

Robust Structure Without Predictability: The “Compass Rose” Pattern of the Stock Market

TIMOTHY FALCON CRACK and OLIVIER LEDOIT*

ABSTRACT

Plotting daily stock returns against themselves with one day's lag reveals a striking pattern. Evenly spaced lines radiate from the origin; the thickest lines point in the major directions of the compass. This “compass rose” pattern appears in every stock. It is caused by discreteness. However, counter-examples demonstrate that the existence of exchange-imposed tick sizes (e.g. eighths) is neither necessary nor sufficient for the compass rose. The compass rose cannot be used to make abnormal profits: it is structure without predictability. Among other consequences, the compass rose may bias estimation of ARCH models, and tests for chaos.

TECHNICAL ANALYSTS ATTEMPT TO identify recurring “patterns” in stock prices (e.g. Edwards and Magee (1992)). Similarly, chaos researchers embed time series of stock returns into multi-dimensional spaces and attempt to identify low-dimensional “attractors” (e.g. Barnett and Chen (1988), Figure 1). The hope is that with structure comes predictability and abnormal profits. In fact, neither of these literatures has yet discovered indisputable structure.

We follow a procedure similar to a chaos researcher's: we look at a simple scatter diagram of daily stock returns plotted against themselves with one day's lag. We document a structural relationship (the “compass rose”) that is strikingly geometrical. In contrast to the products of technical analysis and chaos theory, the structure we find is indisputably present in every stock.

Unfortunately the compass rose cannot be used for predictive purposes by market participants; it is an artifact of market microstructure. It follows that even a successful search for significant structure in stock market data need not enhance predictability nor lead to abnormal profits. An equally important lesson is that microstructure effects can be spectacular in daily, and even weekly, data.

Several areas for future research are given. These include potential biases caused by the compass rose in the estimation of Autoregressive Conditional Heteroskedasticity (ARCH) models, and in a standard test for chaos.

* Crack is from the School of Business, Indiana University; and Ledoit is from the Anderson Graduate School of Management, University of California at Los Angeles. This research was conducted while both authors were at MIT. We thank Paul Asquith, Kenneth French, Roger Huang, and Andrew Lo for advice and encouragement. We thank the editor (René Stulz) and an anonymous referee for constructive criticism. We thank Darrin King for attracting our attention to this area of research. Any errors are ours.




Figure 1. From Brealey and Myers—Weyerhaeuser. This is a reproduction of Figure 13-2 from Brealey and Myers (1991). Compare this to Figure 2 below. The horizontal axis shows return on Weyerhaeuser stock on a given day; the vertical axis shows return on the next day. Pluses (+) are used to symbolize the data. The data span the period January 2, 1986 to December 30, 1988. Out of 758 points, 6 lie outside the bounds of the plot.

Section I describes the compass rose; Section II explains it; Section III demonstrates its pervasiveness; Section IV extends it along different dimensions; and Section V concludes.

I. The Compass Rose

Earlier researchers have come close to discovering the compass rose. While discussing market efficiency, Brealey and Myers (1991, henceforth BM) show a scatter plot of simple daily returns on Weyerhaeuser stock (Figure 13-2, p. 293). The horizontal axis shows return on a given day, and the vertical axis shows return on the next day. They state that “it is obvious from a glance that there is very little pattern in these price movements (p. 291).” However, as shown below, they overlook a *significant* pattern.

BM’s picture is reproduced here as Figure 1. We demonstrate that a striking pattern is obscured by the (confounding) typographic choice of pluses (+) to symbolize the data points, and the choice of a short time series of observations. If the pluses are replaced with dots, and the time series is extended, then




Figure 2. Compass Rose—Weyerhaeuser. This figure is identical to Figure 1, except that we extend the time series, symbolize the data with dots instead of pluses, and adjust the scale. The compass rose pattern appears clearly. The horizontal axis shows return on Weyerhaeuser stock on a given day; the vertical axis shows return on the next day. The data span the period December 6, 1963 to December 31, 1993. Out of 7566 points, 1212 lie outside the bounds of the plot.

Figure 1 becomes Figure 2. In Figure 2, lines radiate from the origin; the rays shooting in the major directions of the compass are the thickest. We call this pattern a “compass rose,” after the navigators’ symbol of similar appearance.¹

Huang and Stoll mention an *intra*-daily version of the compass rose in a footnote, but they neither illustrate nor discuss it: “if (five-minute) price returns at t are plotted against lagged price returns at $t - 1$, one observes clusters of points that radiate from the zero return point at the center of the graph (Huang and Stoll (1994), p. 199).”

The compass rose deserves more than a footnote; Figure 2 is much more impressive than the previous quote suggests. We demonstrate that there is more to the compass rose than meets the eye. It is surrounded by unanswered questions that demand further research.

¹ To be exhaustive, let us mention that the actual pattern is subtly different from a compass rose. For example, North-North-East is almost, but not exactly, the bisector of North and North-East. Furthermore, the next level of subdivision is twice as close to North-North-East as to North. Details are left to the curious reader.

II. Explanation

The compass rose is generated because stock price changes are typically small-integer multiples of some tick size, be it an “official” tick size imposed by the exchange, or a larger “effective” tick size chosen by market participants. In almost every stock, the official and effective tick sizes are identical; in extremely high priced stocks, this need not be so (see Section III).

We now detail the generation of a compass rose in the hypothetical “Stock XYZ.” Let R_t and P_t be respectively the return and the closing price of Stock XYZ at date t . Let h be the tick size.

Suppose that daily price changes for Stock XYZ are small relative to the price level. We may then make the following approximation:

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

where the integer $n_t \equiv (P_t - P_{t-1})/h$ is the price change in ticks at date t .

It follows from Equation (1) that the data point (R_t, R_{t+1}) falls approximately on the ray joining the origin to the point with integer coordinates (n_t, n_{t+1}) .

Suppose in addition that a price change for Stock XYZ is typically a small number of ticks. Then the integer $n_t = (P_t - P_{t-1})/h$ is small in absolute value for most data points: 0; +1 or -1; and so on. In this case, only the major directions of the compass rose are strongly delineated: first North, East, South and West; then North-East, South-East, South-West and North-West; and so on.

If Stock XYZ's price were to remain relatively stable throughout the period, the resulting pattern would be a grid. On a 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. However, as the price P_t of Stock XYZ varies, data points are spread evenly on the ray. This “centrifugal smudging” transforms what would otherwise be a grid into the compass rose.

In summary, the compass rose appears clearly if Stock XYZ satisfies three conditions:

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

These three conditions are necessary and sufficient. We express them in subjective language only because the above statement “the compass rose appears clearly” is itself subjective. All stocks satisfy Conditions 1–3. The more strongly they satisfy them, the clearer the compass rose. Empirical evidence is given in Section III.




Figure 3. Compass Rose—2123 NYSE stocks. This plot shows one point only for each of 2123 NYSE stocks on a randomly selected pair of consecutive trading days; the structure is cross-sectionally coherent. The horizontal axis shows return on October 19, 1993; the vertical axis shows return on October 20, 1993. Out of 2123 points, 469 lie outside the bounds of the plot.

Observing a compass rose does not tell us anything beyond Conditions 1–3. Condition 2 reformulates the institutional rule that stock prices must move in discrete increments and in an “orderly” fashion. Conditions 1 and 3 characterize the intensity of stock price movements in the short run and in the long run, but remain neutral on timing and direction. It is hard to earn abnormal profits based on this type of information. In any case, all of this information—and much more—is contained in decades-old studies on the time-series properties of stock returns (marginal distributions, autocorrelation functions, and so on). The compass rose does not contribute any additional information to the prediction of stock price movements for the purpose of earning abnormal profits.²

III. Pervasiveness

Every stock we have investigated has a compass rose. A striking consequence is that if we select one pair of consecutive trading days and plot a single point for each New York Stock Exchange (NYSE) stock, the compass rose is

² However, as pointed out by the referee, we do not show that there are no other predictable structures present in Figure 2.




Figure 4. Compass Rose—Berkshire Hathaway. The pattern appears clearly even though the stock price moves by hundreds of eighths each day. The horizontal axis shows return on Berkshire Hathaway stock on a given day; the vertical axis shows return on the next day. The data span the period October 14, 1976 to December 31, 1993. Out of 4229 points, 405 lie outside the bounds of the plot.

preserved (see Figure 3). The compass rose is the norm for NYSE stocks and is not peculiar to BM's randomly chosen Weyerhaeuser.³

For high-priced stocks, the official NYSE tick size of $\frac{1}{8}$ is relatively small. Readers might thus expect that high-priced stocks do not have compass roses; in fact they do. For example, Berkshire Hathaway (trading around \$20,000 in late 1994) has a very clear compass rose (Figure 4).⁴ This counter-intuitive result appears because Berkshire Hathaway's stock price typically moves a few \$25 "ticks" each day (thus satisfying Condition 2 of the previous section). Although the official tick size of $\frac{1}{8}$ is irrelevant, the effective tick size of \$25 is not.

For almost every NYSE stock, the effective tick size equals the official one. Berkshire Hathaway is unusual in this respect. With its effective tick size much larger than the official one, Berkshire Hathaway stock is the ultimate in price clustering. This is not surprising, given that clustering increases in price

³ Do note however that for a few stocks (e.g., IBM, not shown here) an abundance of rays shooting out from the origin can make the compass rose slightly more hazy.

⁴ The compass rose is also clear in other stocks whose prices are high but not as extreme as Berkshire Hathaway's (e.g., CBS Inc., and Capital Cities ABC Inc., not shown here).




Figure 5. No Compass Rose—gold returns. This figure plots returns on gold. The returns are calculated using daily U.S. dollar London P.M. (i.e. afternoon) gold fixing prices. There is no compass rose. The horizontal axis shows return on a given day; the vertical axis shows return on the next day. The data span the period January 3, 1978 to November 17, 1994. Out of 4015 points, 339 lie outside the bounds of the plot.

level (Harris (1991)), and Berkshire Hathaway is an extremely high priced stock.

Although the compass rose is remarkably robust across individual stocks, it does not extend to portfolios of more than two stocks. In market indices, it is altogether absent. Portfolios violate Condition 2: their effective tick sizes are negligible with respect to their price changes. For example, the effective tick size of an equally-weighted portfolio is inversely proportional to the total number of stocks.

Robustness of the compass rose in individual stocks does not extend to all individual assets. For example, London gold fixing prices do not generate a compass rose (Figure 5). Although gold price levels do cluster on quarter dollars and dimes (Ball et al. (1985)), the degree of clustering is not sufficient to raise the effective tick size above the official one, a nickel. A nickel is negligible when compared to typical gold price changes, thus Condition 2 is violated.

We do not know whether the violation of Condition 2 in gold is attributable to the nature of gold (a precious metal) or to the nature of the London gold




Figure 6. Lagged returns—Weyerhaeuser. This differs from Figure 2 because we have increased the lag in Weyerhaeuser stock's daily returns from one day to 5 days (approximately one calendar week). The horizontal axis shows daily return on a given day; the vertical axis shows daily return 5 days later. The data span the period December 6, 1963 to December 31, 1993. Out of 7562 points, 1233 lie outside the bounds of the plot.

market (a Walrasian gold fixing).⁵ More generally, it is not clear how asset class (e.g. precious metal versus stock) and market structure (e.g. gold fixing versus NYSE auction) interact with the economic determinants of price discreteness. The first author is currently investigating this line of research.

The irrelevance of the official tick size for Berkshire Hathaway shows that the existence of an official tick size is not necessary for the compass rose. The gold fixing example shows that the existence of an official tick size is not 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.

Indeed, we propose this as a general rule of thumb to determine whether discreteness matters. For example, the absence of a compass rose in returns to gold suggests that, even though there is some clustering on price levels (Ball et al. (1985)), discreteness-induced biases in variance and serial covariance estimators may be negligible in this case. Further research is needed to verify this claim.

⁵ For details about the structure of the gold market in general, and the London gold fixing in particular, see O'Callaghan (1991). For a discussion of the Walrasian nature of the London gold fixing, see Ball et al. (1985).




Figure 7. Compounded returns—Weyerhaeuser. This figure plots overlapping five-day (approximately one calendar week) returns on Weyerhaeuser stock. It should be compared to Figure 2, which plots daily returns. The horizontal axis shows five-day return; the vertical axis shows the following five-day return. The data span the period December 6, 1963 to December 31, 1993. Out of 7558 points, 1518 lie outside the bounds of the plot.

IV. Related Plots

After extending the compass rose to other assets in the previous section, we now extend it along different dimensions. Some of these extensions suggest additional directions for future research.

The compass rose becomes more fuzzy as we increase the lag between the daily returns on the horizontal and vertical axes (see Figure 6 for a week's lag). The fuzziness is because Condition 1 is not satisfied; the price change from one period to the next is no longer negligible with respect to the price level. However, it should be noted that even at one year's lag (not illustrated), some traces of the compass rose structure remain in the plot.

The compass rose also fades away as we increase the compounding period of returns. It is visible in weekly returns (Figure 7), but indistinct in monthly returns, and completely absent in quarterly returns. This is due to the violation of Condition 2; the price change from one month to the next is often many ticks.

Microstructure effects are often associated with high-frequency data. However, we demonstrate in Figures 6 and 7 that the compass rose persists at relatively low frequency. It is essential that researchers be aware that even weekly data display complex microstructure effects.




Figure 8. Squared returns—Weyerhaeuser. We plot squared daily returns on Weyerhaeuser stock against themselves with one day's lag. The slight anti-symmetry in Figure 2 engenders pairs of rays here (see Footnote 1). The horizontal axis shows squared return on a given day; the vertical axis shows squared return on the next day. The data span the period December 6, 1963 to December 31, 1993. Out of 7566 points, 1626 lie outside the bounds of the plot.

In Figure 8, we plot squared stock returns against themselves with one day's lag; pairs of lines radiate from the origin. This pattern is a consequence of the compass rose. It may affect regressions of squared returns on themselves with one day's lag. Such regressions appear in the estimation of ARCH models. ARCH models have become very important in finance over the past decade (see Bollerslev et al. (1992) for a review). The impact of discreteness on the estimation of ARCH models has not yet been investigated and, judging from Figure 8, it is worthy of future research. It should receive at least the same attention as discreteness-induced biases in the estimation of variance of return (Gottlieb and Kalay (1985), Ball (1988), Harris (1990)).

Monte-Carlo simulations of stock prices as random walks truncated to the nearest eighth of a dollar (allowing for stock splits) also exhibit the compass rose. The simulated returns are so well behaved that if we plot them in a three-dimensional space (return versus return with one and two lags), a "star burst" is clearly visible (see Figure 9). In actual data, this three-dimensional version of the compass rose is visible, but it is not as crisp as in Figure 9.

Note that the "star burst" in Figure 9 is obtained by plotting what a chaos researcher would call "three-histories" of returns. That is, a 3-tuple $(R_t, R_{t+1},$




Figure 9. Star burst—simulated \$5 stock. The axes show return versus return with one and two lags. A three dimensional “star burst” pattern appears. The figure uses a simulated time series of 10,000 daily stock returns. Returns are drawn from the normal distribution with mean 0.0005 and standard deviation 0.012. Prices are truncated to the nearest eighth of a dollar. Stock splits (in the ratio 2:1) are enforced if the price doubles from its starting value of \$5. Out of 9998 points, 36 lie outside the bounds of the plot.

R_{t+2}) (see Barnett and Chen (1988)). Some chaos tests look for low-dimensional structure in n -histories (e.g., the “BDS” test in Brock et al. (1991)). However, the star burst in Figure 9 shows that such low-dimensional structure can appear spuriously. Discreteness may therefore bias the BDS test towards finding chaos. Further research is warranted.

V. Conclusion

For any given stock, price changes are typically small-integer multiples of the effective tick size. This discreteness in stock price changes restricts the possible values that the ratio of successive daily returns may assume. As a consequence, a striking geometrical pattern—a compass rose—emerges when daily stock returns are plotted against themselves with one day’s lag.

The compass rose is a robust phenomenon that extends to many other kinds of plots. In this case, we demonstrate clearly that structure and predictability do not necessarily go hand-in-hand, and that microstructure effects can be spectacular in daily, and even weekly, data.

We show that an “official” tick size (imposed by the exchange) is neither necessary nor sufficient for the compass rose. Instead, the key is to compare

the “effective” tick size (chosen by market participants) to typical price changes.

Directions for future research include: estimation of ARCH models; tests of chaos theory; and the interaction of asset class and market structure with the economic determinants of price discreteness.

REFERENCES

- Ball, Clifford A., 1988, Estimation bias induced by discrete security prices, *Journal of Finance* 43, 841–865.
- Ball, Clifford A., Walter N. Torous, and Adrian E. Tschoegl, 1985, The degree of price resolution: The case of the gold market, *Journal of Futures Markets* 5, 29–43.
- Barnett, William A., and Ping Chen, 1988, The aggregation-theoretic monetary aggregates are chaotic and have strange attractors: An econometric application of mathematical chaos, in William A. Barnett, Ernst R. Berndt, and Halbert White, eds.: *Dynamic Econometric Modeling, Proceedings of the Third International Symposium in Economic Theory and Econometrics* (Cambridge University Press, Cambridge, Cambridgeshire).
- Bollerslev, Tim, Ray Y. Chou, and Kenneth F. Kroner, 1992, ARCH modeling in finance: A review of the theory and empirical evidence. *Journal of Econometrics* 52, 5–59.
- Brealey, Richard A., and Stewart C. Myers, 1991, *Principles of Corporate Finance*, 4th ed. (McGraw-Hill, New York).
- Brock, William A., David A. Hsieh, and Blake LeBaron, 1991, *Nonlinear Dynamics, Chaos, and Instability: Statistical Theory and Economic Evidence* (MIT Press, Cambridge, Mass.).
- Edwards, Robert D., and John Magee, 1992, *Technical Analysis of Stock Trends* (New York Institute of Finance, Business Information and Publishing, New York).
- Gottlieb, Gary, and Avner Kalay, 1985, Implications of the discreteness of observed stock prices, *Journal of Finance* 40, 135–153.
- Harris, Lawrence, 1990, Estimation of stock price variances and serial covariances from discrete observations, *Journal of Financial and Quantitative Analysis* 25, 291–306.
- Harris, Lawrence E., 1991, Stock price clustering and discreteness, *Review of Financial Studies* 4, 389–415.
- Huang, Roger D., and Hans R. Stoll, 1994, Market microstructure and stock return predictions, *Review of Financial Studies* 7, 179–213.
- O’Callaghan, Gary, 1991, The structure and operation of the world gold market, Working paper, International Monetary Fund.