

A Statistical Approach to Technical Indicators

by
Michael Gutmann

Robust Technical Indicators

The practitioner of technical analysis views markets based solely on price and volume data. Consistent with this approach, the technical analyst seeks tools that are independent of any particular market or time frame. The goal is to trade any market and time frame relying on the best set of technical indicators available. For example, a reliable trending indicator plus volume study should be applicable whether the asset traded is S&P 500 futures or INTC (Intel Corp.) shares. Technical indicators that are market-independent are viewed as robust and sought after by the analyst.

Some of the best known technical indicators require asset-specific settings. For example, the well-known and widely used MACD indicator outputs the difference between two exponential moving averages, and the moving averages are based on the price of the underlying asset. In this case, using similar 5-minute bar charts, an overbought/oversold (OB/OS) MACD level might be 8 ticks (2 points) for the S&P 500 e-mini ES contract and 5 ticks (\$0.05) for INTC shares. Using the MACD across a variety of assets and time-frames requires the trader figure OB/OS levels for each application, which can be time-consuming and error-prone.

This article describes a simple method of making technical indicators robust. It illustrates a method of converting indicator outputs from asset-specific values to statistical measures of price extension and compression. This results in technical indicators that can be used across markets without modification.

Statistical Indicators

The most common method of estimating the spread, or dispersion, of a data set is standard deviation. The data's mean, or average, referred to with the Greek letter μ ("mu"), is first calculated. The data's standard deviation is then calculated as the average distance of the data set from its mean. The standard deviation, referred to with the Greek letter σ ("sigma"), is easy to work with because it takes values that are the same units as the original, underlying data.

The science of statistics has determined that for many naturally occurring populations (population height, weight, test score, etc.), data is "normally" distributed about its mean. This is the well-known bell curve of population distribution. Interestingly, bell curves are completely defined by their mean and standard deviation. This allows one to say that a normal distribution has approximately 70% of its data contained within one standard deviation of its mean.

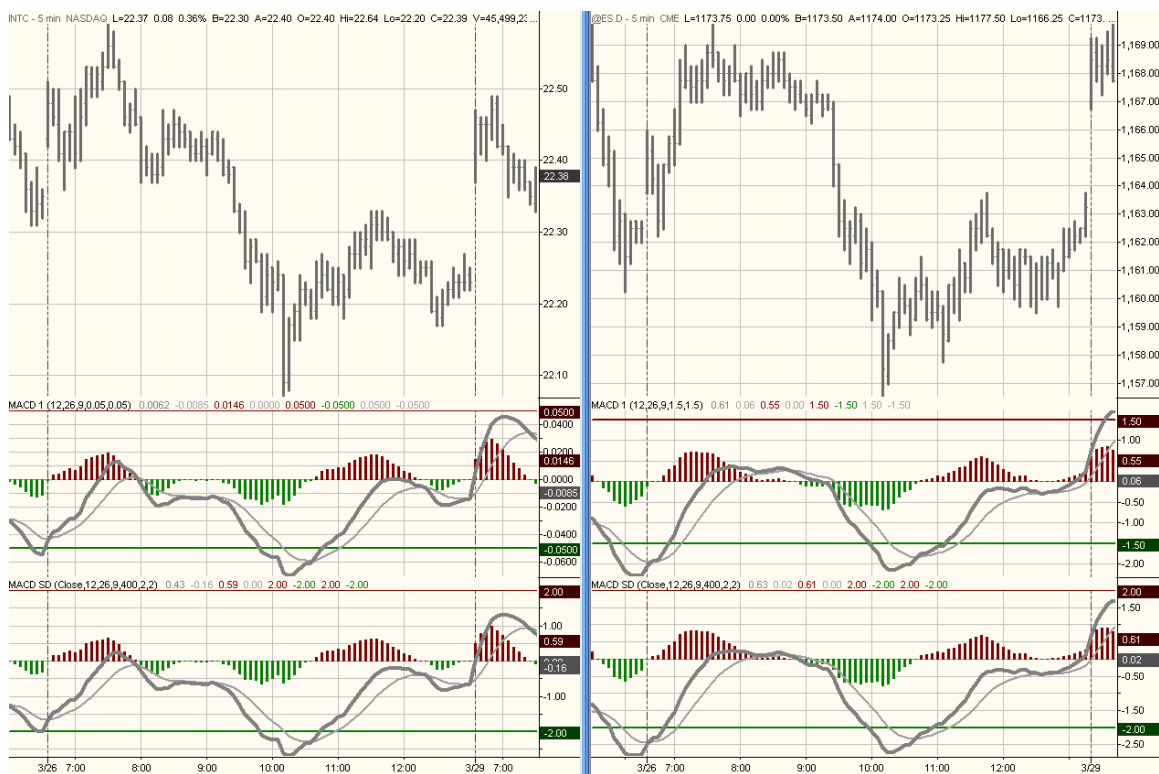
and 95% within two standard deviations, regardless of the values computed for the data's mean and standard deviation. For example, if test scores range from 0 to 100, with a mean of 75 and a standard deviation of 10, then we can predict that 70% of the scores range from 65 to 85 ($75-10$ to $75+10$).

It is questionable whether financial markets can be accurately modeled using a normal distribution. Markets are said to have "fat tails," meaning there may be an unexpected amount of data far from the mean, say outside of the two, or even three, standard deviation levels. Fat tails model market panics and over-exuberance, and there is a large body of financial analysis that uses different assumptions of dispersion than the normal distribution. None-the-less, normal distribution works as a good first approximation that can be used by traders. For example, the Market Profile Value Area is defined as encompassing 70% of price movement beginning from a mode (most common) price. The 70% value represents a one standard deviation variation in price.

Making use of standard deviation with market price and, in particular, bar charts, is straightforward. The standard deviation of recent price history is first calculated using some number of previous bar prices. Picking a history length is similar to selecting the length for a moving average indicator. Then a standard deviation relative position of recent price can be determined. Specifically, the current price, x , is said to be at " $(x - \mu)/\sigma$ standard deviations"; that is, the current price, x , is some number of standard deviations displaced from the mean.

The value $(x - \mu)/\sigma$ may at first seem obscure. But consider the meaning of any fraction or ratio, for example, the fraction one-third. An interpretation of the fraction is, How many 3s are there in 1? In the same way, the ratio $(x - \mu)/\sigma$ asks, How many standard deviations are there in $x - \mu$, the distance of the current price from market mean price? The technical term for the ratio $(x - \mu)/\sigma$ is "z-score." By converting indicator output values to z-scores we are able to move technical indicators to a statistical basis.

Figure "Pennies and Points" shows a side-by-side comparison of 5-minute ES and INTC charts, 3/26/2010, with standard MACD and new "statistical" MACD (MACD SD) plots below the bar charts. Both MACDs input the popular (12, 26, 9) moving average length values. MACD SD takes one additional parameter, the price history length, which will be used to calculate a recent MACD mean and standard deviation. In this example, price history length was set to 400 bars. The standard MACD has used an OB/OS level of 1.5 points for the ES contract and \$0.05 for INTC. The "statistical" MACD SD uses a fixed $2\text{-}\sigma$ OB/OS level (2 standard deviations) for both ES and INTC. The MACD SD can be moved across assets without modification.



Pennies and Points: MACD and MACD SD. 5-minute INTC, ES; 3/26/2010.

Statistical Price Extension and Compression Indicator

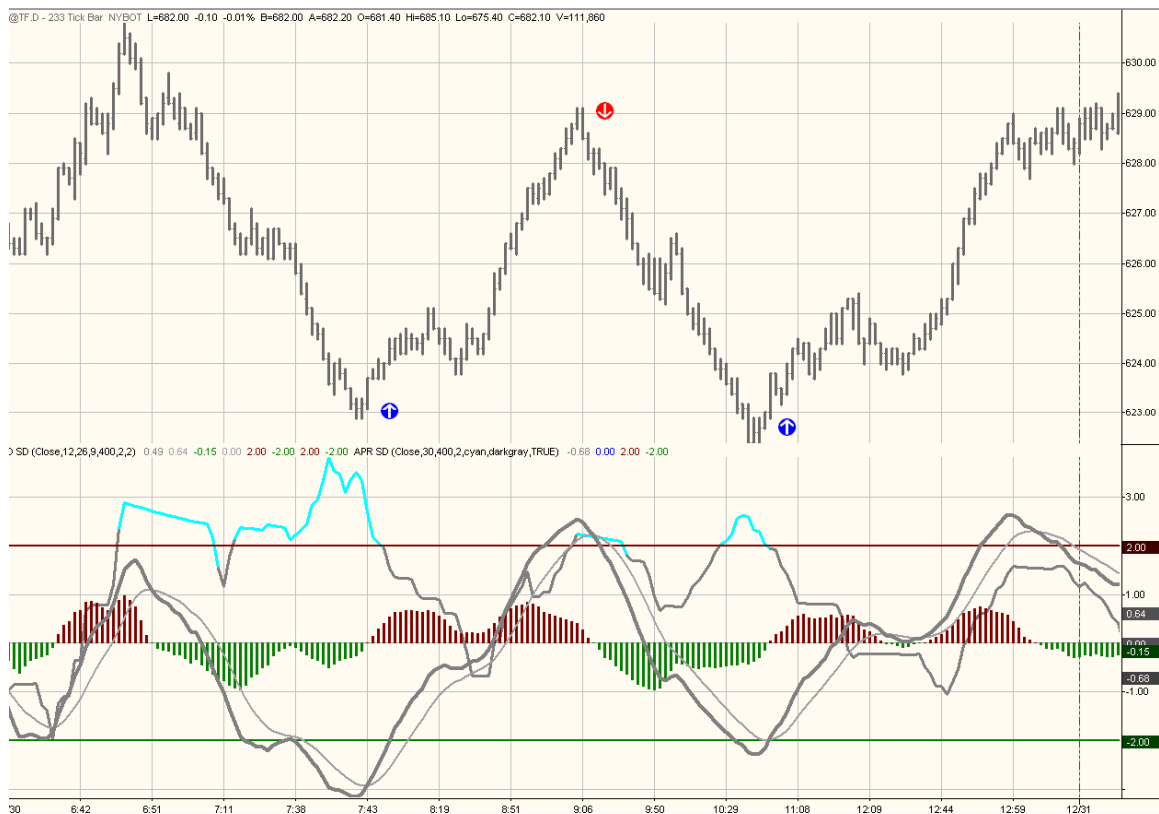
Converting technical indicators to statistical, z-score outputs can be helpful whenever absolute prices are referenced. Another example of this comes from a price extension/compression indicator. In this case, the trader is interested in gauging the extent of recent price action in a rotational market: Has a recent move been significant enough to consider fading the market expecting some reversal in price?

One method of measuring price extension/compression calculates the “highest high – lowest low” difference of the last n bars. If the difference is greater than some average price range, the assumption is that a significant price move has occurred. But when price movement is asset-specific, the measurement is complicated when moved from one chart to the next. Just as in the case of the MACD, a statistical approach is preferred.

The new indicator “Average Price Range Standard Deviation” (APR SD) has been written using the z-score techniques of standard deviation to measure price extension and compression. Figure “MACD Gets Some Help” shows the MACD SD overlaid with the APR SD on the Russell 2000 e-mini futures TF contract (NYBOT), 233-tick, 12/30/2009. APR SD produces a cyan-colored signal line

when recent price extension is above the two standard deviation level. Price compression is indicated when the APR SD indicator is below the negative two standard deviation level.

Using the MACD SD and APR SD in tandem offers the trader possible entry signals. While it must be said that this chart gives a particularly compelling illustration of the indicators, the theory, from a technical perspective, is consistent: Price extension may indicate a point of market rotation and MACD cross-over provides a reasonable entry point. Once again, there are no changes to the z-score based technical indicators when moving them to different assets.



MACD SD Gets Some Help – APR SD. 233-tick TF; 12/30/2009.

Statistical Implementation

The general implementation of a statistical indicator is:

1. Calculate technical indicator values. This is identical to the computation used for the original technical indicator.
2. Save the technical indicator values.
3. At the close of the last bar, calculate the mean and standard deviation from the saved technical indicator values; then calculate the z-score value for the current bar.
4. Plot the z-score value similarly to how the original indicator plotted its output.

There are two points to consider when carrying out a statistical implementation:

1. When saving technical indicator values, make use of an efficient data structure to retain value history. Data arrays, which require shifting previous entries to make room for the next entry are inefficient. A queue structure that provides efficient FIFO insertion is preferred.
2. Do not compute the statistical indicator on intra-bar updates. If lengthy price histories are used for mean and standard deviation calculation, then performing their computation intra-bar can be compute intensive. Most platforms support the option of restricting indicator computation to the close of a bar, and this is adequate to avoid unnecessary computation. The minor restriction does not eliminate the use of tick charts, and statistical indicators continue to work with tick charts. For example, see the TF 233-tick chart, above.

The technical analyst seeks to apply his or her preferred indicators across markets and time frames. Moving technical indicators to a statistically valued format retains the properties of the original indicator while making the new statistical indicators robust and portable. The statistical indicator often eliminates the need for time-consuming and error-prone modifications. This sort of simple statistical technique isn't rocket science but it is very useful.

And the use of statistical indicators can standardize the discussion of trade decisions and may add discipline to a trading program. Trade entry and exit decisions begin to be based on a well-defined measure of general price action, as opposed to varied, and perhaps inaccurate, asset specific values.

§§§

For 20 years Michael Gutmann was a software engineer and manager at Intel Corporation. He trades his system daily and has recently authored a new trading text: "The Very Latest E-Mini Trading: Using Market Anticipation to Trade Electronic Futures." Mike can be reached at mike@vgxcapital.com .