0.066,0.073)
	$I_{0.025}$	-1.538	-0.003	-0.089	-1.621	0.015	-0.041	0.014
	$I_{0.975}$	0.878	0.062	0.018	1.464	0.210	0.012	0.225

rounded and integer rounded series with the original non-rounded series, we get $two \times 1000$ deviation vectors. In Table 2, we look at the distributions of these deviations. The confidence intervals are much wider for $\tilde{\theta}_I$ compared to $\tilde{\theta}_S$. The absolute values of the mean deviations increase with the level of discreteness, and the sign of the mean deviations differ between parameters, but do not change with the level of rounding. According to the 95% confidence intervals of the mean deviations, all means of $\tilde{\theta}_S$ and $\tilde{\theta}_I$ are quite small but significantly different from zero.

In Table 1, further evidence of the effect of discrete prices on model estimation is given by the log-likelihood value statistics from the Monte Carlo simulations. The log-likelihood values are lower for discrete time series, in particular for heavily rounded series with a mean log-likelihood value clearly below the original series. Looking at the distribution of the log-likelihood deviations in Table 2, it becomes even clearer that the rounded series have significantly lower log-likelihood values. In fact, for the integer rounded series, the likelihood values from the original series are never smaller than the corresponding likelihood values from the rounded series. Since we have rounded the series, we would expect the log-likelihood values to change. In this case, they are almost always smaller, which is a consequence of the rounded series no longer fulfilling the assumptions underlying the model in (1).

4 Enhanced Forecasts

Another way of assessing the importance of discrete stock prices is to study whether the compass rose can be used to enhance forecasts. Crack and Ledoit (1996) argue that the compass rose contains no information that can be used for predictive purposes. Chen (1997) on the other hand

tries to contradict this empirically by incorporating the compass rose pattern in his ARMA-GARCH return forecasts (not modifying his estimates, though) and in this way improving his forecasts⁸. In summary, Chen argues that there are feasible regions (the rays radiating from the origin) and unfeasible regions (the white spaces between the rays) in the compass rose. Chen's enhanced procedure simply tries to force his ARMA-GARCH forecasts to lie within the feasible regions. When a forecast is being made, Chen goes through the constantly updated historical compass rose pattern created from all historical return pairs (r_t, r_{t+1}) and replaces the *original forecasted return* pair $(r_{today}, r_{forecast})$ by the pair from the feasible region closest to the actual forecasted pair in an Euclidean sense. Here, we call such enhanced forecasts *rose-enhanced forecasts*.

In addition to Chen's method, we suggest an alternative enhancement considering the fact that prices move in discrete jumps and can only take on certain values. Instead of looking at the historical return pattern, we suggest that more emphasis should be put on the feasibility of the forecasted returns, implied by the feasibility of the forecasted prices. Since, at a certain point in time, only a limited number of discrete prices are feasible forecasts for tomorrow's stock price, we replace our actual forecasted price by one rounded to the nearest possible price⁹. From this forecasted price, we then calculate an associated forecasted return, called *tick-enhanced forecast*, that can be compared to the original forecast and the rose-enhanced forecast as well as the realized out-of-sample return. This is in contrast to Chen's method, which adjusts the forecasts to belong to the set of realized returns, although these are not attainable at all price levels.

Since we do not incorporate the discreteness in our parameter estimates, it can be shown how the enhancements will affect the performance. Suppose that we have estimated our AR-GARCH parameters from a discretized time series. When forecasting, tomorrow's return, r_t , is assumed to be a continuous stochastic variable symmetrically distributed around the conditional mean return, m_t , which is described by the AR parameters. Since r_t is actually discrete, so is m_t . Suppose further that the possible states at time t are $ix, i = 0, \pm 1, \pm 2, \dots$, where, for simplicity, we let $x = 1$ ¹⁰. Let p_j denote the conditional probability that state j occurs. Because of the symmetry of r_t , it follows that if $j < m_t < j + 1/2$, then $p_j > p_{j+1} > p_{j-1} > \dots > p_{j+N} > p_{j-N}$, where we limit the possible states to $2N + 1$, where N states are larger than j , and N states are smaller than j . We can let $N \rightarrow \infty$, or we can truncate the distribution by letting

⁸Chen's results are based on out-of-sample predictions of six stocks during a certain period of time, by using the mean absolute error and the root mean squared error as performance measures. No statistical significance of his results are presented.

⁹The possible prices depend on the level of discreteness; either we allow only integer prices or also the rounded prices mirroring the Stockholm Stock Exchange prices.

¹⁰Since we are using log-returns calculated from discrete prices, this is only approximately correct.




Figure 2: The conditional density function of r_t around m_t , and its discrete counterparts $p_0, p_{\pm 1}$, and $p_{\pm 2}$.

$p_{j \pm N} = \sum_{l=1}^{\infty} p_{j \pm (N+l)}$. In the same way, if $j - 1/2 < m_t < j$, then $p_j > p_{j-1} > p_{j+1} > \dots > p_{j-N} > p_{j+N}$. We disregard the low probability events of $m_t = j$, and $m_t = j \pm 1/2$, although they can be treated in a similar manner. Figure 2 illustrates the notation for $j = 0, N = 2$, and $m_t = 0.25$.

Now, if our forecast is f , the Mean Squared Error will become:

$$\text{MSE} = \sum_{i=j-N}^{j+N} p_i (f - i)^2 = E[(f - i)^2] = V[f - i] + (E[f - i])^2 = V[i] + (f - E[i])^2,$$

which is minimized for $f = E[i]$ which, in turn, should be close to m_t . Thus, the original forecasted return will always have the lowest MSE, although this forecast is not feasible. Let us

consider the Mean Absolute Error on the interval $k \leq f \leq k + 1$:

$$\begin{aligned} \text{MAE} &= \sum_{i=j-N}^{j+N} p_i |f - i| = \sum_{j-N}^k p_i (f - i) + \sum_{k+1}^{j+N} p_i (i - f) = \\ &= f \left(\sum_{j-N}^k p_i - \sum_{k+1}^{j+N} p_i \right) + \sum_{k+1}^{j+N} p_i i - \sum_{j-N}^k p_i i. \end{aligned} \quad (2)$$

If $\sum_{j-N}^k p_i < \sum_{k+1}^{j+N} p_i$, the MAE is minimized when f is chosen as large as possible on the interval and consequently, if $\sum_{j-N}^k p_i > \sum_{k+1}^{j+N} p_i$, then f should be as small as possible. What happens when we move from one interval to the next? If we let $f_1 = k - \epsilon$ and $f_2 = k + \epsilon$, with ϵ being small, it is easily shown that

$$\text{MAE}_2 - \text{MAE}_1 = 2\epsilon \left(\sum_{j-N}^{k-1} p_i - \sum_{k+1}^{j+N} p_i \right). \quad (3)$$

In the case of $p_j > p_{j+1} > p_{j-1} > \dots > p_{j+N} > p_{j-N}$, (3) is negative if $k \leq j$, and positive if $k > j$. The MAE is therefore minimized when $k = j$, and according to (2), f should be chosen as small as possible, that is $f = j$. Similarly, when $p_j > p_{j-1} > p_{j+1} > \dots > p_{j-N} > p_{j+N}$, (3) is negative for $k < j$, and positive otherwise. The minimizing state is then $k = j - 1$, and (2) is, once more, minimized for $f = j$. In a MAE sense, choosing the most probable feasible outcome is favored compared to the expected forecast. Thus, our enhanced forecast will outperform the original one, using the MAE as a performance measure.

Let us examine our theoretical results empirically in the following way: AR-GARCH return forecasts are nothing but the returns forecasted from the mean process in (1). Using the $two \times 1000$ different parameter sets $\hat{\theta}_S$ and $\hat{\theta}_I$, estimated in section 3, we forecast returns by using the different forecasting methods. As a test sample, we use $two \times 1000$ different return series simulated from (1), each 1000 observations long, and rounded in the two different ways¹¹. For each of the $two \times 1000$ parameter sets, Root Mean Squared Errors (RMSE) as well as Mean Absolute Errors (MAE) are calculated over the test sample for the different forecasting methods and different roundings. In our setting, the RMSE and the MAE are:

¹¹Each of the $two \times 1000$ parameter sets (estimated by using the observations in the estimation period) is tested on exactly one of the $two \times 1000$ (out-of-sample) test series. The model parameters in each set are kept constant over the entire test period.

$$\begin{aligned}\text{RMSE} &= \left[\frac{1}{1000} \sum_{i=1}^{1000} (r_i - \hat{r}_i)^2 \right]^{\frac{1}{2}} \\ \text{MAE} &= \frac{1}{1000} \sum_{i=1}^{1000} |r_i - \hat{r}_i|,\end{aligned}$$

where r_i is the actual return at day i , \hat{r}_i is the forecasted return at day i , and the number of days in the test period is equal to 1000.

In Table 3, we present the means and the 95% confidence intervals of the ratios of the enhanced and the original forecast errors. Starting with the MAE statistic, we can see how the mean ratios are generally close to one but that our tick-enhanced method clearly outperforms the original non-enhanced method; the mean of the ratio is significantly smaller than one at the 5% level for both the exchange rounded and the integer rounded series. Even for the RMSE statistics, the mean ratios are close to one, but otherwise the results are somewhat reversed and the original method seems to outperform the tick-enhanced one; for the exchange rounded series, the mean ratio is significantly larger than one at the 5% level, while for the integer rounded series the mean ratio is not significantly different from one at the 5% level. The results are qualitatively the same for the rose-enhanced forecasts; a mean ratio significantly smaller than one for the MAE, and a mean ratio significantly larger than one for the RMSE. These results confirm our theoretical conclusions about the choice of evaluation measure; the MAE measure indicates gains from using our enhanced method while the RMSE seems to favor the original method. It is also worth noticing how the relative performance of the enhanced methods improves when we round more heavily; throughout Table 3, we see that the heavier we round, the smaller are the error ratios. Finally, if the tick-enhanced method were to be directly compared with the rose-enhanced method, significant results indicating a dominance of our tick-enhanced method over Chen's rose-enhanced method would most probably be found; in three cases out of four, the mean ratios are smaller for the tick-enhanced method.

The results depend, a priori, on the choice of performance measure, so which should be chosen? If constraining one's forecasts to belong to the set of possible outcomes is believed to be reasonable, then the forecasts should be adjusted to be feasible, as we have done, and an evaluation measure favoring a feasible forecast, such as the MAE, should be chosen. It is obvious that the continuous AR-GARCH framework is not the proper one to apply to discrete return series, but we see no theoretical or empirical justification why Chen's adjustments should result in better forecasts. However, the discreteness of the returns should already be incorporated in the parameter estimates to fully improve forecasts, an issue treated in Amilon (1999).

Table 3: Mean estimates and 95% confidence intervals of the out-of-sample forecast error ratios. Each statistic is computed from a test sample of the length 1000. Small numbers are 95% confidence intervals of the mean estimates.

	Exchange Rounding	Integer Rounding
$\hat{E} \left[\frac{MAE_{tick-enhanced}}{MAE_{original}} \right]$	0.9951 (0.9949, 0.9953)	0.9738 (0.9728, 0.9748)
$I_{0.025}$	0.9876	0.9319
$I_{0.975}$	1.0014	0.9955
$\hat{E} \left[\frac{MAE_{rose-enhanced}}{MAE_{original}} \right]$	0.9993 (0.9992, 0.9994)	0.9841 (0.9834, 0.9848)
$I_{0.025}$	0.9945	0.9537
$I_{0.975}$	1.0036	1.0000
$\hat{E} \left[\frac{RMSE_{tick-enhanced}}{RMSE_{original}} \right]$	1.0055 (1.0052, 1.0058)	0.9995 (0.9989, 1.0001)
$I_{0.025}$	0.9956	0.9738
$I_{0.975}$	1.0171	1.0151
$\hat{E} \left[\frac{RMSE_{rose-enhanced}}{RMSE_{original}} \right]$	1.0031 (1.0028, 1.0034)	1.0012 (1.0007, 1.0017)
$I_{0.025}$	0.9946	0.9781
$I_{0.975}$	1.0127	1.0144

5 Correlation Integral Statistics

In this section, we will examine how the distribution of some correlation integral based statistics, defined below, are influenced by the fact that prices only move in discrete ticks. We are primarily concerned with two issues: Do the null distributions of the test statistics change, when applied to return series originating from discrete prices? How are the powers of these test statistics affected when trying to detect explicit time series dependences in such return series?

5.1 The BDS Test and Savit and Green's Dependability Index

Grassberger and Procaccia (1983) introduced a quantity called the correlation integral, in order to identify structures and dependences in data series. In the case of time series analysis, the correlation integral is computed by first forming m -histories from the time series considered:

$$\mathbf{x}(t) = (r(t), r(t-1), \dots, r(t-m+1)) = (x_1(t), x_2(t), \dots, x_m(t)).$$

If the distance between the k th components of two m -histories, $\mathbf{x}(t)$ and $\mathbf{x}(s)$, is defined as

$$l_k(t, s) = |x_k(t) - x_k(s)|, \quad k = 1, 2, \dots, m,$$

then the correlation integral at embedding dimension m and tolerance ϵ can be expressed as

$$C_m(\epsilon) = \frac{1}{N_{pair}} n(l_1 \leq \epsilon, \dots, l_m \leq \epsilon), \quad (4)$$

where N_{pair} is the total number of pairs, and $n(l_1 \leq \epsilon, \dots, l_m \leq \epsilon)$ is the number of pairs with all components within ϵ apart.

The BDS test is based on the observation that for an IID sample,

$$C_m(\epsilon) = [C_1(\epsilon)]^m.$$

The identity should be understood in a statistical sense. Brock et al. (1996) derived a normalization factor, $V_m(\epsilon)$, in order to make a correct statistical quantification of the departure from IID. More specifically, they showed that the BDS statistic

$$W_m(\epsilon) = \frac{C_m(\epsilon) - [C_1(\epsilon)]^m}{V_m(\epsilon)}, \quad (5)$$

converges in distribution to $N(0, 1)$, as the time series become infinitely long, for $\epsilon > 0$, and $m > 1$, under the null hypothesis of IID. With the possibility of bootstrapping small sample distributions as in Efron (1979), the asymptotic properties are nowadays of less importance. The BDS test has proved to be quite successful in finding departures from IID in a number of Monte Carlo studies (see, for example, Brock et al. (1991)), but the test gives no hints of at what time lags there are dependences in the data, information that is most useful in time series modeling and analysis.

Savit and Green (1991) filled this gap by introducing a dependability index computed from the correlation integrals in different embedding dimensions. As seen from (4), $C_m(\epsilon)$ is nothing but the joint probability of two m -histories being no more than ϵ apart in all their Cartesian components, that is $C_m(\epsilon) = \Pr(l_1 \leq \epsilon, \dots, l_m \leq \epsilon)$. This holds under the assumption of time-invariance or stationarity, so that we can compare pairs from different parts of the time series. It is also possible to form the conditional probabilities of two observations being close, given that their m -histories are close:

$$\Pr(l_1 \leq \epsilon | l_2 \leq \epsilon, \dots, l_m \leq \epsilon) = \frac{\Pr(l_1 \leq \epsilon, \dots, l_m \leq \epsilon)}{\Pr(l_2 \leq \epsilon, \dots, l_m \leq \epsilon)} = \frac{C_m(\epsilon)}{C_{m-1}(\epsilon)},$$

since $\Pr(l_2 \leq \epsilon, \dots, l_m \leq \epsilon) = \Pr(l_1 \leq \epsilon, \dots, l_{m-1} \leq \epsilon)$ by construction. In the same way,

$$\Pr(l_1 \leq \epsilon | l_2 \leq \epsilon, \dots, l_{m-1} \leq \epsilon) = \frac{C_{m-1}(\epsilon)}{C_{m-2}(\epsilon)}.$$

If $r(t)$ does not depend on $r(t - m + 1)$, then

$$\Pr(l_1 \leq \epsilon | l_2 \leq \epsilon, \dots, l_m \leq \epsilon) = \Pr(l_1 \leq \epsilon | l_2 \leq \epsilon, \dots, l_{m-1} \leq \epsilon),$$

that is, $C_m(\epsilon)/C_{m-1}(\epsilon) = C_{m-1}(\epsilon)/C_{m-2}(\epsilon)$. It is now possible to define a dependability index:

$$\delta_{m-1}(\epsilon) = 1 - \frac{C_{m-1}^2(\epsilon)}{C_m(\epsilon)C_{m-2}(\epsilon)} \quad (6)$$

for $m > 1$ ($C_0 \equiv 1$), which is zero (in a statistical sense), if there is no dependence of $r(t)$ on the lag $t - m + 1$. In contrast to the BDS test, the delta indexes provide information on at what time lags there are dependences causing the rejection of IID. An asymptotic distribution for the δ 's under the null hypothesis of IID is not available, but can easily be estimated by a bootstrap procedure.

5.2 Rounding of IID Series

In order to investigate the ability of the correlation integral based statistics to pick up the micromarket dependences caused by the rounding of prices we do the following Monte Carlo simulation: First, we generate 1000 return series of the length of 2000 from $N(\mu, \sigma)$, with the same unconditional moments as implied by the AR-GARCH coefficients in section 3, that is $\mu = 5.57 \cdot 10^{-4}$, and $\sigma = 0.019$. After exponentiating to prices, rounding and calculating log-returns as described in section 3, we have *three* $\times 1000$ return series originating from non-rounded prices, realistically discretized prices, and heavily rounded prices. An integer rounding of prices around and below 50 must be regarded as unrealistic, at least for modern financial data. Henceforth, we denote these series R, RS, and RI.

In Table 4, we report some distributional statistics of the W 's and the δ 's, such as the mean, standard deviation, skewness, kurtosis, minimum, maximum, and the 2.5% and 97.5% percentiles. The embedding dimensions are $m = 2, 3, 4$, and 5, corresponding to time lags 1 to 4, and the tolerance parameter ϵ equal to the standard deviation of each time series, as is often suggested.¹²

For R, the distribution of the BDS statistics is close to $N(0, 1)$, as expected for sample sizes of this magnitude. The δ 's also seem to be normally distributed around zero, indicating no dependences on past lags. In the case of RS, the picture is very much the same. Both the BDS statistics and the deltas are normally distributed, but the distribution appears to be shifted somewhat upward, which is especially notable for W_5 , and δ_2 . The percentiles and the mean for W_5 are -1.74, 2.35, and 0.16 as compared to -1.87, 2.10, and -0.03 for the R statistics. For realistically discretized data, the rounding mechanism is picked up by the correlation integral statistics, but the changes in the distributions are not very severe. This is certainly not true when examining RI. Neither the W 's nor the δ 's are normally distributed, and the means as well as the critical values are heavily shifted upward for all m , although most heavily when $m = 2$.

The δ 's are all large, indicating dependences, but our simulated series have no explicit time dependences. How is this possible? The reason is that the rounded series are not time-invariant,

¹²Throughout this paper, we report our results for $\epsilon = \sigma$ only. We have also examined the cases of $\epsilon = 0.25\sigma, 0.5\sigma, 1.5\sigma$, and 2.0σ , with similar findings.

Table 4: Distributional properties of the BDS statistics $W_m(\epsilon)$, and Savit and Green's dependability index $\delta_{m-1}(\epsilon)$ for the R, RS, and RI series, for $m = 2, 3, 4$, and 5, and with $\epsilon = \sigma$, the standard deviation of each time series. The α -percentile is denoted by I_α . Note that the entries corresponding to the δ 's are multiplied by 100, *except* for the entries corresponding to skewness and kurtosis.

		W_2	W_3	W_4	W_5	δ_1	δ_2	δ_3	δ_4
R	mean	-0.04	-0.02	-0.02	-0.03	-0.02	0.01	-0.02	-0.03
	std.dev.	0.99	0.99	1.01	1.01	0.46	0.49	0.53	0.54
	skewness	0.10	0.13	0.23	0.29	0.07	0.18	0.12	0.08
	kurtosis	3.11	2.97	3.10	3.27	3.12	3.08	2.97	2.95
	minimum	-3.57	-3.01	-2.97	-2.82	-1.70	-1.51	-1.48	-1.56
	maximum	3.34	3.55	4.14	4.60	1.51	1.67	1.89	1.93
	$I_{0.025}$	-1.88	-1.88	-1.92	-1.87	-0.89	-0.90	-1.02	-1.06
	$I_{0.975}$	2.07	1.97	2.05	2.10	0.92	1.04	1.06	0.98
RS	mean	0.07	0.13	0.15	0.16	0.03	0.06	0.03	0.01
	std.dev.	1.00	1.00	1.02	1.04	0.46	0.50	0.53	0.56
	skewness	0.08	0.11	0.22	0.30	0.06	0.17	0.16	0.06
	kurtosis	3.12	2.94	2.95	3.07	3.10	2.90	2.90	3.16
	minimum	-3.20	-2.76	-2.57	-2.38	-1.43	-1.52	-1.51	-1.60
	maximum	3.32	3.46	4.04	4.41	1.48	1.63	1.82	2.30
	$I_{0.025}$	-1.88	-1.74	-1.84	-1.74	-0.85	-0.85	-0.96	-1.10
	$I_{0.975}$	2.12	2.08	2.32	2.35	0.99	1.08	1.09	1.11
RI	mean	2.42	3.06	3.48	3.84	1.47	1.14	0.91	0.81
	std.dev.	2.43	2.98	3.42	3.85	1.69	1.35	1.16	1.04
	skewness	0.96	1.13	1.27	1.42	1.41	1.42	1.39	1.17
	kurtosis	3.46	3.81	4.26	4.80	4.47	5.01	5.69	5.10
	minimum	-2.01	-2.01	-1.97	-1.76	-1.00	-1.23	-1.42	-1.51
	maximum	11.08	13.80	17.14	19.99	8.47	7.35	7.11	5.43
	$I_{0.025}$	-0.91	-0.69	-0.60	-0.58	-0.45	-0.56	-0.74	-0.82
	$I_{0.975}$	8.24	10.48	12.23	14.42	5.83	4.84	3.96	3.65

which is assumed in the derivation of the correlation integral statistics. At a given price, a stock will only jump a certain amount of ticks, giving rise to a finite number of return realizations. As the stock price evolves, the number of probable realizations will change, increasing if the stock goes up and decreasing otherwise (if the tick size is the same for all prices). In addition, the non-stationarity in returns is linked to the price level, which creates new complicated dependences. Their appearance is best illustrated by an example.

Let $r_1, r_2 \in N(0, \sigma)$ denote the returns at times $t = 1$ and $t = 2$, let $\sigma = 0.019$, and let the starting price be $P_0 = 100$. The prices at $t = 1$ and $t = 2$ are then $P_1 = 100 \exp(r_1)$ and $P_2 = 100 \exp(r_1 + r_2)$, the rounded prices $\bar{P}_1 = [100 \exp(r_1)]$ and $\bar{P}_2 = [100 \exp(r_1 + r_2)]$, where $[\cdot]$ denotes integer rounding, and the resulting rounded returns are $\bar{r}_1 = \ln(\bar{P}_1/P_0)$ and $\bar{r}_2 = \ln(\bar{P}_2/\bar{P}_1)$. Suppose we want to calculate the probabilities of $\bar{r}_1 = 0$, and $\bar{r}_2 = \ln(101/100)$. The probability that $\bar{r}_1 = 0$ is equal to:

$$\begin{aligned} \Pr(\bar{r}_1 = 0) &= \Pr(\bar{P}_1 = 100) = \Pr\left(100 - \frac{1}{2} \leq P_1 < 100 + \frac{1}{2}\right) = \\ &= \Pr\left(\ln\left(\frac{100 - \frac{1}{2}}{100}\right) \leq r_1 < \ln\left(\frac{100 + \frac{1}{2}}{100}\right)\right). \end{aligned}$$

The joint probability that $\bar{r}_1 = 0$ and $\bar{r}_2 = \ln(101/100)$ is given by:

$$\begin{aligned} \Pr\left(\bar{r}_1 = 0, \bar{r}_2 = \ln\left(\frac{101}{100}\right)\right) &= \Pr(\bar{P}_1 = 100, \bar{P}_2 = 101) = \\ &= \Pr\left(\ln\left(\frac{100 - \frac{1}{2}}{100}\right) \leq r_1 < \ln\left(\frac{100 + \frac{1}{2}}{100}\right), \ln\left(\frac{101 - \frac{1}{2}}{100}\right) \leq r_1 + r_2 < \ln\left(\frac{101 + \frac{1}{2}}{100}\right)\right). \quad (7) \end{aligned}$$

However, the probability that $\bar{r}_2 = \ln(101/100)$ is given by the following:

$$\Pr\left(\bar{r}_2 = \ln\left(\frac{101}{100}\right)\right) = \Pr\left(\frac{\bar{P}_2}{\bar{P}_1} = \frac{101}{100}\right) = \sum_k \Pr\left(\bar{P}_1 = k, \bar{P}_2 = k \frac{101}{100}\right). \quad (8)$$

Since we use integer rounding, the probabilities in the sum are zero unless $k = 100, 200, \dots, \infty$. Besides, the probability that $\bar{P}_1 \geq 200$ is almost zero, so (8) reduces to

$$\Pr\left(\bar{r}_2 = \ln\left(\frac{101}{100}\right)\right) \simeq \Pr\left(\bar{P}_1 = k, \bar{P}_2 = k \frac{101}{100}\right) \Big|_{k=100} = \Pr\left(\bar{r}_1 = 0, \bar{r}_2 = \ln\left(\frac{101}{100}\right)\right).$$

If the rounded returns \bar{r}_1, \bar{r}_2 were independent, then

$$\Pr\left(\bar{r}_1 = 0, \bar{r}_2 = \ln\left(\frac{101}{100}\right)\right) = \Pr(\bar{r}_1 = 0) \times \Pr\left(\bar{r}_2 = \ln\left(\frac{101}{100}\right)\right),$$

Table 5: BDS test statistics and autocorrelations in returns and squared returns for lags 1-4, for series no. 347 in R, RS, and RI, denoted r , rs , and ri .

	AC in returns								AC in squared returns			
	W_2	W_3	W_4	W_5	ac_1	ac_2	ac_3	ac_4	ac_1	ac_2	ac_3	ac_4
r	0.18	0.33	0.64	0.69	0.02	-0.01	0.02	-0.03	-0.01	0.01	0.03	0.01
rs	0.37	0.56	0.79	0.85	0.01	-0.01	0.02	-0.04	-0.01	0.02	0.03	0.02
ri	3.59	3.59	4.08	4.47	-0.04	-0.00	0.01	-0.02	-0.02	-0.01	0.01	0.01

which is obviously not satisfied here. It is these kind of dependences, which could be more or less pronounced in different time series, that are detected by the correlation integral statistics¹³. The rounded returns are no longer IID. In Table 5, we show the BDS test statistics, and the autocorrelation in returns and squared returns, for a typical series in R, RS, and RI. We see that the BDS test strongly rejects integer rounded IID returns, a finding not discovered by just examining the autocorrelations, as is the common approach in econometric analysis.

To summarize, the rounding of prices has two effects. It makes the rounded return series time-variant, and it introduces complicated dependences in the series. The BDS test correctly rejects the rounded returns as IID variables, but the lag-dependences identified by the δ 's may be spurious, since the rounded series are nonstationary.

5.3 Rounding of AR-GARCH Series

We turn to exploring the power of our test statistics when applied to our simulated AR-GARCH return series from section 3, denoted G, GS, and GI. We have established that the correlation integral statistics are distorted by the nonstationarities and dependences introduced by price rounding. Here, we are interested in the ability of detecting explicit time series dependences. In Table 6, we report the minimum, the maximum, and the 95% confidence interval of the test statistics, together with the frequencies of rejecting the null hypothesis of rounded IID observations, at the 5% significance level, when it is false. In the power tests, we are using the critical values from the simulations of R, RS, and RI in Table 4.

Somewhat surprisingly, no sign of upward shifts is distinguishable in the confidence intervals

¹³The probability in (7) can be calculated by numerically integrating

$$\frac{1}{\sqrt{2\pi}} \int_{r_1 = \ln \frac{100-1/2}{100}}^{\ln \frac{100+1/2}{100}} \exp\left(-\frac{r_1^2}{2\sigma^2}\right) \left(\Phi\left(\frac{\ln\left(\frac{101+\frac{1}{2}}{100}\right) - r_1}{\sigma}\right) - \Phi\left(\frac{\ln\left(\frac{101-\frac{1}{2}}{100}\right) - r_1}{\sigma}\right) \right) dr_1.$$

With the numbers chosen here, $\Pr(\bar{r}_1 = 0) \simeq 0.2076$ and $\Pr(\bar{r}_1 = 0, \bar{r}_2 = \ln(\frac{101}{100})) = \Pr(\bar{r}_2 = \ln(\frac{101}{100})) \simeq 0.0703$.

Table 6: Distributional properties of the BDS statistics $W_m(\epsilon)$, and Savit and Green's dependability index $\delta_{m-1}(\epsilon)$ for the G, GS, and GI series, for $m = 2, 3, 4$, and 5, and with $\epsilon = \sigma$, the standard deviation of each time series. The α -percentile is denoted by I_α . Note that the entries corresponding to the δ 's are multiplied by 100, *except* for the power_{0.05}'s.

		W_2	W_3	W_4	W_5	δ_1	δ_2	δ_3	δ_4
G	minimum	2.99	3.28	4.23	4.69	1.39	0.16	0.25	0.06
	maximum	14.54	16.41	18.30	20.54	7.10	4.81	4.19	3.76
	$I_{0.025}$	4.08	5.28	6.19	6.85	1.93	1.18	0.94	0.53
	$I_{0.975}$	10.81	13.23	14.96	16.74	5.22	4.07	3.33	2.96
	power _{0.05}	1.000	1.000	1.000	1.000	1.000	0.985	0.961	0.904
GS	minimum	2.93	3.48	4.17	4.86	1.34	0.26	0.37	-0.12
	maximum	14.14	15.98	17.30	19.24	6.92	4.60	4.09	3.75
	$I_{0.025}$	3.80	5.09	5.99	6.80	1.82	1.14	0.90	0.52
	$I_{0.975}$	10.48	12.84	14.99	16.48	5.10	4.02	3.35	3.00
	power _{0.05}	1.000	1.000	1.000	1.000	1.000	0.977	0.948	0.851
GI	minimum	2.18	3.25	4.06	4.31	1.16	0.05	0.37	-0.74
	maximum	13.70	18.16	22.28	26.37	9.81	9.66	8.46	7.62
	$I_{0.025}$	3.64	5.02	6.09	6.99	1.95	1.21	1.11	0.69
	$I_{0.975}$	10.63	13.62	16.30	18.67	7.19	6.07	5.31	4.64
	power _{0.05}	0.228	0.233	0.233	0.192	0.076	0.093	0.089	0.071

of the BDS statistics when we increase the level of discreteness. The δ 's, on the other hand, show more of the previous positive shifts, at least for the integer rounded GI. Because of the upward shifts present in RS and RI, in combination with the tendency of downward shifts in GS and GI, the power of the test statistics deteriorates, due to the rounding effects. The empirical rejection frequencies of the BDS statistics are 100% for both G and GS. The δ 's clearly also identify dependences at all time lags, as should be the case for GARCH series. The ability of detecting lagged dependences weakens as m increases, and also when comparing G and GS. This should not be confused with any superiority of the BDS test over Savit and Green's dependability index. The former is a true portmanteau test, answering the question if a time series is IID or not. The latter, on the other hand, determines whether including additional lags raises the conditional probability of two observations being close, given that their m -histories are close. Deciding whether a time series is IID or not just by examining a certain lag, rather than using information at all lags, is quite a different task. The BDS test and the dependability indexes therefore give different, but complementary, information usable in time series analysis.

When examining GI, the power of the test statistics falls dramatically to around 23% for

the W 's and to 8% for the δ 's. Trying to determine whether an integer rounded time series (at the price levels examined here) has any explicit time series dependences based on the null distributions of the test statistics would most likely lead to the wrong conclusion. It is somewhat surprising that the power of W_2 and δ_1 are so different, since the numerator of W_2 in (5) is equal to $C_2\delta_1$. Obviously, this is related to the influence of the denominator of the BDS test statistic¹⁴.

Let us examine the standardized residuals from the maximum likelihood estimations of the AR-GARCH series, denoted S, SS, and SI. As shown in Brock et al. (1996), the asymptotic properties of the BDS test are the same, whether the test is applied to IID series or residuals from linear (and some nonlinear) stochastic models. This nuisance parameter free property of the BDS test is not valid for GARCH residuals, which is clearly visible in Table 7. The test statistics are still normally distributed, but with smaller variances, resulting in narrower confidence intervals (see Hsieh (1989)).

The distributional properties of S and SS are again very similar, confirming that realistic rounding of prices does not change the null distribution of the correlation integral based statistics to a very large extent. Once again, these similarities are lost when examining SI. We have upward shifts in the distributions, giving rise to totally different confidence intervals for both the W 's and the δ 's. Furthermore, there is no longer only a scaling factor of the standard deviation separating the distributions of RI and SI, as is the case between the distributions of R and S. If one wishes to investigate the standardized residuals from a heavily rounded return series, there seem to be no shortcuts. A proper Monte Carlo simulation, similar to ours, must be performed in order to extract the critical values to compare with the test statistics of the residuals of one's GARCH model. Even so, the power of such a test would most probably be quite low.

6 Summary and Conclusions

Throughout this chapter, we try to shed some light on the consequences of the trade induced compass rose pattern in stock returns. We examine how AR-GARCH parameter estimates change with the level of rounding in stock prices, and find that the distribution of the estimates differ, in particular for the AR(1) parameter. The differences arise from two effects. First, when rounding the series, we may change the dynamics of the processes, and second, the rounded series no longer fulfil the assumptions underlying the continuous-state AR-GARCH model. Obviously, the higher the level of discreteness, the more pronounced will either of the two effects be.

Further, we show theoretically that incorporating discretization in return forecasts (not in

¹⁴When investigating the power of $C_2\delta_1$, compared to the integer rounded null distribution of the same quantity, the power raises to 12.1%, but is still far from 22.8%.

Table 7: Distributional properties of the BDS statistics $W_m(\epsilon)$, and Savit and Green's dependability index $\delta_{m-1}(\epsilon)$ for the S, SS, and SI series, for $m = 2, 3, 4$, and 5, and with $\epsilon = \sigma$, the standard deviation of each time series. The α -percentile is denoted by I_α . Note that the entries corresponding to the δ 's are multiplied by 100, *except* for the entries corresponding to skewness and kurtosis.

		W_2	W_3	W_4	W_5	δ_1	δ_2	δ_3	δ_4
S	mean	-0.02	-0.08	-0.09	-0.10	-0.01	-0.01	-0.01	-0.02
	std.dev.	0.87	0.75	0.68	0.63	0.40	0.45	0.49	0.55
	skewness	0.17	0.16	0.15	0.20	0.13	0.12	0.15	0.01
	kurtosis	2.75	2.88	2.98	3.07	2.71	2.79	3.00	2.83
	minimum	-2.45	-2.07	-2.02	-1.82	-1.15	-1.30	-1.75	-1.78
	maximum	3.22	2.87	2.18	2.07	1.31	1.29	1.56	1.63
	$I_{0.025}$	-1.61	-1.46	-1.37	-1.30	-0.75	-0.93	-0.93	-1.08
	$I_{0.975}$	1.68	1.40	1.33	1.23	0.77	0.86	1.00	1.05
SS	mean	-0.01	-0.07	-0.09	-0.10	-0.01	-0.07	-0.02	-0.04
	std.dev.	0.88	0.77	0.69	0.64	0.41	0.45	0.50	0.54
	skewness	0.10	0.18	0.22	0.25	0.06	0.04	0.11	0.09
	kurtosis	2.83	3.02	3.16	3.19	2.79	2.73	3.04	2.87
	minimum	-2.31	-2.00	-2.22	-1.84	-1.14	-1.35	-1.67	-1.49
	maximum	3.38	3.31	2.72	2.49	1.41	1.17	1.55	1.77
	$I_{0.025}$	-1.71	-1.48	-1.37	-1.32	-0.79	-0.92	-0.95	-1.09
	$I_{0.975}$	1.69	1.43	1.31	1.20	0.79	0.82	1.02	1.02
SI	mean	1.10	1.53	1.86	2.12	0.63	0.72	0.70	0.57
	std.dev.	1.63	1.95	2.23	2.48	0.99	1.08	1.03	0.99
	skewness	0.49	0.82	0.97	1.09	0.84	1.18	1.30	1.27
	kurtosis	3.33	3.44	3.62	3.87	4.72	4.53	5.08	5.74
	minimum	-4.21	-2.50	-2.21	-2.00	-3.77	-1.37	-1.36	-1.37
	maximum	6.89	8.32	10.61	12.71	4.31	5.23	4.95	5.68
	$I_{0.025}$	-1.65	-1.46	-1.17	-1.02	-0.80	-0.76	-0.69	-0.89
	$I_{0.975}$	4.85	6.24	7.25	8.29	3.11	3.43	3.44	3.05

estimation), as outlined in section 4, improves the performance in an MAE sense, while the opposite holds using the RMSE as an evaluation measure. Simulations reveal that the out-of-sample performance is better, significantly at the 5% level, when using the MAE measure, and we argue for the use of this measure if wishing to favor feasible forecasts.

We also investigate how the distributions of some correlation integral statistics change when applied to rounded return series, and residuals from such series. Our findings suggest that the effects on the distributions are small, provided the return series come from realistic roundings of prices, such as those present at the Stockholm Stock Exchange. Only when we investigate heavy rounding, most likely uncommon in modern financial markets, do the null distribution of the test statistics change remarkably. This is contrary to the findings of Krämer and Runde (1997), who discovered large changes in the null distribution of the BDS test already at low roundings (tick size of 0.1), due to their highly unrealistic simulated stock prices. In section 5, we explain why these changes occur. The rounding of prices makes the corresponding return series time-variant, and introduces dependences in the series. The rounded IID returns are no longer IID. The distributions of the correlation integral statistics therefore change, because of the dependences and the nonstationarities in the rounded data.

The main conclusion is that the effects of discrete prices are small, at least for a discretization comparable to the one present at the Stockholm Stock Exchange. Investigating series with higher tick size to price ratios, such as those present in low-priced stocks, the use of statistical models and tests based on state-continuity can be questioned.

Our results are based on time series of the length 2000. It may be the case that the rounding effects are more severe when examining shorter time series. Caution should also be taken when less traded stocks are examined. The "effective" tick size chosen by market participants can then be larger than the official tick size.

References

- Amilon, H. (1999), "GARCH Estimation and Discrete Stock Prices", Working paper, Lund University.
- Atchison, M., and White, M. (1996), "Disappearing Evidence of Chaos in Security Returns: A Simulation", *Quarterly Journal of Business and Economics*, 35(2), 21-37.
- Brock, W., Dechert, W., Scheinkmann, J., and LeBaron, B. (1996), "A Test for Independence Based on the Correlation Dimension", *Econometric Reviews*, 15(3), 197-235.
- Brock, W., Hsieh, D., and LeBaron, B. (1991), *Nonlinear Dynamics, Chaos and Instability:*

Theory and Evidence, M.I.T. Press, Cambridge.

- Chen, A.-S. (1997), "The Square Compass Rose: The Evidence from Taiwan", *Journal of Multinational Financial Management*, 7, 127-144.
- Crack, T., and Ledoit, O. (1996), "Robust Structure Without Predictability: The "Compass Rose" Pattern of the Stock Market", *Journal of Finance*, June, 751-762.
- Efron, B. (1979), "Bootstrap Methods: Another Look at the Jackknife", *Annals of Statistics*, 7, 1-26.
- Engle, R. (1982), "Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of U.K. Inflation", *Econometrica*, 50, 987-1007.
- Grassberger, P., and Proccacia, I. (1983), "Measuring the Strangeness of Strange Attractors", *Physica D*, 9, 189-208.
- Hsieh, D. (1989), "Testing for Nonlinear Dependence in Daily Foreign Exchange Rates", *Journal of Business*, vol. 62, no. 3, 339-368.
- Krämer, W., and Runde, R. (1997), "Chaos and the Compass Rose", *Economics Letters*, 54, 113-118.
- Savit, R. and Green, M. (1991), "Time Series and Dependent Variables", *Physica D*, 50, 95-116.
- Szpiro, G. (1998), "Tick Size, the Compass Rose and Market Nanostructure", *Journal of Banking and Finance*, 22, 1559-1569.