Death Of A Chartsman

"Past performance does not predict future results." Future archaeologists will puzzle over this incantation, which they will discover in many apparently unread – but nevertheless expensively printed – financial documents, and doubtless will conclude it was the misguided mantra from some bizarre sect. Otherwise, why would it be ignored so often in daily life, and for good reason: If past performance does not predict future results, then just what does?

Markets are driven by human behaviors common across cultures and instruments, such as eagerness to take small profits and distaste for small losses. It does not matter whether you are an American trading soybeans, or an Italian trading Eurolira or a Japanese trading silk cocoons, the behavior of traders should be the same in all markets. One corollary to this postulate is trading patterns are independent of the level of technology used in a market: It does not matter whether we are using clay tablets, rice paper, an abacus, or the Internet, the resulting footprints of a market are the same. We can read a cotton price chart from the Civil War, stock price charts from the 1920s, and grain price charts from the early 1970s on the same basis. A second corollary is fractal scalability; the comparability of five-minute bar charts to daily bar charts to weekly bar charts, for example.

Pattern recognition, for better or worse, is part of one's socialization process as a trader. We look at charts and see pennants, double tops, spikes, flags, and we learn to associate these patterns with an underlying economic story. Modification of a standard bar chart to include aspects of intraday structure yields Japanese candlesticks. As useful as candlesticks are, they are not a mutually exclusive and collectively exhaustive classification system, one capable of providing a unique label for all days.

The key to developing such a classifier is normalization of the key identifiers of a day's structure – its open, high, low, close, and midpoint – by locating them on a stochastic distribution of the day's range. Since there is no specific notation for this concept, a standard relational notation will be used. For example, if the range between a day's open and close – corresponding to the body on a candlestick – exceeds the mode of the day's stochastic distribution, it will be designated as "O>=C." In the table below, the first classification would be for a day where the open/close range, the open/midpoint range, and the midpoint/close range all exceeded the stochastic criterion.

O >= C		O < C	
O >= M	O < M	O > M	O <= M
(**1**) M >= C	(**3**) M < C	(**5**) M > C	(**7**) M <= C
(**2**) M < C	(**4**) M >=C	(**6**) M <=C	(**8**) M > C

Candlestick examples of each of the eight intraday structures are shown in the graph below. For those readers unfamiliar with candlesticks, the black bodies represent closes less than the open, the white bodies represent closes greater than the open, and the bars are extensions higher and lower than the maximum and minimum, respectively, of the open and close.

Intraday Structure Types

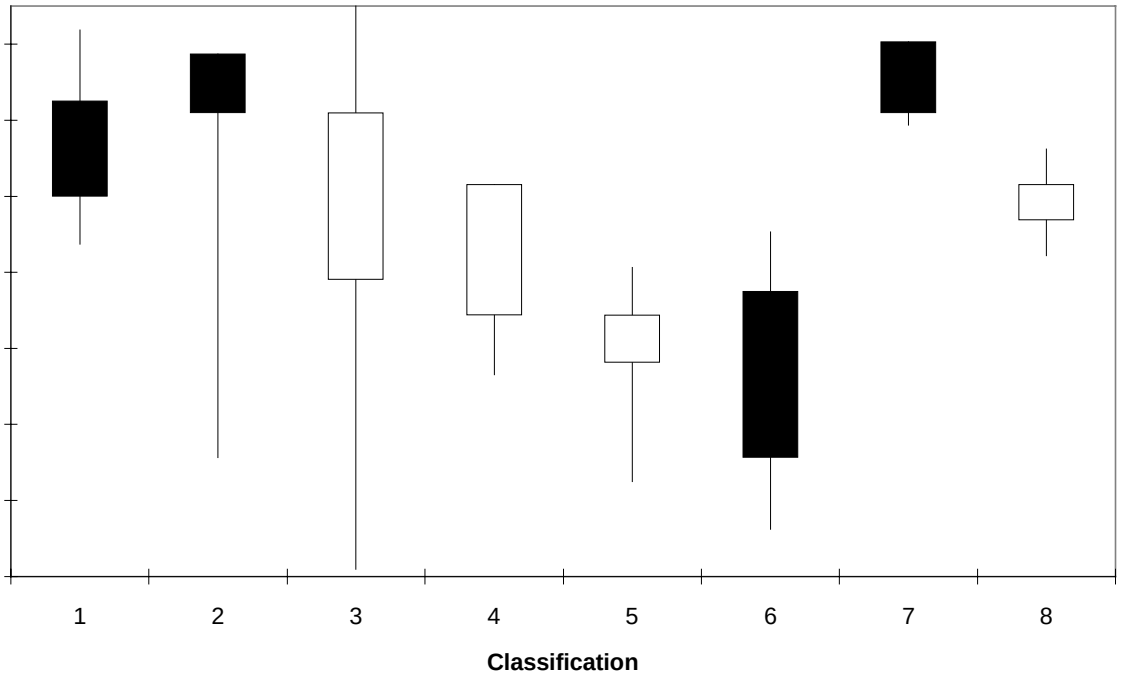

The utility of such an intraday classification lies in its information regarding the relative anxieties of buyers and sellers. For example, a day structure such as No. 2 demonstrates the willingness of buyers to step in and stop a price slide; the low of this day will be recognized by all as an important support point. A day structure such as No. 3, however, indicates the presence of overhead resistance.

Taken in isolation, these observations may or may not have any sort of predictive capability for interday price change or for the next day's intraday structure. The intraday structures need to be placed in the context of price trend; any trader recognizes instinctively the significant difference between a spike bottom occurring after several days of decline and one occurring as the reaction to a news-related development within an uptrend or a consolidation.

The interday path of price also can be classified on a mutually exclusive and collectively exhaustive basis. We can compare each day's open, high, low, close, and midpoint to those of the previous day, as shown in the table below.

1: Greater than the previous high
4: Less than the previous low
2: Greater than or equal to the midpoint
3: All else

Anatomy Of A Top
The chart below of June 1988 Treasury Bonds contains readily-visible patterns, including the double top, the diamond in the first top, and the ascending triangle in the second top.

A Double Top
June 1988 Treasury Bonds

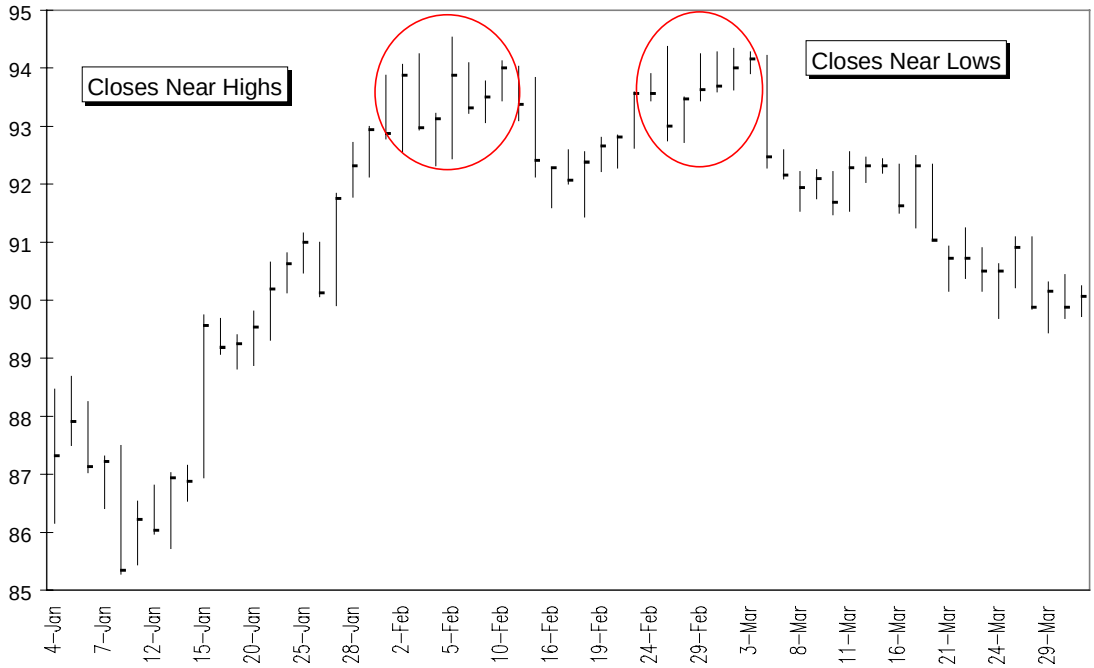

This rich information content contained in this graphical image is not in a form useful for universal analysis. First, we must classify this information using the combined intraday/interday patterns; the classifications of the two tops are presented below as an example. Classification allows data from different markets and time segments to be pooled; nearly 500 different pattern combined patterns exist. The pooled data base then can be "data-mined" to search for predictive trading patterns.

Integrated Classification

First Top		Second Top	
4-Jan-88	741343	19-Feb-88	821221
5-Jan-88	721222	22-Feb-88	821232
6-Jan-88	622344	23-Feb-88	321121
7-Jan-88	843443	24-Feb-88	411122
8-Jan-88	611444	25-Feb-88	441344
11-Jan-88	222333	26-Feb-88	333343
12-Jan-88	421232	29-Feb-88	721121
13-Jan-88	331341	1-Mar-88	431233
14-Jan-88	821222		
15-Jan-88	221121		

Data mining is predicated upon past performance predicting future results. We can have fifty historical observations of a certain day structure always leading to a higher close on the next day, but there is nothing deterministic in this observation: Any unforeseen event on the following day can lead to a significantly lower close. You will now have a condition that leads to a higher close on 98% of subsequent days, and you must then decide whether this is good enough for you to base a trading decision. The tradeoff will be between accepting a larger number of trades versus accepting a lower success rate. Obviously, these decisions are easier with a larger number of

observations of a particular pattern than with just a few. A statistical test for whether the pattern's apparent-predictive powers are significantly different from zero is given by the formula below. Place the mean of the observations in the numerator, and divide by its standard deviation multiplied by the square root of the number of observations. Then look up the location of the Z value in a table of the Student's, or "T," distribution. This test is available using the TDIST function in Microsoft Excel.

$$Z = \frac{\overline{X} - \mu_o}{\sigma / \sqrt{N}}$$

Changes in how markets operate physically are threatening the comparability of past data. The very concepts of opens and closes are being eroded by electronic trading systems that either never close or operate in such adjacent time frames to pit sessions as to render them virtual extensions of the pit. The proliferation of electronic trading already has relegated the "gap" on a chart, once one of the more reliable signposts for technicians, to the endangered-species list: How much can any price gap when its electronic session opens up thirty minutes after the pit closes?

Since the trading world clearly is becoming more electronic – and at a faster pace than anyone imagined – the traditional methods of organizing price histories will become less useful. Two broad alternatives have been presented. The first is to organize price over equal volume, as opposed to equal time, segments. This alternative has particular merit in foreign exchange markets, which are the most continuous on a global basis. However, all volume is not created equally: The volume occurring in the active overlap between mid-day in London and early morning in New York may be more significant for determining market direction than a comparable volume total made during late North American hours, for example. Moreover, volume is affected by externalities such as holidays, time of year, news events, and arbitrage strategies, while time moves along in a known direction at a constant pace. Finally, only the foreign exchange and, by extension, short-term interest rate, markets are suited for this sort of classification.

A second proposed method of data organization is demarcation by type of price action itself. Prices do not move at a uniform pace, they surge quickly toward newly-perceived zones of underlying economic value and then oscillate slowly around fairly-priced economic value. Several problems emerge from such a proposal. First, what are the boundaries between a consolidation and a price distribution; when does each begin and end? Second, how do we distinguish between different types of price movements and consolidations? Finally, do we need to make a distinction between similar type of price movements and the underlying time and volume conditions in which they occurred?

Even with the problems of price-demarcation, it is close enough to the day-dependent integrated classification scheme discussed above to suggest that it could form a basis for pattern recognition when trading days disappear altogether: The integrated classification system defines whether a period has been in a consolidative or distributive structure, and provides a relational measure of its extent. Time-dependence must remain a constant feature of any classification scheme because of its constancy and ease of sub-division; the day may become a meaningless division, but the five-minute bar, for example, will never go out of style.

The addition of other dimensions, such as volume or trade count, time of day, day-of-week, or volatility to a combined intraperiod/interperiod classification will result in a much higher information density than heretofore available. This should result in significantly more sophisticated trading systems for those able to develop them first.