Discrete Dynamics in Nature and Society, Vol. 6, pp. 101–120
Reprints available directly from the publisher
Photocopying permitted by license only

The Angular Distribution of Asset Returns in Delay Space*

ROGER KOPPL[a,†] and CARLO NARDONE[b,‡]

[a]*Department of Economics and Finance, Fairleigh Dickinson University, Madison, NJ 07940, USA;*
[b]*Sun Microsystems Italia SpA, via G. Romagnosi 4, I-00196 Roma, Italy*

(Received 14 May 2000)

Plotting asset returns against themselves with a one-period lag reveals the "compass rose" pattern of Crack and Ledoit (1996). They describe the pattern, caused by discreteness, as "subjective". We develop a new and original set of "objective" statistical procedures to quantify the compass rose and detect changes in it. Comparing empirical and bootstrapped "theta histograms" permits hypothesis testing. Simulations suggest that intertemporal statistical dependence skews the compass rose in ways that mimic ARCH phenomena. Using our techniques on "credit ruble" data, we test the hypothesis that "Big Players" influence the degree of this "X-skewing" and, therefore, apparent ARCH behavior.

Keywords: Compass rose; Theta histograms; X-skewing; ARCH phenomena

Crack and Ledoit (1996) plot daily stock returns against themselves with one day's lag. Doing so produces the "compass rose" pattern of Figures 1 and 2. The points of the graph are concentrated along several evenly spaced rays from the origin. The rays corresponding to the major directions of the compass accumulate the most points. This "strikingly geometrical" pattern "is indisputably present in every stock." The existence of a non-zero tick size produces discreteness in the data which, in turn, generates the compass rose.

The compass rose biases some standard statistical tools of financial analysis. The Monte Carlo study of Kramer and Runde (1997) proves that the BDS test will falsely indicate chaotic structure when applied to discrete data. The larger the (simulated) tick size, the greater this false propensity of the BDS test. Crack and Ledoit conjecture that discreteness biases the standard tests for

*We thank Geoffery Booth, Rainford Knight, Richard Langlois, Maria Minniti, Sangeeta Pratap and Soren Tuluca for useful comments. Special thanks to John Broussard for help with SAS progamming, for comments, and for encouragement. Only we are to blame for any errors in the paper.
†Corresponding author. e-mail: koppl@alpha.fdu.edu
‡e-mail: carlo.nardone@italy.sun.com

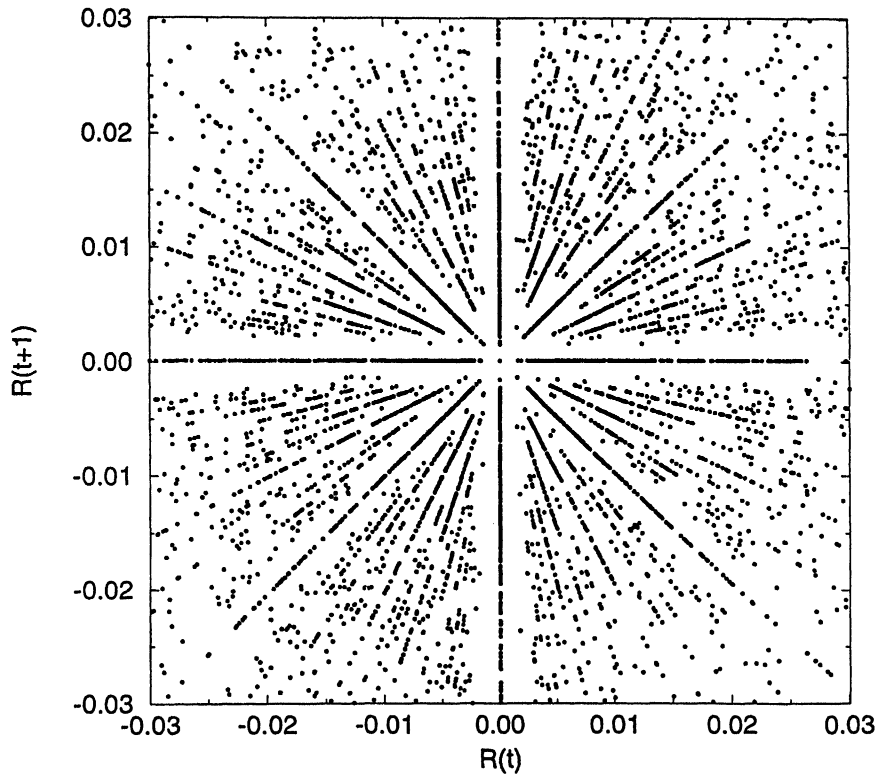

FIGURE 1 The horizontal axis shows a given day's return for Weyerhauser stock on the New York Stock Exchange. The vertical axis shows the next day's return. The data cover the period from December 6, 1963 to December 31, 1993. Out of 7558 points, 1212 fall outside the graph.

autoregressive conditional heteroskedasticity (ARCH). Apparently, financial researchers cannot rely on the statistical tools most commonly used to study volatility dynamics.

Surprisingly, Crack and Ledoit do not call for new tools of time-series analysis specifically suited to the existence of ticks and of the compass rose pattern. They seem to have assumed objective tools of analysis cannot be created for use on the compass rose. Their explanation of the pattern uses "subjective language", they report, "because the above statement 'the compass rose appears clearly' is itself subjective" (p. 754). We argue, however, that if our current tools are biased, we should try to create new "objective" tools that are not biased by discreteness.

Chen (1997) takes a step in that direction by using the information in the compass rose to construct improved forecasts within an ARMA-GARCH framework. He reports improved forecasts in all cases. If Chen's result is sustained by future studies, it will contradict Crack and Ledoit's conjecture that the compass rose "cannot be used for predictive purposes", because "it is an artifact of market microstructure" (p. 751). Chen's forecasting algorithm corrects for discreteness, at least partly. His GARCH coefficient estimates, however, do not. Chen's ARMA-GARCH models, therefore, are subject to the same discreteness-induced biases likely to affect other ARCH techniques.

We show that new objective techniques can be devised which are not biased by discreteness and are suitable for the compass rose. We plot the number of points along a given ray of the compass rose against the angle of that ray. This creates a

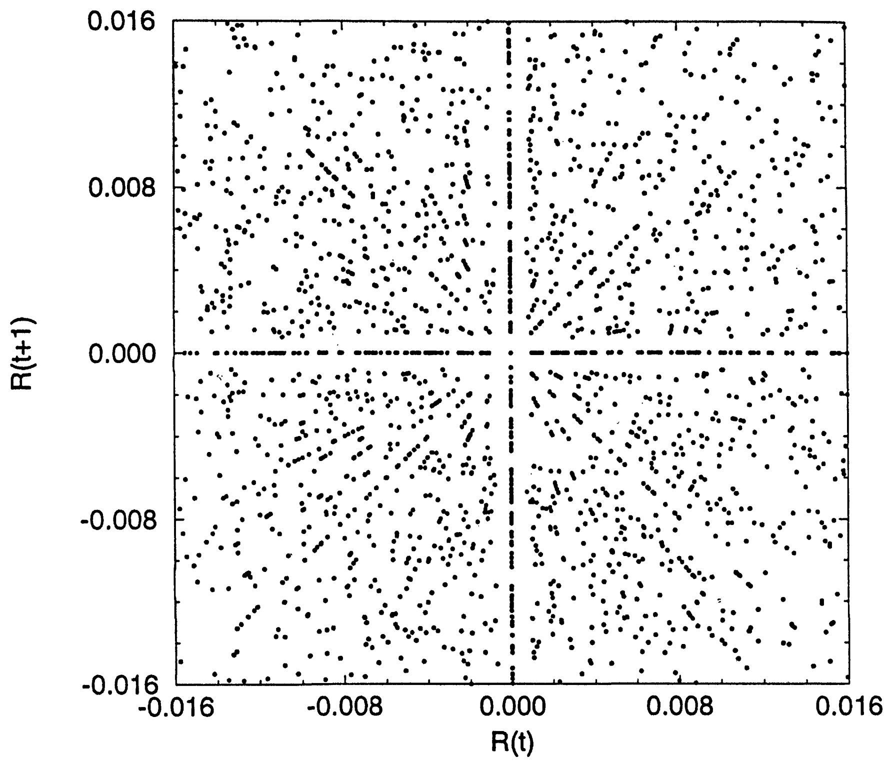

FIGURE 2 Crack and Ledoit's "compass rose" graph for IBM daily stock returns, January 1, 1980 to October 8, 1992.

"theta histogram" which describes the angular distribution of the points in delay space. We compare this distribution to a standard theta histogram created by a simple bootstrap procedure. The χ^2 test is then performed in order to estimate quantitatively the consistency of the actual data with the standard theta histogram. We test other hyphotheses using the Bernoulli distribution.

Our tests suggest that returns on stocks and other assets have a greater tendency to accumulate along the main diagonals of the compass rose than they would if they were statistically independent. An X pattern is embedded within the compass rose. In this sense, there are typically "too many" points along the main diagonals of the compass rose. This X pattern skews the compass rose away from the pattern that would exist if returns were statistically independent. We call this "X-skewing." A simple simulation shows that X-skewing of the compass rose can cause standard tests to indicate autocorrelation in conditional heteroskedasticity even when the underlying stochastic process is not consistent with existing ARCH models. X-skewing of the compass rose is a product of human actions that are subject to economic influences. The policy regime governing an asset market helps determine how much the compass rose is skewed. An important episode in Russian monetary history illustrates the point.

In the late nineteenth century, the "credit ruble" was a floating currency unlinked to precious metals. Generally, the finance ministry actively intervened to influence the ruble exchange rate. The one exception was during Nicolai Bunge's tenure as finance minister. Bunge's successor, Ivan Vyshnegradsky, was an unusually vigorous interventionist. The shift in regime from Bunge the non-interventionist to Vyshnegradsky the interventionist produced a marked change in the behavior of the ruble exchange rate. The angular distribution in delay space of the ruble's return against the German mark shifted dramatically under Vyshnegradsky. Hypothesis tests using our new techniques support the view that Vyshnegradsky's activism caused X-skewing. It caused a disproportionate number of points of the compass rose to accumulate along the main diagonals in delay space. The theory of "Big Players" (Koppl and Yeager, 1996) provides a plausible explanation.

Section I explains the compass rose and its history. Section II explains and illustrates our new statistical techniques. Section III uses a simple simulation to show how the X-skewing we discover causes the appearance of ARCH behavior. In Section IV we use data on the credit ruble to study the possibility that X-skewing, and thus apparent ARCH phenomena, are influenced by the policy regime. In particular, we argue that an increase in Big Player influence, or any other policy change that increases the difficulty of trading on fundamentals, will increase X-skewing and, therefore, the appearance of ARCH effects. Section V is the conclusion.

I. THE COMPASS ROSE

Crack and Ledoit chronicle the (brief) history of the compass rose in finance. Huang and Stoll explain in a footnote that a graph of intra-day returns plotted against themselves with a five-minute lag shows "clusters of points that radiate from the zero" (Huang and Stoll, 1994, p. 199, as quoted in Crack and Ledoit, 1996, p. 753). They do not pursue the issue, however.

Brealy and Meyers (1991) plot daily Weyerhauser stock returns against themselves with one day's lag (13-2, p. 293). They said "it is obvious from a glance that there is very little pattern in these price movements" (p. 291). Crack and Ledoit showed, however, that Brealy and Meyers did "overlook a *significant* pattern", namely, the compass rose (p. 752; emphasis in original). Brealy and Meyers would have seen the compass rose if they had used more data and a better graphical package. In the next edition of their textbook, Brealy and Meyers (1996) replace their earlier graph with Crack and Ledoit's Figure 2 (our Fig. 1) and include a footnote explaining the compass rose. They still say, however, "it is obvious from a glance that there is very little pattern in these price movements" (p. 326). The sentence continues "but we can test this more precisely by calculating the coefficient of correlation between each day's price change and the next" (p. 326). Brealy and Meyers thus neglect the difference between correlation and statistical dependence. As we show below, they overlook a significant pattern, as do Crack and Ledoit. The pattern we find is X-skewing. Like the ARCH behavior to which it is related, X-skewing is a case of statistical dependence without linear correlation.

Crack and Ledoit list three conditions for the compass rose pattern to emerge:

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

The explanation of these three conditions is straightforward. Following their notation, let P_t and R_t be the price and return of some stock on day t. If price changes are small relative to price level $((P_t - P_{t-1}) \ll P_t)$, and ignoring dividends and splits, the following approximation holds:

$$R_{t+1}/R_t = \frac{(P_{t+1} - P_t)/P_t}{(P_t - P_{t-1})/P_{t-1}} \approx \frac{P_{t+1} - P_t}{P_t - P_{t-1}} = \frac{n_{t+1}h}{n_t h} = \frac{n_{t+1}}{n_t}, \qquad (1)$$

where h is the tick size and $n_t = (P_t - P_{t-1})/h$ is the day-t price change calculated in ticks. Equation (1) shows that the ordered pairs (R_t, R_{t+1}) will be close to the rays through the origin that pass through (n_t, n_{t+1}). If prices usually change by a small number of ticks, then most points will accumulate along the major directions of the compass rose. The ticks that induce discreteness need not be official. As Crack and Ledoit explain, an official tick size is neither necessary nor sufficient for the compass rose. "The correct criterion for the existence of the compass rose is whether the effective tick size is of the same order of magnitude as typical price changes" (p. 758).

Finally, Crack and Ledoit explain, if the price varied only slightly around the value P_t, a grid pattern would result, not the compass rose. "On any given ray (m, n) data points would cluster at discrete distances from the origin: $(mh/P_t, nh/P_t)$, $(2mh/P_t, 2nh/P_t)$, and so on" (p. 754). Price variations produce "centrifugal smudging" which, in turn, produces the compass rose pattern.

II. HYPOTHESIS TESTING WITH THETA HISTOGRAMS

Crack and Ledoit describe the compass rose pattern as "subjective". It is possible, however, to transform the data of the compass rose and apply objective techniques to them. The transformation we propose is the result of a two step procedure. First, express the points of the compass rose in polar coordinates. The point (R_t, R_{t+1}) becomes (r_t, θ_t) where

$$r_t = \sqrt{R_t^2 + R_{t+1}^2}$$

$$\theta_t = \begin{cases} \arctan(R_{t+1}/R_t) & \text{if } R_t \geq 0 \\ \arctan(R_{t+1}/R_t) + \pi & \text{if } R_t < 0,\ R_{t+1} \geq 0 \\ \arctan(R_{t+1}/R_t) - \pi & \text{if } R_t, R_{t+1} < 0 \end{cases} \tag{2}$$

(arctan conventionally ranges from $-\pi/2$ to $\pi/2$).

Second, associate each θ_t, not with any of the corresponding r_t values, but with the number of such values corresponding to a narrow interval $\theta \pm \delta\theta$. Finally, normalize by π in order to plot histograms in the interval $[-1, 1]$. We call the result a "theta histogram". A theta histogram represents the angular distribution of asset returns in delay space.

Figure 3 illustrates. The horizontal axis shows the value of θ/π. The θ/π values have been

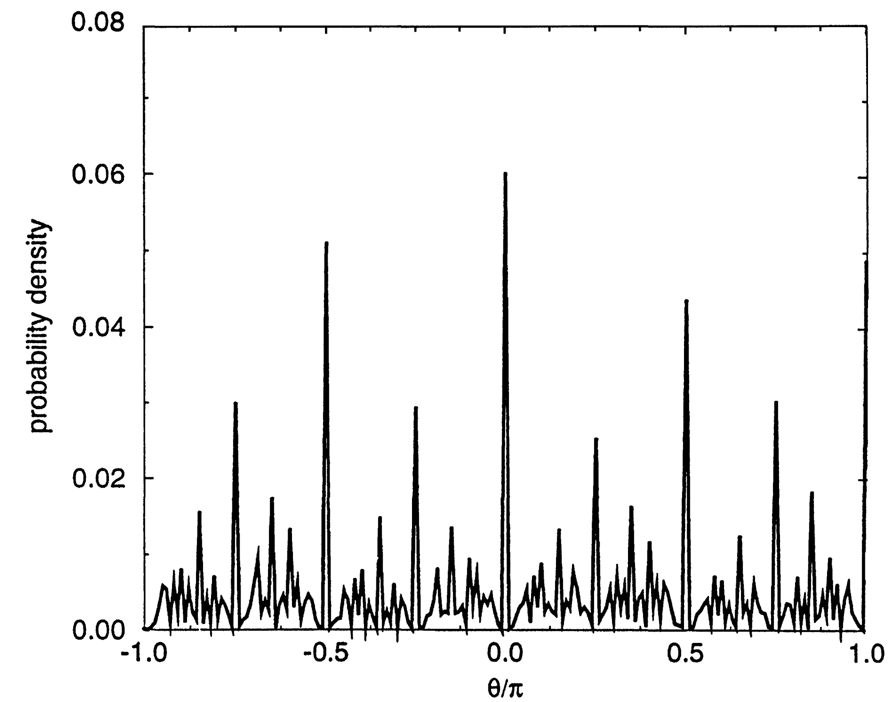

FIGURE 3 "Theta histogram" (angular distribution in delay space) of Weyerhauser stock returns.

partitioned into 201 bins of width 0.01. (Details are in a footnote.) The vertical axis shows the relative frequency of points in each bin. For each θ/π value, it shows the number of points in that bin divided by the total number of points.[1]

The theta histogram just described is an empirical theta histogram. Before we can engage in hypothesis testing, we need a benchmark with which to compare it. We propose a simple bootstrap to create such a benchmark. To construct a bootstrapped theta histogram, one takes the observed relative frequency of each return in the data under study. Assume each period's return was draw from this distribution, and assume every period's return is independent of every other period. Repeated sampling (with replacement) from the empirical distribution of asset returns allows one to generate a bootstrapped theta histogram. Figure 4 illustrates. (We have also constructed empirical and bootstrapped "extensional histograms" showing the number of points at each distance from the origin. These are not shown here.)

Hypothesis tests can be conducted by comparing the empirical and bootstrapped theta histograms. In the tests we have devised, the null hypothesis, H_0, is that the R_t are statistically independent. Under the null hypothesis, some values of R_{t+1}/R_t are more likely than others. If we had an infinite pool of identically, independently distributed returns, each ratio R_{t+1}/R_t would have a given relative frequency. Thus, each

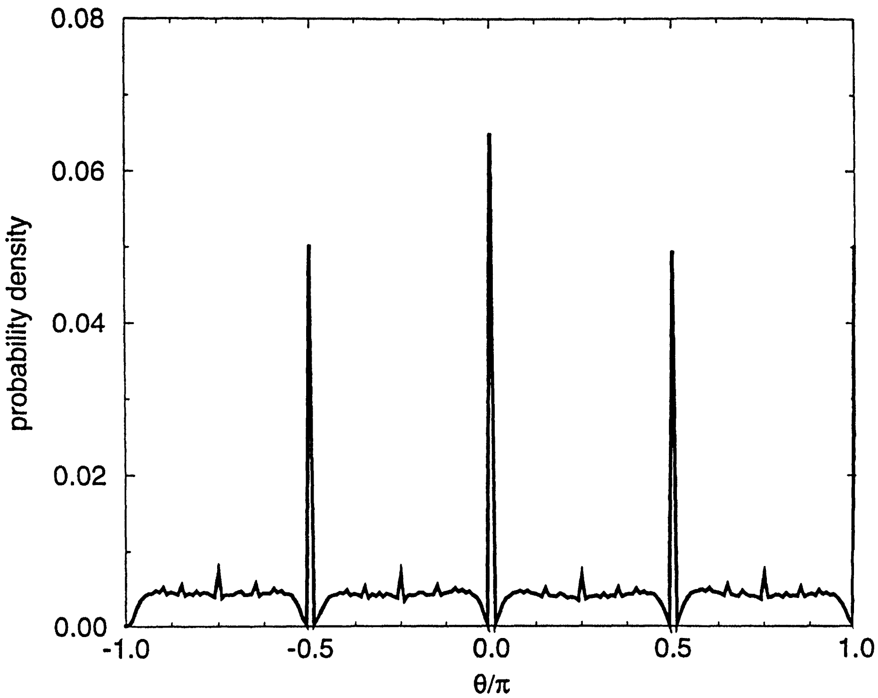

FIGURE 4 Bootstrapped theta histogram created using the same data used to construct the empirical theta histogram of Figure 3.

[1] Since the θ/π values 1.0 and -1.0 correspond to the same point on the compass rose, it was necessary to split the bin straddling this point in order to avoid double counting. (Other solutions would have created programming difficulties.) Thus, the partition contains an initial bin straddling zero, $[-0.005, 0.005]$. It includes 99 bins of width 0.01 arrayed on to the right and another 99 bins of width 0.01 arrayed on the left. The 100th bin on the right, $[0.995, 1.0]$, has width 0.005. Similarly, the 100th bin on the left, $[-1.0, -0.995]$, has width 0.005. This makes a total of 201 bins.

corresponding θ/π value would have its relative frequency. For a given distribution, statistical independence implies a certain relative frequency of "hits" for each ray of the compass rose; it implies a certain profile for the theta histogram. If the observed relative frequencies of a sample are close to the hypothetical values, we have no reason to reject the null hypothesis of statistical independence. If, on the other hand, the observed relative frequencies of some sample were sufficiently far from the hypothetical values, we should reject the null hypothesis.

An unrealistically simple hypothetical example provides an easy illustration. Assume we had a sample of 100 returns in which one half of the returns were 0.001 and the other half were 0.002. Under the null hypothesis of statistical independence, we would expect about half the points of the compass rose to fall along the 45-degree line, *i.e.*, the ray bisecting the positive quadrant. For this ray, $\theta/\pi = 0.25$. The remaining points would be about evenly split between the rays at $\theta/\pi = 0.125$ and $\theta/\pi = 0.375$, *i.e.*, those at 22.5 degrees and 67.5 degrees. The bootstrapped theta histogram of our sample would have three spikes. The relative frequency for the spike at $\theta/\pi = 0.25$ would be one half. The relative frequency for the other two spikes would be one fourth each.

Since our hypothetical sample contains 100 returns, it gives us 99 points of the compass rose. Assume we observed, say, 45 points at $\theta/\pi = 0.25$, 26 at $\theta/\pi = 0.125$, and 28 at $\theta/\pi = 0.375$. In this case, the empirical theta histogram would be similar to the bootstrapped theta histogram. We would have no cause to reject the null hypothesis. The tests we describe below would not reject H_0 for this sample. A different sample, however, might yield a different result. We would reject the null if the original returns had come in the following sequence: 0.001, 0.002, 0.001, 0.002, ... In this case statistical dependence is obviously present. The empirical theta histogram would have one spike of approximate length one half at $\theta/\pi = 0.125$ and another of the about same length at $\theta/\pi = 0.375$. (Details are in a footnote.) The compass rose would have no points along the 45-degree line instead of the 50 or so to be expected under the null hypothesis. The tests we describe below would reject H_0 for this sample.[2]

In the example just given, statistical dependence caused us to reject H_0. The statistical dependence considered would also have shown up as (negative) linear correlation. But correlation and dependence are distinct. If a sufficiently large fraction of the points of the compass rose accumulate along the main diagonals, forming an X pattern, there is statistical dependence, but zero correlation. We show below that X-skewing exists in the data. We now explain our tests more carefully.

Let n be the number of observations. That is, n is the number of R_{t+1}/R_t ratios in our sample and $n+1$ is the number of returns. (We are no longer using n_t to denote the number of ticks by which price moved on day t.) Consider a narrow interval of a ray θ/π, namely, $\omega \pm \delta\omega$. Let p be the relative frequency of points in that interval under H_0. One reads p off of the bootstrapped theta histogram. Given the sample size n and the null hypothesis of independence, the expected number of points in $\omega \pm \delta\omega$ is np. Let k denote the observed number of points in the interval. Define χ^2_{obs} as follows:

$$\chi^2_{\text{obs}} \equiv \sum_{\omega} \frac{(k - np)^2}{np} \tag{3}$$

where $\sum_{\omega} k = n$. Assuming H_0, in the limit of a large number ν of partitions, the complement cumulative distribution of χ^2_{obs} is $Q(\chi^2|\nu)$, an incomplete gamma function (Press *et al.*, 1992). Selecting the customary confidence level of 0.05,

[2] In this example there are 99 points of the compass rose. Thus, they cannot be split evenly between the two rays. The precise value for $\theta/\pi = 0.125$ would be 50/99, slightly more than one half; the precise value for $\theta/\pi = 0.375$ would be 49/99, slightly less than one half. These values would be reversed, of course, if the sample had been 0.002, 0.001, 0.002, ...

we reject H_0 if

$$P(\chi^2 \geq \chi^2_{\text{obs}}) \equiv Q(\chi^2|\nu) = \frac{1}{\Gamma(\chi^2_{\text{obs}})} \int_{\nu}^{\infty} e^{-t} t^{(\chi^2_{\text{obs}}-1)} dt < 0.05, \tag{4}$$

where $\Gamma(x)$ is the gamma function.

This χ^2 test probes the distribution as a whole. One may wish to know, however, if a particular ray or subset of rays of the compass rose has more (or fewer) points than would be expected under the null hypothesis of statistical independence. For example, to test for X-skewing we want to know, in effect, if the four rays making up the main diagonals of the compass rose have collected "too many" points. Let us consider the test as it would apply to one ray. The extension to groups of rays will be obvious. Consider, then, the interval $\omega \pm \delta\omega$. We can easily determine the relative frequency of theta histograms of size n for which the number of hits in that interval, k, is at least equal to the number, h, observed in our sample. Assuming H_0, we have a sequence of Bernoulli trials in which the probability of a hit is p. For every integer k such that $0 \leq k \leq n$, there are $C(k,n) \equiv n!/[k!(n-k)!]$ ways you could get k hits. The probability of exactly k hits is $C(k,n)p^k (1-p)^{(n-k)}$. Thus, the probability of $k \geq h$ is

$$P(k \geq h) \equiv B(h) = \sum_{k=h}^{n} C(k,n)p^k(1-p)^{(n-k)}. \tag{5}$$

We reject H_0 if $P(k \geq h) < 0.05$.

We can use our tests on the Weyerhauser data studied by Crack and Ledoit. We used CRSP data on daily returns from the period studied by Crack and Ledoit, namely, December 6, 1963 to December 31, 1993. This gives us 7,559 returns and 7,558 points of the compass rose. As explained earlier, our data are partitioned into 201 bins of width 0.01. Comparing Figures 3 and 4 suggests that the empirical and bootstrapped theta histograms are not the same. Figure 5 superimposes the two

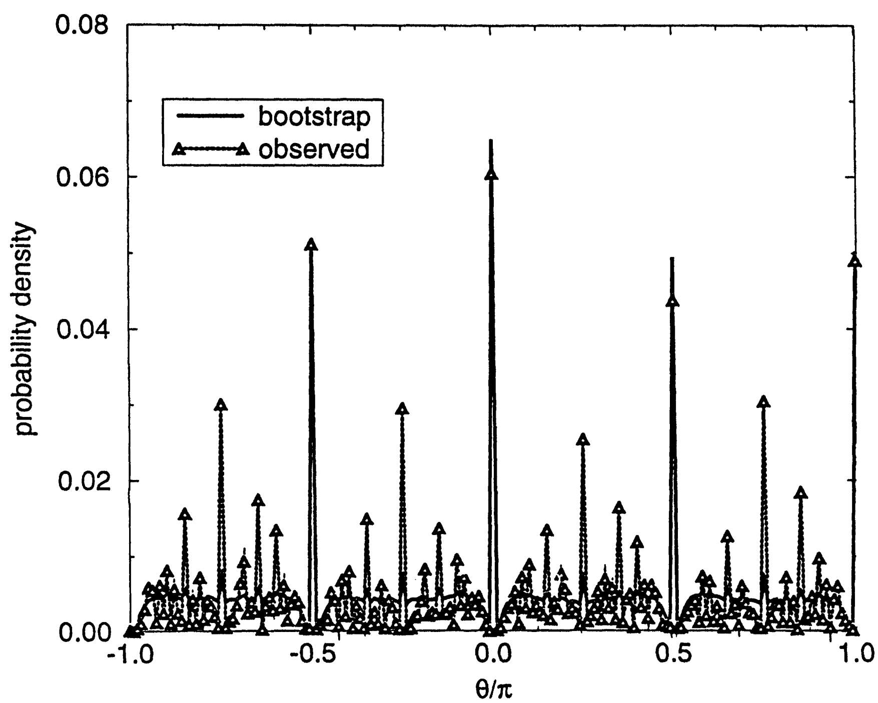

FIGURE 5 Figures 3 and 4 superimposed.

histograms. It appears that the two histograms are significantly different and that the empirical histogram shows more points at the θ/π values ± 0.75 and ± 0.25. These are the values that correspond to the main diagonals of the compass rose.

The χ^2 test shows that Weyerhauser returns are not independent. The number of degrees of freedom, ν, is equal to the number of bins, 201. Since ν is large, we could rely on the asymptotic distribution to carry out our test. We have

$$P(\chi^2 > \chi^2_{\text{obs}}) \equiv Q(\chi^2_{\text{obs}} \mid \nu) \approx Q(x) \equiv 1/\sqrt{2\pi} \int_x^\infty \exp(-t^2/2)dt \quad (6)$$

where $x = \sqrt{2\chi^2_{\text{obs}}} - \sqrt{2\nu - 1}$ (Abramovitz and Stegun, 1972, formula 26.4.13). The χ^2 test is reported in Table I. For this data, $\chi^2_{\text{obs}} = 5019.2$ *i.e.*, much greater than the expected value of 201. The reduced normal variable $x = 80.2$, then we can again use an asymptotic formula for large x (Abramovitz and Stegun, 1972, formula 26.2.12), $Q(x) \approx 1/\sqrt{(2\pi)}\exp(-x^2/2)/x$, which corresponds to an astronomically negligible value for $Q(x)$. Therefore we reject the null hypothesis of statistical independence.

TABLE I χ^2 hypothesis test for Weyerhauser data. The empirical series includes $n = 7558$ observations histogrammed in $\nu = 201$ bins. $x = \sqrt{2\chi^2_{\text{obs}}} - \sqrt{2\nu - 1}$ is the reduced variable which is asymptotically distributed normally, $Q(x)$ is the corresponding probability $P(x' \geq x)$

Series	χ^2	x	$Q(x)$
Weyerhauser	5019.2	80.2	$\approx 10^{-1392}$

This result of the χ^2 test is not entirely surprising. If the Weyerhauser data seem to exhibit ARCH phenomena, they are probably not statistically independent even if first-order autocorrelation is zero.[3] Our χ^2 test is not guaranteed to pick up all forms of statistical dependence. Information contained in the extensional histogram, for instance, will not show up in the theta histogram. Nevertheless, if standard tests indicate that the data are autoregressive in conditional heteroskedasticity, it is not surprising that our χ^2 test should indicate statistical dependence.

We noted that Figure 5 seemed to indicate a heavy accumulation of points at $\theta/\pi = \pm 0.75$ and ± 0.25. If these rays collect significantly more points than would be consistent with the null hypothesis of independence, then we have X-skewing in the Weyerhauser data. Results reported in Table II confirm the presence of such a pattern. The compound probability of a point being with the intervals $\pm 0.25(\pm 0.005)$ and ± 0.75 (± 0.005) is $p = 0.027$ under H_0. The expected number of points in those intervals is therefore $np = 206.5$. The observed number is $h = 872$. From Eq. (5) we obtain $B(h) < 10^{-271}$, a value below our confidence level of 0.05. We reject the null. It seems as if the main diagonals are accumulating "too many" points. Weyerhauser returns are X-skewed.

Crack and Ledoit repeat in a footnote their referee's comment that they have not shown "that there are no other predictable structures present" in the compass rose. X-skewing is an example. The existence of this pattern shows that Crack and

TABLE II Bernoulli hypothesis test for Weyerhauser data, $\theta/\pi = \pm 0.25(\pm 0.005)$ and $\pm 0.75(\pm 0.005)$. In this and the following Bernoulli tables, n is the number of points in the empirical distribution, p is the probability that θ/π belongs to the interval considered according to the bootstrapped histogram, $k = np$, h is the observed number of points in the interval considered, $\Delta k/\sigma = |h - k|/\sigma$ is the normalized fluctuation observed, and finally $B(h) = \sum_{k=h}^{n} P_k$ (where P_k is the Bernoulli probability of k observation in the bin) is the probability $P(k \geq h)$

Series	n	p	k	h	$\Delta k/\sigma$	$B(h)$
Weyerhauser	7558	0.02732 (± 0.00013)	206.5(± 1.0)	872	665.5	$< 10^{-271}$

[3] If ARCH tests are biased by discreteness, we cannot rule out the possibility that standard ARCH procedures may sometimes falsely indicate that statistically independent returns are autoregressive in conditional heteroskedasticity.

Ledoit were mistaken to claim that all of the information of the compass rose "is contained in decades-old studies on the time-series properties of stock returns" (p. 755). These studies make distributional assumptions which may be false and may mask regularities more easily detected by other techniques of analysis.

X-skewing is distinct from the ARCH behavior of standard models. But it induces standard tests to indicate ARCH behavior. The next section reports on a simulation that produces X-skewing and gives the (false) appearance of standard ARCH behavior.

III. A SIMPLE SIMULATION OF X-SKEWING

One of our tests indicated X-skewing for Weyerhauser returns. Tests on other data suggest this X-skewing is common. This pattern may explain the widespread phenomenon of ARCH behavior in return data. To test this possibility, we ran a simple simulation. Our simulation was designed to see if X-skewing can cause standard procedures to indicate the presence of ARCH. It was not designed to model stock returns accurately. It seems that the X-skewing we found in Weyerhauser data can indeed induce the false appearance of standard ARCH behavior.

Our simulation started with an initial price, $p(0)$, and return $r(0)$. It then drew from a normal distribution centered about $r(0)$. If the value drawn exceeded a predetermined threshold level, it was discarded and a new value is drawn from the same distribution. (Without a cutoff value, the variance would grow without bound. Since our simulation is merely illustrative, we simply imposed a cutoff.) With a probability of 0.5, the sign of the value drawn was reversed. This value became a tentative return. The tentative return was added to the old price to get a new price. The new price was rounded up or down to simulate the existence of a postive tick size. From the new price a new return was calculated. The simulation then used these new values to calculate, in the same way, a price and return for the following period. And so on.

Each period's return was gotten by first drawing from a normal distribution centered about last period's return, then changing the sign with probability one half, and, finally, adjusting to create discreteness in the associated price series. If the return in period t, $r(t)$, was large and positive, the absolute value of the return in period $t+1$, $|r(t+1)|$, was likely to be large too. The return, $r(t+1)$, was just as likely to be negative, however, as positive. The distribution of the $r(t+1)$, given $r(t)$, is bimodal. One mode occurs at $r(t)$, the other at $-r(t)$. The distance between the modes is greater the larger $|r(t)|$. Only when $r(t)=0$ is the conditional distribution of $r(t+1)$ unimodal.

Since the sign of each period's return has an equal chance of being positive or negative, the expected value of $r(t)$ is zero for all t. The returns generated by this simulation have no autocorrelation in first moments. There is autoregression in conditional variance. But the series is not generated by the processes described in standard ARCH models, which assume a unimodal distribution.

Figure 6 shows the compass rose for this simulated series.

The series generated by this simulation passes standard tests for the existence of ARCH behavior. Results of the Q and LM tests reject the null of no autocorrelation of conditional variance at the 0.0001 confidence level. We fit a GARCH(1, 1) model to the data. The results, reported in Table III, seem to indicate ARCH behavior. The SAS subroutine we ran indicated a statistically significant ARCH(1) coefficient and a statistically insignificant GARCH(1) coefficient. An unwary researcher might conclude that the series is an ARCH(1) process.

The purpose of our simulation is not to model stock returns, but to illustrate two points. First, X-skewing may be generated by a stochastic process which is not in the customary family of ARCH models. Second, X-skewing nevertheless causes standard tests to indicate ARCH behavior. The relationship between ARCH and X-skewing merits further study.

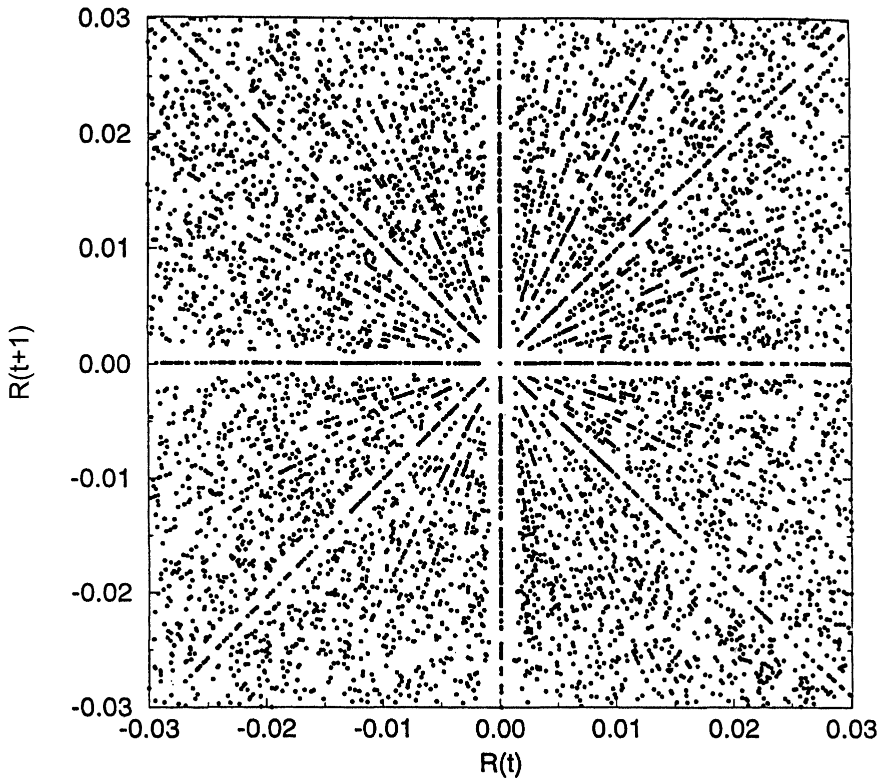

FIGURE 6 Compass rose graph for a simulated series.

TABLE III Results of SAS fit of our simulation data to a GARCH(1, 1) model. The results falsely suggest that the data follow an ARCH(1) process

Variable	DF	*B* Value	Std. error	*t* ratio	Approx. prob.
Intercept	1	0.000256	0.000183	1.399	0.1618
ARCH(0)	1	0.000224	0.000048	4.670	0.0001
ARCH(1)	1	0.128261	0.0283	4.537	0.0001
GARCH(1)	1	0.011617	0.1852	0.063	0.9500

X-skewing is an unintended consequence of human action. The human actions that generate it are responses to the prevailing environment. If the regime governing the market process should change, the degree of X-skewing may change too. We show in next section that certain changes in government policy cause changes in the interpretive environment of asset markets and thus in the amount of X-skewing.

IV. BIG PLAYERS INDUCE X-SKEWING

We apply our techniques to an important episode in nineteenth-century Russian monetary history.[4] This episode has been studied in the past as a test case of the theory of "Big Players" (Koppl and Yeager, 1996; Broussard and Koppl, 1999). We show here that Big Players induce X-skewing by corrupting the interpretive environment of the

[4] The story we tell is related in Koppl and Yeager, 1996, which we follow closely. The data we analyze were gathered by Yeager. They has been studied by Broussard and Koppl, 1999; Koppl and Yeager, 1996 and Yeager, 1969 and 1984.

market. (Our results also show that the compass rose is present in a class of assets other than stocks, namely, foreign exchanges. We do not know if it is always present or only sometimes. Here is another area for further study.) Some discussion of theory and terminology is need before we can turn to our case study.

When traders observe a price change, earnings announcement, or any other potentially relevant event, they must interpret the signal. The signal must be given a meaning. This interpretted meaning encourages the trader to buy, sell, or hold. The meaning given is a function of many factors including the peculiarities of each individual trader. Among the factors is the general market environment. The general environment may be influenced by government policy. An ill-chosen policy may cause the general market environment to produce poor signals. The signals become hard to decipher. When this happens, we may say that the policy has corrupted the interpretive environment of the market. We argue presently that Big Players have this corrupting effect. Big Players create changes in the interpretive environment of asset markets that encourage herding and contra-herding. Contra-herding is simply the contrarian policy of acting against the trend.

Big Players are defined by three properties. First, they are big. Their decisions influence the course of market events. Second, they are insensitive to profit and loss. A nation's central bank, for instance, cannot be weeded out of the market by losses resulting from bad choices. Third, they act on their discretion. They are not bound by any fixed rule of behavior. A central bank is a Big Player when it practices discretion in monetary policy. If it follows a rule to, say, expand the money supply at a fixed annual rate, then it is not a Big Player. It is not a Big Player because it chooses rules over discretion. An activist central bank is the proto-typical Big Player. Private actors are normally sensitive to profit and loss. They are not likely to be Big Players.

Building on Scharfstein and Stein (1990); Koppl and Yeager show that Big Players may create bandwagon effects. (See Fig. 7) Scarfstein and Stein argued that portfolio managers have an incentive to imitate one another. If things go well, fine. If losses are incurred, one may share the blame with others. A manager's reputation

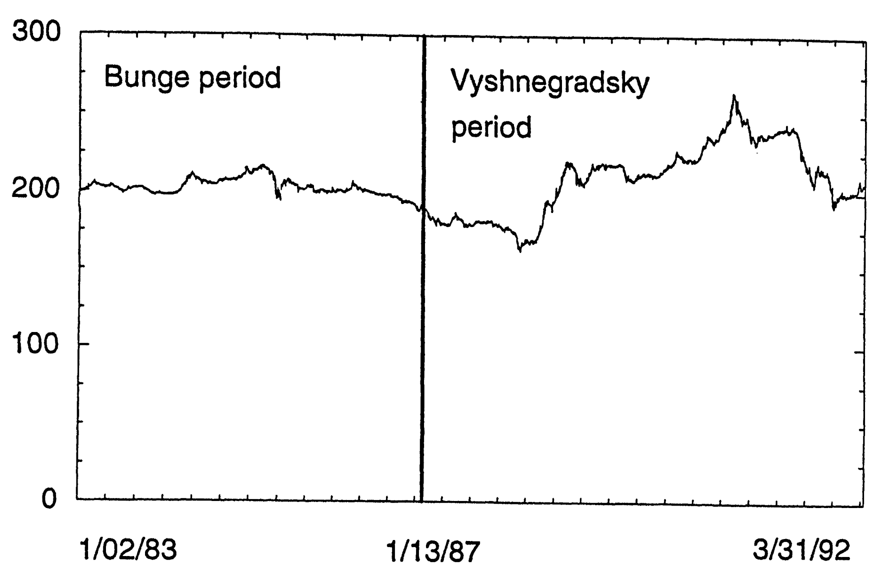

FIGURE 7 From Koppl and Yeager, 1996 (ruble's exchange rates in German marks per 100 rubles).

depends more upon his relative performance than his absolute performance. Thus, acting idiosyncratically is risky, herding is safer. Koppl and Yeager point out that this incentive to herd is stronger when traders' knowledge of fundamentals has been clouded by Big Player interventions. The Big Players' actions often override fundamentals. (One might prefer to say "other fundamentals".) These actions must be anticipated to the extent possible. But because Big Player actions are based on discretion rather than any fixed rule, they are hard to anticipate. Big Players act in surprising ways. Big Players thus increase the chances of failure for those who do not follow the herd. The same is true for those who do follow the herd; but they can share the blame.

Without Big Players it is easier for some traders to made idiosyncratic choices yielding returns that are satisfactory or above average. When Big Players enter, fundamentals are harder to read and indiosyncratic choices grow more dangerous. When Big Players tred, prudence more strongly suggests following the crowd and sharing the blame should things go awry. Thus, Big Players induce herding.

The root cause of the increased herding under Big Players is ignorance. By scrambling market signals, Big Players reduce the value of the market information little players use to make their choices. In this reduced state of knowledge, little players are more likely to study past price changes as a clue to the future course of prices.

Broussard and Koppl (1999) extend the point. They argue that ARCH-like effects are likely to be stronger under Big Players. "By reducing the value of all information", they explain, "Big Players increase the relative value of information about an asset's recent price behavior". An unusually large price change becomes the object of competing interpretations. "Some will see a trend, others will expect reversal. Whichever view happens to gain more adherents, the exaggerated attention paid to the price movement encourages another large movement" to follow. Broussard and Koppl argue, in effect, that when Big Players induce herding, some traders will become contrarians who always expect the trend to be reversed. This is contra-herding.

Big Players reduce the value of all information traders might use except, perhaps, information about the Big Players themselves. When the value of all sources of information is reduced in this way, the relative value of information about past price behavior increases. More trading is conducted on the basis of such information. Herding and contra-herding result. X-skewing results from the simultaneous presence of herding and contra-herding. The history of the credit ruble gives us a good test case for the predictions the Big Players theory.

From the Crimean War (1853–56) to 1897, Russian had a paper currency which floated against other currencies, including the Germany mark, a gold-standard currency. This was the period of the "credit ruble". During most of this period, the Russian finance ministry actively intervened in the foreign exchange market, hoping to influence the ruble's exchange rate. A notable exception was the period of Nicolai Bunge's tenure as finance minister. Bunge served from May 18, 1881 to January 13, 1887. He was a principled non-interventionist.

Bunge's successor, Ivan Vyshnegradsky, was very different. He served from January 14, 1887 to September 11, 1892. Vyshnegradsky was a highly active interventionist who meddled frequently in the Berlin market. Vyshnegradsky seemed to derive great pleasure from getting the better of the Berlin speculators in the ruble.

The contrast between Bunge and Vyshnegradsky is unusually clear case of a change in regime from a simple policy rule of non-intervention to an activist, discretionary policy. It is thus a test case for the "Big Player" theory. Nicolai Bunge was not a Big Player because he maintained a principled non-interventionist stance; he did not exercise his discretion. Vyshnegradsky did use his discretion and was thus a Big Player.

Our data is constructed from a series created by Leland Yeager. Yeager used two contemporary German newspapers to find the ruble exchange rate in German marks per 100 rubles of bank notes. His data cover the period from January 02, 1883 to March 31, 1892. The ruble's price moved in discrete jumps of 0.05 or multiples thereof. We do not know if this effective tick size corresponds to an official tick. Nevertheless, it exists and is of the same order of magnitude as price changes. Thus, returns for this data exhibit the compass rose pattern. We created a return series from this data by taking the forward first differences of prices and normalising by the price (*i.e.*, $R_i \equiv (P_i - P_{i-1})/P_{i-1}$).

Koppl and Yeager give evidence from R/S analysis to show that herding increased under Vyshnegradsky. (See Fig. 7) Broussard and Koppl (1999) fit a modified GARCH(1, 1) model to the data and find an (apparent) increase in ARCH effects under Vyshnegradsky. (See Fig. 8) Using the same data, we find X-skewing under Vyshnegradsky, but no X-skewing under Bunge. This is and other results are reported below.

From our return series we calculated empirical and bootstrapped theta histograms. The total number of samples in the bootstrapped histograms was 1000 times the length of the original series. The data points have been binned into 201 partitions with a resolution of $\theta/\pi = 0.01$. The probability associated with each bin is estimated to be $p = k/n$ with an error given by the corresponding (approximate) bernoullian standard deviation $\sigma \approx \sqrt{np} = \sqrt{k}$.

Figures 9 and 10 show the empirical theta histograms from the Bunge and Vyshnegradsky periods. Figures 11 and 12 show their corresponding bootstrapped theta histograms. For the Bunge period, the empirical and bootstrapped histograms are almost identical. For the Vyshnegradsky period, they differ. The asymmetry in the Vyshnegradsky period is evident. Especially evident is the large number of points accumulated at $\theta/\pi = -0.5$. Days in which the ruble's exchange rate did not change tended to be followed by days in which its value fell. We don't know why.

For each period, we tested the hypothesis that the probability of a point at $\theta/\pi = -0.5$ is equal to the relative frequency of such points when returns are independent. The results are reported in Table IV.

For the Bunge period, the number of points n in the sample was 1224. If returns were independent, the probability of a point at $-0.5(\pm 0.005)$ would be $p = 0.0307$ and the expected value of the

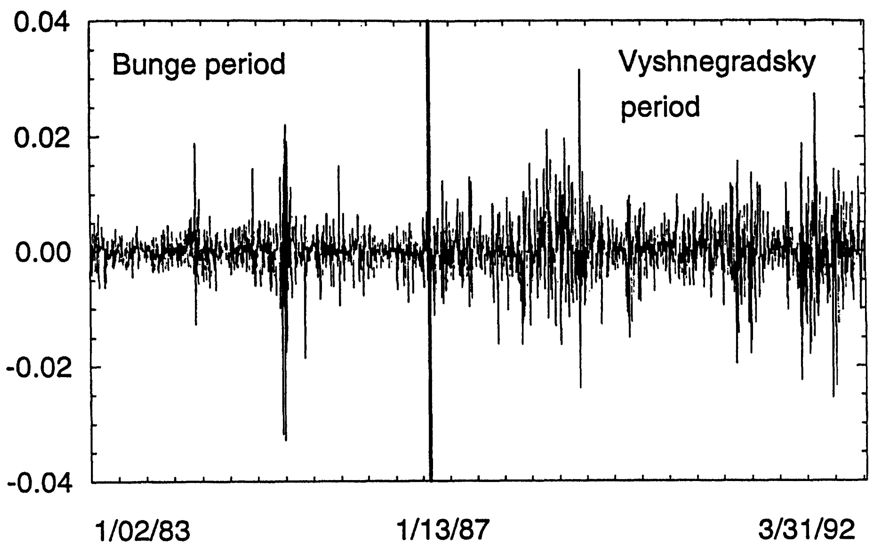

FIGURE 8 From Broussard and Koppl, 1996 (ruble's price percent changes).

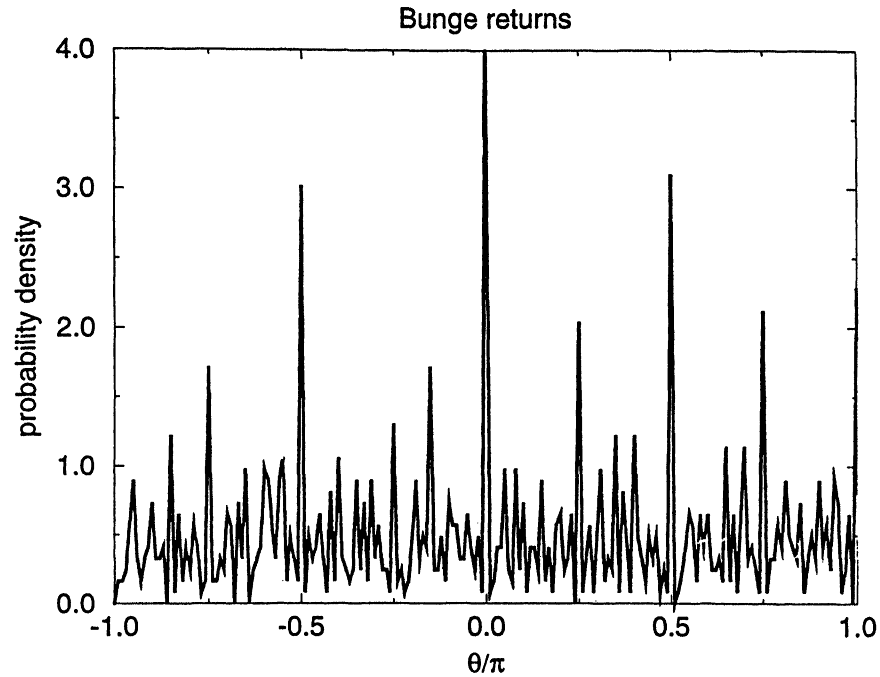

FIGURE 9 Empirical theta histogram, Bunge period.

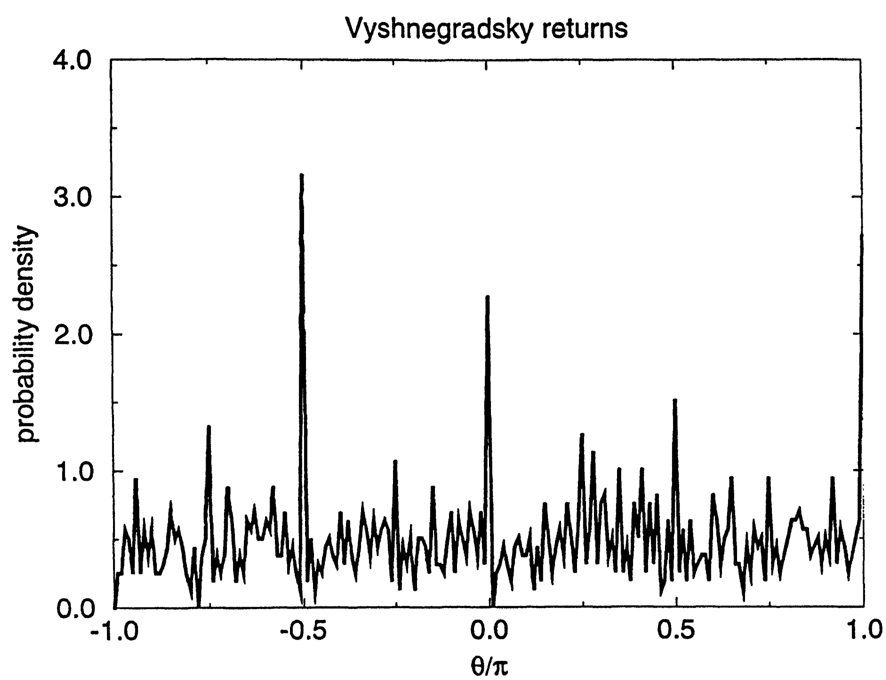

FIGURE 10 Empirical theta histogram, Vyshnegradsky period.

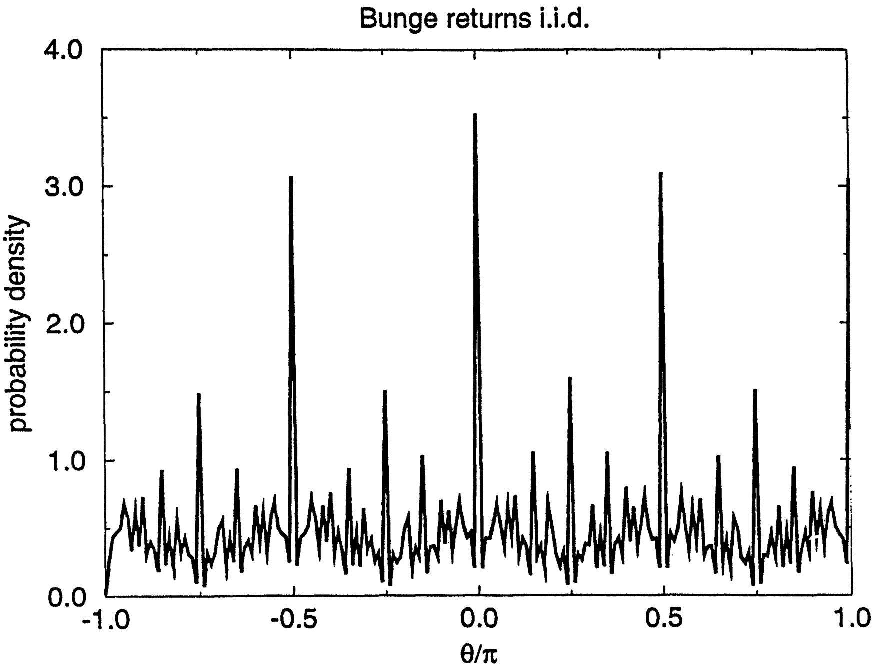

FIGURE 11 Bootstrapped theta histogram, Bunge period.

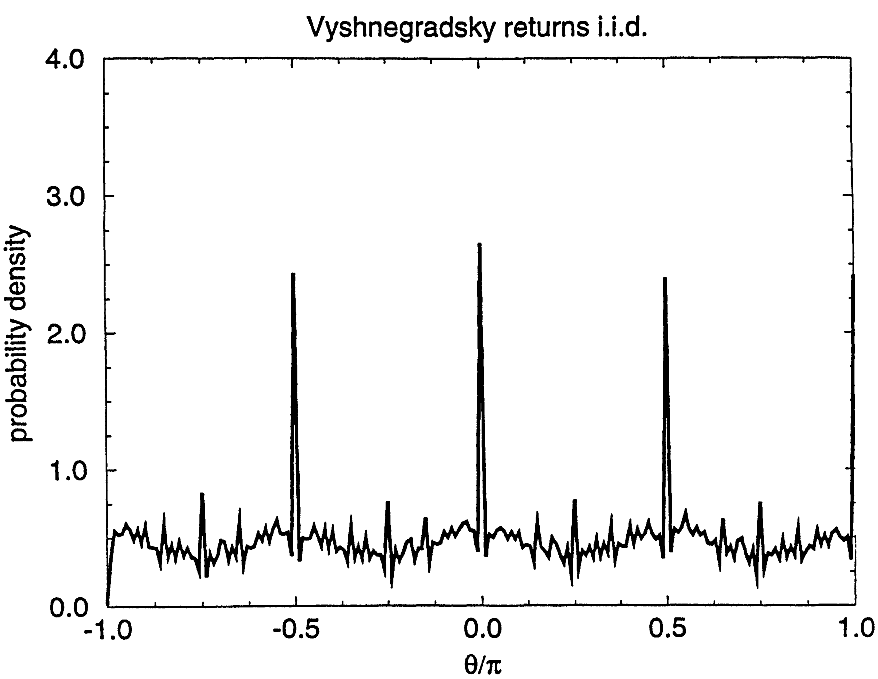

FIGURE 12 Bootstrapped theta histogram, Vyshnegradsky period.

TABLE IV Bernoulli hypothesis tests for ruble data, $\theta/\pi = -0.5(\pm 0.005)$

Series	n	p	k	h	$\Delta k/\sigma$	$B(h)$
Bunge	1224	0.03072 ($\pm$ 0.00016)	37.60($\pm$ 0.20)	37	0.10	> 0.5
Vyshnegradsky	1581	0.02436 ($\pm$ 0.00012)	38.51($\pm$ 0.19)	50	1.85	0.041

number of points would be $k = np = 37.6$. The actual number was $h = 37$, well within the standard deviation for a Bernoullian ($\sigma \approx \sqrt{np} = 6.1$). We are therefore unwilling to reject the hypothesis of independence. In other words, for the Bunge period we do not have so many points accumulating at -0.5 that we wish to reject the null hypothesis of independence in returns.

For the Vyshnegradsky period, the number of points in the sample was $n = 1581$. If returns were independent, the probability of a point at -0.5 (± 0.005) would be $p = 0.0244$ and the expected value of the number of points would be $k = 38.5$. The actual number was $h = 50$. From Eq. (5) we calculate the probability to get such k or higher. This probability is $B(h) = 0.041$. Since $B(h)$ is less than 0.05, our confidence level, we reject the hypothesis of independence. For the Vyshnegradsky period, under the assumption of independence, we have an improbably large number of points accumulating at -0.5. Our hypothesis test supports the conclusion one is likely to draw from looking at Figure 10. Days in which the ruble's exchange rate did not change tended to be followed by days in which its value fell.

The Big Players theory suggests we should find another difference between the Bunge and Vyshnegradsky periods. Under Vyshnegradsky, there should be X-skewing. There should be a greater tendency for points to accumulate at $\theta/\pi = \pm 0.25$ and $\theta/\pi = \pm 0.75$. These are the values corresponding to the two main diagonals of the compass rose pattern. We expect the Big Player influence of Vyshnegradsky to encourage traders to pay more attention to price history, because all sources of information have been degraded by the Big Player's discretionary interventions. A large price change today will become the subject of interpretation in which some see a trend and other expect "correction". Whichever theory becomes more popular, a large-magnitude return today is likely to be followed by a return of similar magnitude, though not necessarily in the same direction.

Our confidence in this result is strengthened by an inspection of Figures 13 and 14. These figures show the compass rose pattern for the absolute value of returns. Since absolute values are non-negative, all points appear in the positive quadrant. Under Bunge, no ray is obviously accumulating too many or too few points. Note that the graph shows something close to a grid, with little centrifugal smudging. This is because the ruble exchange rate did not vary widely during Bunge's tenure as finance minister. (This is exactly the result predicted by Crack and Ledoit.) Under Vyshnegradsky, the 45-degree line seems to have collected more points than it would have if returns were independent. This difference between the Bunge and Vyshnegradsky periods is confirmed by hypothesis tests reported in Table V.

In this case, in the Bunge period the compound probability for θ/π being within the intervals $\theta/\pi = \pm 0.25(\pm 0.005)$ and $\theta/\pi = \pm 0.75(\pm 0.005)$ is $p = 0.0610$, if returns were independent. Accordingly, the expected number of points is $k = 74.6$, while we observe $h = 88$. From Eq. (5) we calculate the probability to get such h or higher. This probability is $B(h) = 0.065$. Since $B(h)$ is more than 0.05, we are unwilling to reject the hypothesis of independence. For the Bunge period, under the assumption of independence, we do not have an improbably large or small number of points along the 45-degree line of Figure 13.

There is no X-skewing under Bunge. Broussard and Koppl do find statistically significant ARCH

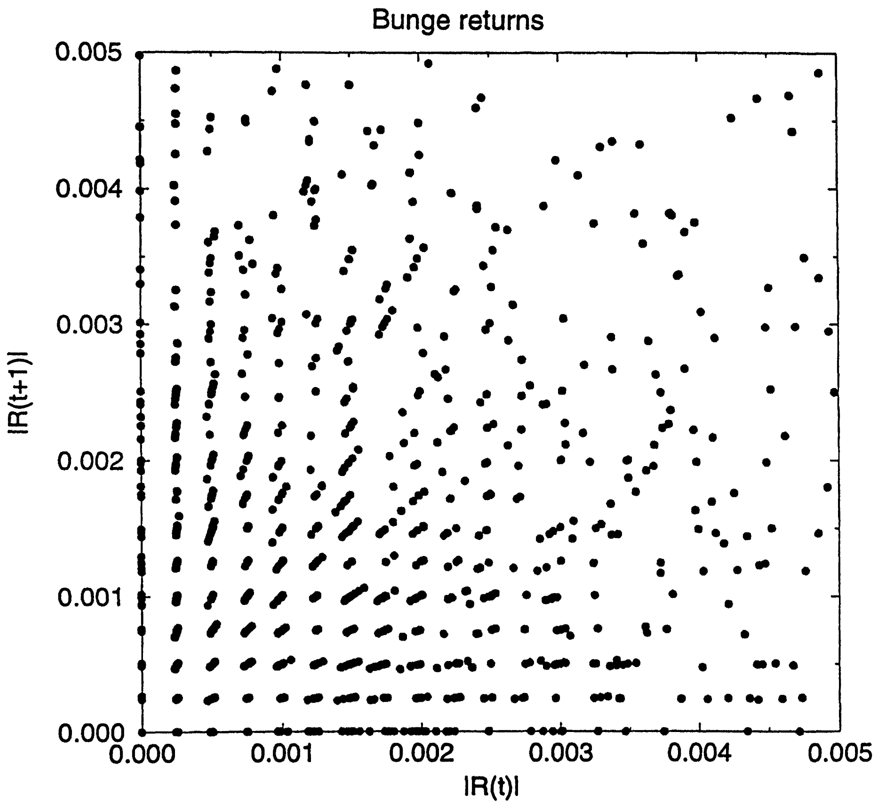

FIGURE 13 Compass rose in absolute values, Bunge period.

effects during this period. This pair of results is futher evidence that X-skewing and ARCH behavior are distinct even though they are related. We hope to clarify their relationship in future research.

For the Vyshnegradsky period, $p = 0.0312$ and $k = 49.4$, while we observe $h = 73$. Again from (5) we obtain $B(h) = 0.0008$, a value much smaller than our level of confidence. We can therefore reject the hypothesis of independence. For the Vyshnegradsky period, under the assumption of independence, we have an improbably large number of points accumulating at ± 0.25 and ± 0.75. Our hypothesis test supports the conclusion one is likely to draw from looking at Figure 14. Large changes in the exchange rate on one day tend to be followed by similarly large changes the next day, though not necessarily in the same direction. This tendency is consistent with the theory of Big Players.

Finally, for each period, we tested the hypothesis of independence among returns using the χ^2 test. In order to avoid effects related to the number of points in the sample, we choose $n = 998$ for each series. A "mixed" series has been studied, taking the end of the Bunge period and the beginning of the Vyshnegradsky period of tenure, in equal proportions. The number of degrees of freedom ν coincides in our case with the number of bins, *i.e.*, 201. We use again the asymptotic formula 6 (Abramovitz and Stegun, 1972).

The χ^2 test results are reported in Table VI. The test is in agreement with the previous tests on specific θ/π values. For the Bunge period $\chi^2_{obs} = 203.7$, very close to the expected value $\langle\chi^2\rangle = \nu = 201$. Therefore $Q(x)$ is quite large

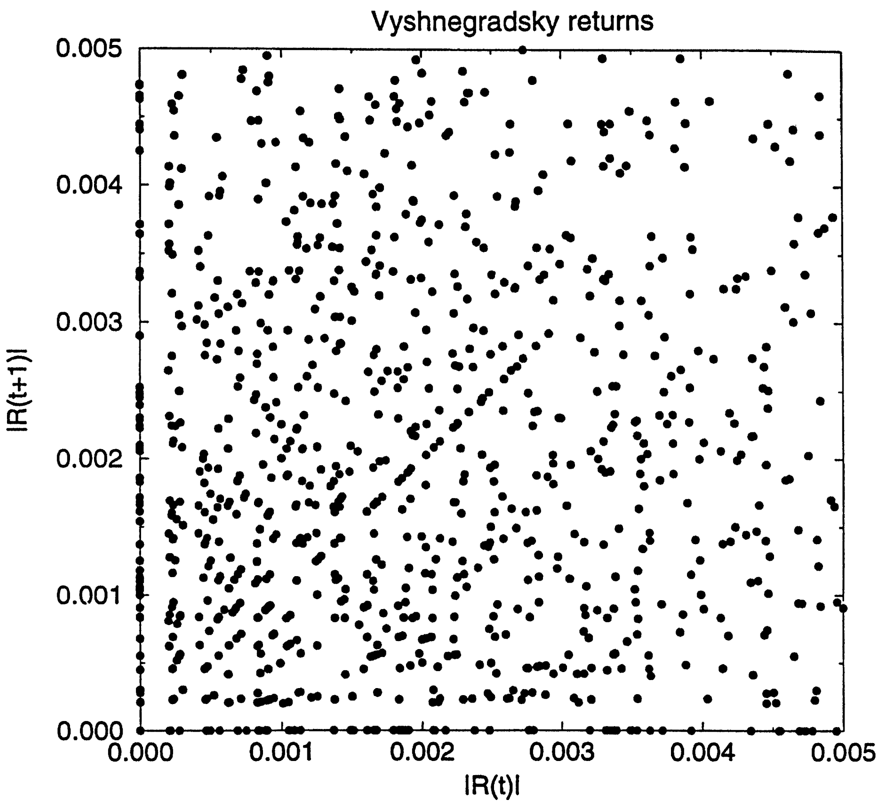

FIGURE 14 Compass rose in absolute values, Vyshnegradsky period.

TABLE V Bernoulli hypothesis tests for ruble data, $\theta/\pi = \pm 0.25(\pm 0.005)$ and $\pm 0.75(\pm 0.005)$

Series	n	p	k	h	$\Delta k/\sigma$	$B(h)$
Bunge	1224	0.06099(± 0.00022)	74.66(± 0.27)	88	1.54	0.065
Vyshnegradsky	1581	0.03123(± 0.00014)	49.38(± 0.22)	73	1.85	0.0008

TABLE VI χ^2 hypothesis tests for ruble data. The empirical series include $n = 998$ observations histogrammed in $\nu = 201$ bins. The reduced variable x^* corresponding to the 0.05 confidence level, *i.e.*, $Q(x^*) = 0.05$, is 1.645

Series	χ	x	$Q(x)$
Bunge	203.696	0.159	0.436
Vyshnegradsky	251.501	2.403	0.0082
Mixed	291.102	4.104	0.00002

and well within the confidence level. For the Vyshnegradsky period $\chi^2_{\rm obs} = 251.5$ and $Q\ (x) = 0.008 < 0.05$, therefore we can reject the hypothesis that Vyshnegradsky's empirical theta histogram is distributed as the corresponding i.i.d. returns histogram. This conclusion is even more probable for the "mixed" series, with $\chi^2_{\rm obs} = 291.1$ and $Q(x) = 2 \times 10^{-5}$.

Our tests show that the angular distribution of credit-ruble returns in delay space shifted dramatically under the Big Player influence of Ivan Vyshnegradsky. We could find no adequate evidence of X-skewing under Bunge. X-skewing was clearly present under Vyshnegradsky. Thus, we have evidence that Big Players can induce or increase X-skewing by changing the interpretive environment of an asset market. This X-skewing, in turn, can induce the appearance of ARCH behavior.

V. CONCLUSION

Building on Crack and Ledoit (1996), we have described and applied some new techniques of time-series analysis. We believe several considerations suggest that our techniques and others like them may be of fairly general interest in economics and finance.

Our tests and techniques are designed specifically for discrete data and the compass rose. They are not biased by discreteness. (They could be applied, however, to data that is not discrete.) Although the compass rose has been described as "subjective", our procedures are perfectly objective.

Our discovery of X-skewing gives us another way to look at volatility dynamics. We have found evidence that apparent ARCH effects may be due at least partly to X-skewing of the compass rose. This gives us information on the dependence among returns that is not reflected in ARCH coefficients. Crack and Ledoit suggest that ARCH estimates may be biased by discreteness. If this is true, then it may be desirable to have another technique of analysis capable of getting at volatility dynamics.

Our simulation is evidence that X-skewing may be inconsistent with standard ARCH models even though it induces the appearence of ARCH behavior under standard tests. Further study is needed on the relationship between X-skewing and ARCH. In particular, the conditional distribution of the return on day t, given the return on day $t-1$ should be studied further. We conjecture that this distribution is typically not unimodal and normal as assumed in standard ARCH models.

Finally, our examination of the credit ruble supports the view that the degree of X-skewing may be a partial function of the interpretive environment facing traders. In particular, Big Players encourage X-skewing by scrambling market signals and corrupting the knowledge traders rely on. One way policy regimes influence asset markets is by altering the interpretive environment facing traders. Big Players change the interpretive environment of the market by throwing traders into a state of ignorance. This ignorance ecourages herding and contra-herding, which together induce X-skewing and the appearance of standard ARCH behavior. Thus, the appearance of standard ARCH behavior depends in part on the interpretive environment, which, in turn, depends on the policy regime governing asset markets. We conjecture that the links between policy, interpretation, and the behavior of asset markets is a fruitful area for future research. Theta histograms and other non-traditional tools may be required to further explore these links.

References

Abramowitz, Milton and Irene A. Stegun (Eds.) (1972) *Handbook of Mathematical Functions*, New York: Dover Publ.

Brealy, Richard A. and Stewart C. Meyers (1991) Principles of Corporate Finance. New York: McGraw-Hill, 4th edn.

Brealy, Richard A. and Stewart C. Meyers (1996) Principles of Corporate Finance. New York: McGraw-Hill, 5th edn.

Broussard, John and Roger Koppl (1999) "Big Players and the Russian Ruble: Explaining Volatility Dynamics", *Managerial Finance*, **25**(1), 49–63.

Chen, An-Sing (1997) "The square compass rose: the evidence from Taiwan", *Journal of Multinational Financial Management*, **7**(2), 127–144.

Crack, Timothy F. and Olivier Ledoit (1996) "Robust Structure Without Predictability: The 'Compass Rose' Pattern of the Stock Market", *Journal of Finance*, **51**, 751–762.

Koppl, Roger and Leland Yeager (1996) "Big Players and Herding in Asset Markets: The Case of the Russian Ruble", *Explorations in Economic History*, **33**, 367–383.

Kramer, Walter and Ralf Runde (1997) "Chaos and the compass rose", *Economics Letters*, **54**(2), 113–118.

Press, W. H., Flannery, B. P., Teukolski, S. A. and Vetterling, W. T. (1992) *Numerical Recipes*, Cambridge: Cambridge University Press, 2nd edn.

Yeager, Leland B. (1969) "Fluctuating Exchange Rates in the Nineteenth Century: The Experiences of Austria and Russia", In: Mundell, R. A. and Swoboda, A. K. (Eds.), *Monetary Problems of the International Economy*, Chicago: University of Chicago Press.

Yeager, Leland B. (1984) "The Image of the Gold Standard", In: Michael D. Bordo and Anna J. Schwartz (Eds.), *A Retrospective on the Classical Gold Standard*, 1821–1931. Chicago: University of Chicago Press.