Computational Statistics & Data Analysis 39 (2002) 299–310

COMPUTATIONAL STATISTICS & DATA ANALYSIS

www.elsevier.com/locate/csda

The compass rose and random walk tests

Yue Fang *

Lundquist College of Business, 1208 University of Oregon, Eugene, OR 97403, USA

Received 1 August 2000; accepted 1 July 2001

Abstract

The recent discovery of the compass rose pattern (Crack and Ledoit J. Finance 51(2) (1996) 751) has sparked considerable interest among researchers. This paper explores the significance of the effect of the compass rose pattern on random walk tests and measures to what extent its influence may limit the performance of test statistics. We show that in general, the asymptotic theory of test statistics is invalid for transactions data. However, Monte Carlo simulations indicate that the impact of the pattern, measured by the empirical size, is visible for moderate size samples only when the tick/volatility ratio is above some threshold, a condition that is readily met with intraday but not daily or weekly returns.

Keywords: Autocorrelation; Heteroscedasticity; Monte Carlo; Price discreteness; Random walk; Variance ratio

1. Introduction

Crack and Ledoit (1996) first documented the existence of the compass rose pattern from an analysis of phase diagrams, where daily returns are plotted against themselves with a one-day lag. They found that the graphical representation takes the form of evenly spaced rays emanating from the origin with the most prominent rays pointing in several major directions, a structure which they termed the compass rose. Crack and Ledoit postulated that the pattern is caused by the discrete jumps in prices. The notion that the compass rose indeed is a consequence of price discreteness is further supported by studies of Kramer and Runde (1997) and Szpiro (1998)

* Tel.: +1-541-346-3265.
E-mail address: yfang@darkwing.uoregon.edu (Y. Fang).

0167-9473/02/$ - see front matter
PII: S 0167-9473(01)00063-9

using different data sets. In this paper, we investigate the implication of the compass rose pattern on random walk tests.

The random walk model is fundamental to modern financial economics and many investment strategies. A great deal of research has been devoted to developing various random walk tests since the late 1980s (see, for example, French and Roll, 1986; Fama and French, 1988; Lo and MacKinlay, 1988; Poterba and Summers, 1988; Jegadeesh, 1991, among others). While exploring the information contained in past stock prices in different ways, those tests are attained by assuming that stock prices are modeled by continuous-state stochastic processes. Since major stock exchanges require that quotes and transaction prices be stated as some multiple of a minimum price variation (or trading tick), individual stocks often display compass rose patterns, suggesting that continuous-state models are only approximations to actual market prices. We anticipate that such compass rose patterns can induce spurious serial autocorrelations in return. Consequently, the asymptotic theory of test statistics based on the continuous-state argument is in question for transaction prices.

To address the problem of testing the random walk model in the presence of the discrete jumps in prices, we focus our attention to variance-ratio tests; in particular to the one proposed by Lo and MacKinlay (1988). As one of the most widely used random walk tests, Lo and MacKinlay's variance-ratio test is heteroscedasticity consistent. In view of the growing consensus among financial economists that volatilities do change over time (Merton, 1980; Poterba and Summers, 1986), the testing results based on heteroscedasticity-consistent tests are of more practical interest in comparison with those from tests without accounting for heteroscedasticity. However, to provide a perspective on test statistics without heteroscedasticity adjustment, the standard Z test (a simple homoscedastic variance-ratio test) and Ljung and Box's (1978) Q statistic are also studied, and both the Gaussian independently and identically distributed (i.i.d.) null hypothesis and a simple heteroscedastic null are investigated. We also note that our results do not, generally speaking, apply only to Lo and MacKinlay's variance-ratio test. In fact, as will be shown, many random walk tests (including Lo and MacKinlay's variance-ratio test) can be expressed as linear combinations of consistent estimators of autocorrelations. Since those autocorrelation estimates are biased by price discreteness regardless of the time periods in calculating returns, the results reported in this paper should be of general interest.

We show that in general, the asymptotic theory of test statistics is invalid for transactions data. For transactions data with high levels of tick/volatility, the true rejection probability exceeds the nominal size by wide margins under both homoscedastic and heteroscedastic null hypotheses. For transactions data with relatively low levels of tick/volatility, discreteness is, however, considerably less problematic and the actual size of the test is well approximated by the continuous-state asymptotic theory. The results are consistent across different nominal significance levels.

The structure of the paper is as follows. Section 2 describes test statistics and the compass rose pattern. Section 3 examines the compass rose effect measured by the empirical size of the test under both homoscedastic and heteroscedastic null hypotheses. Section 4 contains a discussion.

2. Test statistics and the compass rose pattern

2.1. Test statistics

It is assumed throughout that the logarithm of the equilibrium stock price at the end of the ith period, p_i, satisfies the random walk model:

$$\log p_i = \mu + \log p_{i-1} + \varepsilon_i, \tag{1}$$

where μ is the drift term and the disturbance ε_i is serially uncorrelated with mean zero, but could be heteroscedastic.

The variance-ratio test is based on the fact that if the price series p_i follows a random walk formulated as in (1), then, the variance of the q-differences $[\log(p_i) - \log(p_{i-q})]$ grows linearly with the size of q (q is an integer > 1). Define the variance-ratio of q observations, $VR(q)$, as

$$VR(q) = \frac{\sigma_c^2(q)}{\sigma_a^2} \tag{2}$$

where

$$\sigma_c^2(q) = \frac{1}{q(nq - q + 1)(1 - 1/n)} \sum_{j=q}^{nq} (\log p_j - \log p_{j-q} - q\hat{\mu})^2$$

and

$$\sigma_a^2 = \frac{1}{nq - 1} \sum_{j=1}^{nq} (\log p_j - \log p_{j-1} - \hat{\mu})^2$$

with $\hat{\mu} = (1/nq)(\log p_{nq} - \log p_0)$. It is well known that the standardized $VR(q)$ can be used to test the random walk null hypothesis under homoscedasticity and heteroscedasticity. If the ε_is in (1) are i.i.d., the Z test statistic, given by

$$Z(q) = \frac{VR(q) - 1}{[2(2q - 1)(q - 1)/(3q(nq))]^{1/2}}, \tag{3}$$

follows the standard Gaussian distribution $N(0,1)$. If the null hypothesis assumes that p_i possesses uncorrelated increments, but allows for heteroscedasticity, the refined Z^* test statistic,

$$Z^*(q) = \frac{VR(q) - 1}{[\hat{\theta}/(nq)]^{1/2}}, \tag{4}$$

is asymptotically a standard Gaussian. The quantity $\hat{\theta}$ is the asymptotic variance of the variance-ratio that is heteroscedasticity consistent (Lo and MacKinlay, 1988). Note that the fundamental difference between Z and Z^* is that Z^* relies on the calculation of standard errors of lagged serial correlation coefficients using White's (1980) heteroscedasticity-consistent covariance matrix estimator, but Z does not.

It can be shown that $VR(q)$ satisfies the following asymptotic equality:

$$VR(q) - 1 \approx \sum_{j=1}^{q-1} \frac{2(q-j)}{q} \hat{\rho}(j), \tag{5}$$

where $\hat{\rho}(k)$ denotes the kth-order autocorrelation coefficient of continuously compounded returns. As a consequence, the sign of Z (Z^*) depends primarily on whether returns are positively or negatively autocorrelated, or equivalently, whether there is mean reversion in the return series (Summers, 1986; Fama and French, 1988).

Note the similarity between Z (Z^*) and the Ljung-Box (1978) Q statistic of order $q-1$,

$$Q(q-1) = T(T+2) \sum_{j=1}^{q-1} \frac{\hat{\rho}^2(j)}{T-j}, \tag{6}$$

where $T = nq + 1$. As a portmanteau statistic, the Q statistic is designed to detect departures from zero autocorrelations at all lags by summing the squared autocorrelations.

Although we report results for the Z and Q statistics under a heteroscedastic null, we emphasize that results are only for illustrative purposes since these two tests have been designed with the homoscedastic null hypothesis in mind. The inclusion of Z and Q tests in the study of a heteroscedasticity null does, however, reveal interesting results on the combined effects of price discreteness and heteroscedasticity.

2.2. The compass rose

Suppose that the observed stock price at time i, p_i^*, is subject to price discreteness effects. For example, consider the rounding model:

$$p_i^* = \left[\frac{p_i}{d} + \frac{1}{2} \right] d, \tag{7}$$

where the floor function $[x] \equiv$ greatest integer less than or equal to x. The parameter d is the tick size. In this section, we use an eighth of a dollar ($d = 0.125$) as the benchmark — the minimum price movement of stocks with prices greater than or equal to \$1 — the longstanding practice on the NYSE before June 24, 1997, when the NYSE began to trade stocks in sixteenths.

In the rounding model (7), p_i is rounded to the nearest multiple of d. The discreteness is specified as a grid on which outcomes must lie, but no distinctive properties are attributed to particular points on the grid. When the tick size is $\frac{1}{8}$, for instance, if the current price is 60 $\frac{1}{8}$, the price change is equally likely to be 60 or 60 $\frac{1}{4}$.

To develop some intuition for properties of rounded prices in relation to the tick/volatility ratio, we generate a time series of 1000 observations based on (1) with the i.i.d. zero mean[1] Gaussion disturbance and round it according to the method described above. Two levels of standard deviation σ_ε of ε_i are used for the purpose

[1] Price series with non-zero means are considered in Section 3.

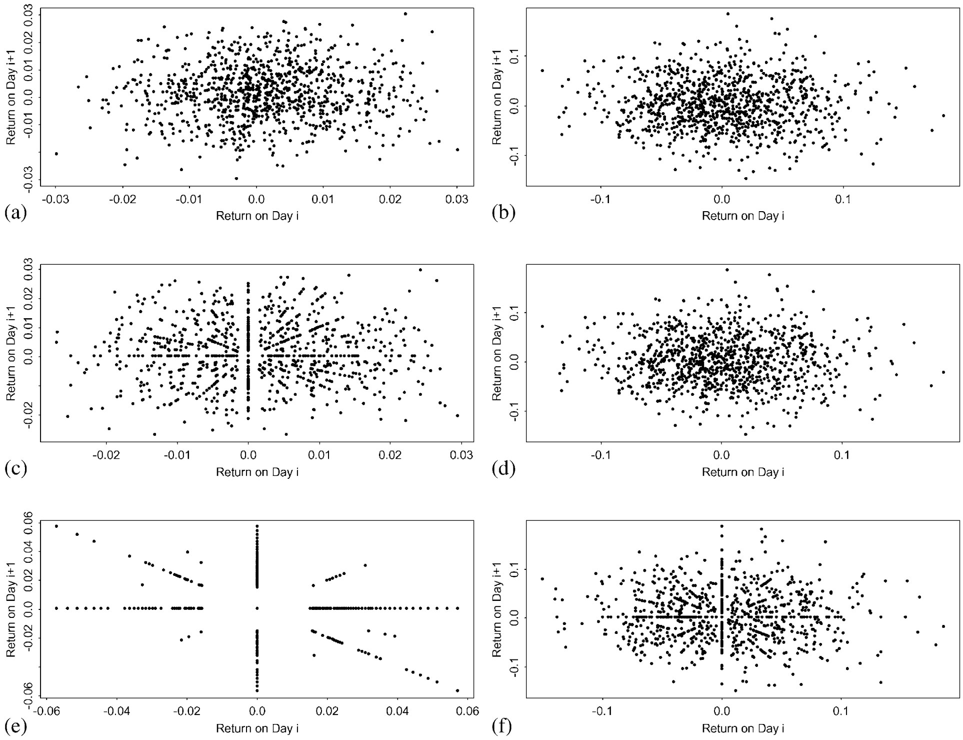

Fig. 1. Compass rose patterns. (a), (c) and (e) are based on $\sigma_\varepsilon = 0.01$; while (b), (d) and (f) are based on $\sigma_\varepsilon = 0.05$. (a) and (b) are obtained from original prices with machine accuracy (16 decimal digits); (c) and (d) are from the prices rounded by an eighth of a dollar; and (e) and (f) are from prices rounded to the nearest dollar.

of evaluating the impact of the tick/volatility ratio on test statistics. To be representative, σ_ε is taken to be 0.01 and 0.05, which represent standard deviations for daily returns of low- and high-volatilities, respectively. This range of daily stock volatility is based on the summary statistics for daily returns of Center for Research in Securities Prices (CRSP) equal- and value-weighted stock indexes and various individual stocks in different size deciles, which are continuously listed over the time period from July 3, 1962 to December 30, 1994 (Campbell et al., 1997). The generated price series is either rounded by an eighth of a dollar or to the nearest dollar.

Fig. 1 presents six scatter plots for daily returns rounded with different levels of d, as returns $(p_i^*/p_{i-1}^* - 1)$ at time i are plotted against lagged returns $(p_{i-1}^*/p_{i-2}^* - 1)$ at time $i-1$. As can be seen, there is a considerable structure in returns if prices are rounded, especially to the nearest dollar. The radically symmetric structure (the so-called *compass rose*) is solely attributable to price discreteness. For low-volatility returns (Fig. 1a, c and e), the structure is evident when the price is rounded by an eighth of a dollar and is much more striking when the price is rounded to the nearest dollar. In contrast, no structure is apparent for high-volatility returns as the price is rounded by an eighth of a dollar (Fig. 1d). Low-volatility returns display less structure than that for low-volatility returns when the price is rounded to the nearest dollar (Fig. 1f).

The results emerged from Fig. 1 suggest that the tick/volatility ratio is an important determinant of the compass rose pattern. It is apparent that the pattern depends intimately on the tick size. It can also be seen from Fig. 1 that the pattern counts on volatility likewise. In fact, the tick/volatility ratio is a more appropriate measure than the tick size in determining whether the pattern is observed. For example, Lee et al. (1999) have shown that the pattern does not exist for daily foreign exchange (FOREX) data, while Szpiro (1998) found the pattern in transactions FOREX data. Noting that the tick/volatility ratio is very small for daily FOREX data but substantially larger for transactions data, [2] the analysis based on the tick/volatility ratio sheds light on the seemingly conflicting results (Lee et al., 1999; Szpiro, 1998) which is difficult to explain by only the tick size.

The notion that the tick/volatility ratio accounts for the presence of the compass rose patterns has been implicitly shared by many researchers by viewing the fact that most empirical studies have been based on daily or weekly data, while the use of intraday data on testing random walk model is rarely found in the literature. For example, weekly stock prices were used in testing the random walk hypothesis for the US stock market in Lo and MacKinlay (1988). As the authors argued, weekly data is the ideal compromise, minimizing the biases inherent in the daily data while yielding a large number of observations which is important since the sampling theory of the test statistic is based on asymptotic approximation. We will provide a formal justification for this argument and quantify the compass rose effect by calculating the empirical sizes of test statistics in the next section.

3. Properties of test statistics

In the setup given in Section 2, the finite-sample properties of test statistics are largely unknown. The asymptotic theory based on p_i^* cannot be derived due to intractable nonlinearity. To evaluate the compass rose effect on test statistics, we calculate empirical sizes of Z^*, Z and Q statistics via simulation experiments under both the Gaussian i.i.d. null hypothesis and a simple heteroscedastic null hypothesis. All simulations are based on 5000 replications and performed in single-precision FORTRAN using random number generators of the IMSL subroutine library. The nominal significance level is chosen to be 1%, 5% and 10%, while the sample size is taken to be 1000 with the standard deviation of ε_i chosen to be 0.01 and 0.05. Test statistics are computed in four different ways: from original prices with machine accuracy (16 decimal digits), price series rounded to the nearest cent ($d = 0.01$), by an eighth of a dollar ($d = 0.125$), or to the nearest dollar ($d = 1.0$). To check the

[2] The FOREX data are quoted in multiples (ranging from 0 to 20) of one point, with one point being equal to, for instance, 0.0001 in the Deutsche mark and US dollar, and 0.01 in the Japanese yen and US dollar.

sensitivity of the test statistic to the truncation lag length, all statistics are computed for several different values of q.

To investigate the effect of the parameter μ, we consider different levels of μ based on the empirical estimates of μ for the Standard & Poor's stocks (French et al., 1987). Since a compass rose should only be apparent if the effective tick size is sufficiently large as compared to volatility regardless of the levels of μ (Crack and Ledoit, 1996), the sizes of the tests are not noticeably affected by μ, especially when the price series are rounded to the nearest cent or by an eighth of a dollar. Hence, for brevity, we will report results only for $\mu = 2.5 \times 10^{-4}$. For illustrative purposes, however, we will also provide results on $\mu = 5 \times 10^{-4}$ in Panel A2 of Table 1.

3.1. The Gaussian i.i.d. null hypothesis

Table 1 reports results of the simulation experiment conducted under the Gaussian i.i.d. null hypothesis. Results for low- and high-volatility returns are organized in Panels A1 and A2, and B, respectively. Overall, the tail behaviors of the three test statistics are comparable. The results show that if original prices are rounded to the nearest cent or by an eighth of a dollar, the rejection rates are inflated, but are close to their nominal values for both low- and high-volatility returns. In particular, when the volatility is high, the tick/volatility ratio is relatively low and the compass rose bias is essentially not observed. On the other hand, when prices are rounded to the nearest dollar, the true rejection probabilities exceed the nominal size by wide margins. For example, about two thirds of the entries for prices rounded to the nearest dollar in Panels A1 and A2 are above 95%. This implies that for low-volatility returns, no matter what significance levels are used, the prices rounded to the nearest dollar almost always yield a rejection of the correct null hypothesis of i.i.d. returns. Although the compass rose bias is considerably smaller for the high-volatility case, the empirical sizes still range from about 20% to 40%, depending on the nominal significance levels (Panel B). It is also interesting to see that the compass rose effect does not appear to depend on q when prices are rounded by an eighth of a dollar. However, when prices are rounded to the nearest dollar, the bias on the size of the test decreases with q for low-volatility returns while the bias increases with q for high-volatility returns.

Finally, as mentioned earlier, the levels of μ do not have a noticeable effect on the size of the tests, as can be seen by contrasting the values in Panels A1 and A2.

To gain more insight on these results, we plot histograms of Z based on the simulated series. For illustrative purposes, six histograms of $Z(q)$ with $q = 4$ are presented in Fig. 2. Consistent to the findings in Table 1, no clear shift in the shape of the empirical distribution of Z is observed when prices are rounded to the nearest cent (which is very similar to those based on original prices and is not reported here to conserve space) or by an eighth of a dollar. [3] However, if prices are rounded to the nearest dollar, it becomes apparent that the mean values of Z in Fig. 2e and f move in a negative direction and are about -4.25 and -0.80 for low- and high-volatility

[3] It may be also verified that the normality cannot be rejected at any commonly used significance levels.

Table 1
Empirical rejection probabilities of the Z^*, Z and Q tests of the random walk null hypothesis with homoscedastic disturbances (quoted in percent)

	Nominal size	Z^*			Z			Q		
		10%	5%	1%	10%	5%	1%	10%	5%	1%
Panel A1: $\sigma_\varepsilon = 0.01$ *and* $\mu = 2.5 \times 10^{-4}$										
Original prices	$q = 4$	10.52	5.50	1.11	10.64	4.82	1.01	10.16	4.56	1.02
	16	10.94	5.66	1.08	10.48	4.78	0.97	10.88	4.82	0.94
	64	11.41	6.08	1.44	9.36	4.84	0.85	9.74	4.68	0.90
Prices rounded to	$q = 4$	10.38	5.48	1.06	10.38	4.99	0.98	10.14	4.56	0.98
the nearest cent	16	10.96	5.58	1.22	10.82	4.98	1.08	10.78	4.72	0.96
	64	11.82	5.86	1.29	9.64	5.04	1.30	9.81	4.89	0.92
Prices rounded	$q = 4$	11.80	5.48	1.45	10.99	5.74	1.51	11.08	5.84	1.67
by an eighth	16	11.22	6.04	1.87	11.01	5.27	1.44	10.86	5.64	1.51
of a dollar	64	11.89	6.66	1.64	10.50	5.46	1.11	11.02	5.52	1.18
Prices rounded to	$q = 4$	99.96	99.61	99.10	99.78	99.71	99.02	99.10	99.08	98.70
the nearest dollar	16	97.50	97.00	95.21	99.66	99.10	95.58	99.87	99.85	97.70
	64	95.24	90.12	78.12	90.12	81.12	71.24	94.78	91.54	84.78
Panel A2: $\sigma_\varepsilon = 0.01$ *and* $\mu = 5 \times 10^{-4}$										
Original prices	$q = 4$	10.52	5.56	1.08	10.58	4.48	0.99	10.15	4.56	1.05
	16	10.62	5.66	1.08	10.48	4.70	0.97	10.78	4.82	0.97
	64	10.82	5.74	1.25	9.36	4.84	0.87	9.48	4.70	0.91
Prices rounded to	$q = 4$	10.84	5.47	1.05	10.72	4.78	0.98	9.98	4.68	1.00
the nearest cent	16	11.18	5.44	1.22	11.05	5.18	1.02	10.88	4.88	0.98
	64	11.83	5.72	1.25	10.01	5.09	1.44	9.87	4.95	0.91
Prices rounded	$q = 4$	11.46	5.87	1.44	10.98	5.44	1.41	10.10	5.68	1.66
by an eighth	16	11.38	5.68	1.61	10.87	5.25	1.14	10.56	5.80	1.18
of a dollar	64	12.18	6.32	1.64	10.50	5.28	1.08	10.88	5.50	1.16
Prices rounded to	$q = 4$	99.45	99.12	98.55	99.56	99.12	98.77	99.05	98.47	97.23
the nearest dollar	16	97.12	95.40	94.41	98.89	97.10	94.51	99.54	99.04	95.55
	64	94.55	87.55	73.77	87.12	80.01	69.88	94.14	90.14	81.45
Panel B: $\sigma_\varepsilon = 0.05$ *and* $\mu = 2.5 \times 10^{-4}$										
Original prices	$q = 4$	10.12	5.01	1.01	10.04	4.57	1.02	10.04	4.98	1.05
	16	10.52	5.21	1.47	10.44	4.98	0.94	9.98	4.77	0.99
	64	10.64	5.88	1.48	10.24	4.87	0.97	10.51	5.14	0.94
Prices rounded to	$q = 4$	10.14	4.89	1.12	10.21	4.91	1.05	10.12	4.98	1.14
the nearest cent	16	10.28	5.14	1.32	10.03	5.02	1.08	10.04	4.85	1.01
	64	10.54	5.86	1.27	10.51	4.98	1.14	10.27	5.25	1.08
Prices rounded	$q = 4$	10.78	5.12	1.14	10.66	4.99	1.11	10.78	5.12	1.24
by an eighth	16	11.02	5.87	1.51	11.02	5.24	1.27	10.54	5.08	1.12
of a dollar	64	10.96	6.01	1.47	10.89	5.78	1.54	10.52	5.58	1.07
Prices rounded to	$q = 4$	36.55	28.99	21.78	35.99	30.14	20.58	36.41	30.04	20.55
the nearest dollar	16	36.21	32.45	20.43	36.78	32.02	23.45	36.78	31.24	24.14
	64	37.41	32.14	21.12	40.12	29.01	22.14	35.87	30.78	25.06

returns, respectively. The negative mean of the distribution of Z suggests that some of the lower-order autocorrelations of returns of prices p_i^* turn out to be negative, which have had serious size implications, as demonstrated by the results in Table 1.

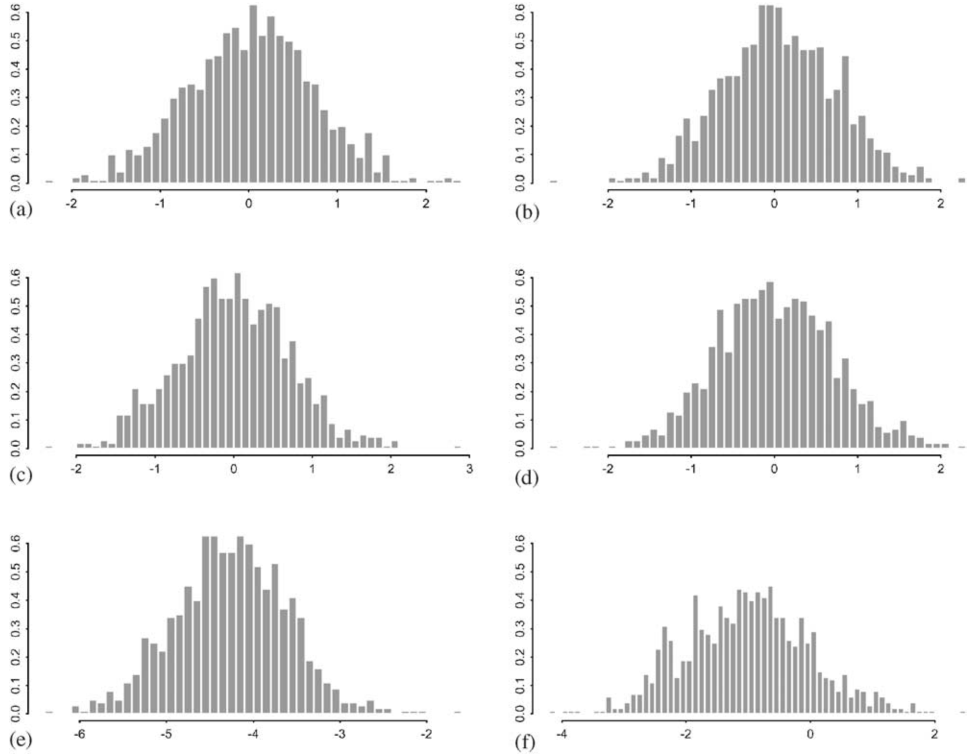

Fig. 2. Histograms of $Z(4)$ under the i.i.d. $N(0,\sigma_\varepsilon^2)$ null. (a), (c) and (e) are based on $\sigma_\varepsilon = 0.01$; while (b), (d) and (f) are based on $\sigma_\varepsilon = 0.05$. (a) and (b) are obtained from original prices with machine accuracy (16 decimal digits); (c) and (d) are from the prices rounded by an eighth of a dollar; and (e) and (f) are from prices rounded to the nearest dollar.

3.2. A heteroscedastic null hypothesis

To evaluate the compass rose effect on test statistics in the case of heteroscedastic returns, we perform simulation experiments under the null hypothesis that the disturbance ε_i in (1) is serially uncorrelated but heteroscedastic in the following manner. Let the random walk disturbance ε_i satisfy the relation $\varepsilon_i \equiv \sigma_i \xi_i$, where ξ_i is i.i.d. $N(0,1)$ and σ_i satisfies the following first-order general autoregressive conditional heteroscedasticity (GARCH(1,1)) (Bollerslev, 1986) process:

$$\sigma_i^2 = \gamma + \alpha\sigma_{i-1}^2 + \beta r_{i-1}^2, \tag{8}$$

where parameters α, β and γ are chosen to yield the desired levels of σ_ε, and r_i is the return based on p_i.[4] The results are displayed in Table 2.

Not surprisingly, the Z^* behaves very similarly as under the i.i.d. null because Z^* is robust to fairly general forms of conditional heteroscedasticity including that in (8). However, results for both the Z and Ljung-Box Q statistics are evidently different

[4] The parameters $\alpha = 0.8$, and $\beta = 0.1$ and $\gamma = \sigma_\varepsilon^2(1-\alpha-\beta)$.

Table 2
Empirical rejection probabilities of the Z^*, Z and Q tests of the random walk null hypothesis with heteroscedastic disturbances (quoted in percent)

	Nominal size	Z^*			Z			Q		
		10%	5%	1%	10%	5%	1%	10%	5%	1%
Panel A: $\sigma_\varepsilon = 0.01$ *and* $\mu = 2.5 \times 10^{-4}$										
Original prices	$q = 4$	10.03	4.68	0.99	34.22	25.58	13.99	34.28	25.61	14.21
	16	10.88	5.31	1.08	32.05	23.54	12.04	40.11	30.34	17.54
	64	11.12	5.61	1.14	27.05	18.54	9.01	36.25	28.54	14.56
Prices rounded to	$q = 4$	10.12	4.88	1.06	34.88	25.87	14.11	34.29	25.64	14.35
the nearest cent	16	10.57	5.36	1.24	32.14	23.55	12.14	40.24	30.35	17.66
	64	11.13	5.57	1.07	34.67	18.65	8.87	34.56	26.01	15.02
Prices rounded	$q = 4$	10.54	5.03	1.14	35.54	26.14	14.57	34.47	26.17	14.78
by an eighth	16	11.41	5.58	1.68	33.14	23.12	13.04	40.05	30.41	17.66
of a dollar	64	11.87	5.98	1.87	27.08	19.04	9.14	36.88	28.74	15.65
Prices rounded to	$q = 4$	80.57	74.15	57.96	91.55	88.13	80.05	91.54	87.65	79.86
the nearest dollar	16	70.05	60.89	43.47	81.87	76.78	65.13	84.21	80.15	70.04
	64	45.63	36.02	19.78	59.68	52.45	36.54	72.14	64.37	50.24
Panel B: $\sigma_\varepsilon = 0.05$ *and* $\mu = 2.5 \times 10^{-4}$										
Original prices	$q = 4$	10.06	4.97	0.91	33.56	25.87	13.56	33.64	25.64	14.52
	16	10.88	5.21	1.04	32.01	23.14	11.64	39.16	30.51	16.31
	64	10.67	5.68	1.25	27.05	19.31	8.97	36.21	27.35	15.12
Prices rounded to	$q = 4$	10.05	5.04	1.05	33.54	25.67	13.64	34.17	26.47	15.03
the nearest cent	16	11.02	5.63	1.14	32.14	23.34	11.85	39.64	31.58	17.45
	64	10.99	5.24	1.36	26.97	19.63	9.05	36.58	28.54	16.33
Prices rounded	$q = 4$	10.77	5.34	1.15	34.86	26.87	15.14	35.04	26.98	15.64
by an eighth	16	11.64	5.89	1.31	35.47	27.63	16.76	43.21	33.42	20.68
of a dollar	64	11.05	5.87	1.54	34.55	26.04	16.84	43.11	34.57	22.84
Prices rounded to	$q = 4$	31.57	24.04	17.54	51.37	43.14	33.14	50.86	42.54	33.14
the nearest dollar	16	37.87	33.50	24.14	53.08	47.96	36.58	59.14	52.98	41.05
	64	41.25	35.17	27.86	52.45	45.63	36.68	59.84	52.84	43.51

due to the inconsistent heteroscedastic nature of these two tests. For example, if low-volatility prices are rounded with the minimum price change of an eighth of a dollar, the sizes of $Z(4)$ and $Q(4)$ for the nominal level 1% could be as large as about 14%. In view of the results in Table 1, it becomes apparent that these high rejection ratios are attributed almost solely to the effects of price heteroscedasticity. When prices are rounded to the nearest dollar, the higher rejection rates of the Z and Q statistics are due to the combination effects from both heteroscedasticity and price discreteness. Note that when prices are rounded to the nearest dollar, sizes of all three test statistics for low-volatility returns decline compared to those in the i.i.d. case (Table 1).

We remark that although we restrict our discussion to daily data, the results apply to data with higher or lower sampling frequencies with appropriately transformed volatility levels. The volatility or the tick/volatility ratio for the return series with a different sampling frequency can be easily calculated using the well-known temporal aggregation results. For example, weekly returns with a standard deviation of 0.05 can be viewed also as a daily return series with a standard deviation of 0.022. [5]

In summary, our results show that the compass rose effect is observed only if the tick/volatility ratio exceeds some threshold level. For the case that $\sigma_\varepsilon = 0.01$, the threshold is relatively small. Hence, we see that the empirical sizes of test statistics are inflated even when prices are rounded to the nearest cent. With a greater value for σ_ε, such as 0.05, the threshold is high and the compass rose effect is very limited if prices are rounded to the nearest cent. Since intraday returns fall into the former category of a smaller volatility while daily or weekly returns fall into the later, the size of the test is well approximated by the continuous-state asymptotic theory if the test is based on daily or weekly returns but not if based on intraday data.

4. Discussion

We have demonstrated in Section 2 that both Z and Z^* can be rewritten as weighted averages of the autocorrelations. We anticipate that our results hold, at least qualitatively, for a general class of random walk tests which are linear combinations of consistent estimators of autocorrelations. Richardson and Smith (1994) have shown that this class captures many test statistics studied in the recent finance and macroeconomics literature. For example, the test considered in Fama and French (1988) is based on the statistic:

$$\tilde{Z}(q) = \frac{(1/n)\sum_{i=1}^{n}\left[\sum_{l=1}^{q}(\log p_{i+l} - \log p_{i+l-1} - \hat{\mu})\sum_{l=1}^{q}(\log p_{i+l-q} - \log p_{i+l-q-1} - \hat{\mu})\right]}{(1/n)\sum_{i=1}^{n}\left(\sum_{l=1}^{q}(\log p_{i+l} - \log p_{i+l-1} - \hat{\mu})\right)^2}, \tag{9}$$

which can be rewritten asymtotically in terms of consistent autocorrelation estimators as

$$\tilde{Z}(q) \approx \sum_{i=1}^{2q-1} \min(i, 2q-i)\hat{\rho}(i)/q. \tag{10}$$

The limiting distribution of (9) can be derived from (10) and the standard result on the asymptotic distribution of autocorrelation estimators (for example, Theorem 7.2.1 of Brockwell and Davis, 1991). Although the test statistic based on $\tilde{Z}(q)$ is different from Lo and Mackinlay's test (albeit with different weights on autocorrelations of various lags), the standardized $\tilde{Z}(q)$ will not be asymptotically a standard Gaussian if the rounded price p_i^* is used in the calculation of the test statistic. This observation is supported by empirical evidence from a Monte Carlo simulation. Since the simulation results of $\tilde{Z}(q)$ are similar to those for Lo and MacKinlay's test in terms of the robustness to the compass rose pattern, they are not reported here.

[5] The standard deviation of daily returns is $0.05/\sqrt{5} = 0.022$, assuming that there are 5 trading days per week.

Acknowledgements

The author thanks the referees for their valuable comments on earlier versions of this article.

References

Bollerslev, T., 1986. Generalized autoregressive conditional heteroskedasticity. J. Econometrics 31, 307–327.

Brockwell, P., Davis, R., 1991. Time Series: Theory and Methods. Springer, New York.

Campbell, J., Lo, A., MacKinlay, A., 1997. The Econometrics of Financial Markets. Princeton University Press, Princeton.

Crack, T., Ledoit, O., 1996. Robust structure without predictability: the 'compass rose' pattern of the stock market. Journal of Finance 51, 751–762.

Fama, E., French, K., 1988. Permanent and temporary components of stock prices. Journal of Political Economy 96, 246–273.

French, K., Roll, R., 1986. Stock return variances: the arrival of information and the reaction of traders. Journal of Financial Economics 17, 5–26.

French, K., Schwert, W., Stambaugh, R., 1987. Expected stock returns and volatility. Journal of Financial Economics 19, 3–29.

Jegadeesh, N., 1991. Seasonality in stock price mean reversion: evidence from the U.S. and U.K. Journal of Finance 46, 1427–1444.

Kramer, W., Runde, R., 1997. Chaos and the compass rose. Economics Letters 54, 113–118.

Lee, C., Gleason, K., Mathur, I., 1999. A comprehensive examination of the compass rose pattern in futures markets. Journal of Futures Markets 19, 541–564.

Ljung, G., Box, G., 1978. On a measure of lack of fit in time series models. Biometrika 66, 67–72.

Lo, A., MacKinlay, C., 1988. Stock market prices do not follow random walks: evidence from a simple specification test. The Review of Financial Studies 1, 41–66.

Merton, R., 1980. On estimating the expected return on the market: an exploratory investigation. Journal of Financial Economics 8, 323–361.

Poterba, J., Summers, L., 1986. The persistence of volatility and stock market fluctuations. American Economic Review 76, 1142–1151.

Poterba, J., Summers, L., 1988. Mean reversion in stock prices: evidence and implications. Journal of Financial Economics 22, 27–59.

Richardson, M., Smith, T., 1994. A unified approach to testing for serial correlation in stock returns. Journal of Business 67, 371–399.

Summers, L., 1986. Does the stock market rationally reflect fundamental values? Journal of Finance 41, 591–600.

Szpiro, G., 1998. Tick size, the compass rose and market nanostructure. Journal of Banking and Finance 22, 1559–1569.

White, H., 1980. A heteroscedasticity-consistent covariance matrix estimator and a direct test for heteroscedasticity. Econometrica 48, 817–838.