

Stationarity: The Best-Kept Secret in Trading Success

Brett N. Steenbarger, Ph.D.

www.brettsteenbarger.com

Note: The following article was written on 12/28/03 specifically for the Trading Psychology website. I believe the concepts in the article help to explain why 90% of the material written for traders and/or presented in workshops is invalid.

Every trader is familiar with what Victor Niederhoffer calls the “ever-changing cycles” within the market. Just as a pattern makes itself evident in the market, the pattern shifts to a new configuration. For example, the market may trade for a while within a range, offering nice buy and sell signals with a 14 period RSI. Then, abruptly, the market will break out and an overbought RSI will stay overbought or oversold for a prolonged period as the market makes a trending move.

It is because of these ever-changing cycles that traditional tools of technical analysis cannot be successfully applied in a purely mechanical fashion. The website Barchart.com has a nice feature where they track trading signals from such standard tools as moving averages. Over time, it is clear from their tracking that the signals do not perform better than random chance. For a while the signals will prove profitable, only to degrade once the cycles change.

It is ironic that traders spend considerable time researching better indicators and models while giving little thought to the time frame over which these trading tools might be valid. If, indeed, the market consists of ever-changing cycles, then any system or indicator is apt to degrade in its performance over time. In fact, if one waits for an indicator or system to develop a fine historical track record, the odds are good that their useful life are limited.

What can a trader do in the face of such uncertainty?

Stationarity

The statistician’s term for ever-changing cycles is stationarity. A number series is stationary if the process that generated the series has been constant. Clifford Sherry, in his excellent text The Mathematics of Technical Analysis explains, “A stationary time series is one in which the underlying *rules* that generate the time series do not change over time.” (p. 9).

My favorite example is the Las Vegas casino. Let’s say that you are playing blackjack and think that you have a superior card-counting strategy that will help you make money. By counting the number of picture cards vs. other cards that have been

dealt, you can assess the probability of drawing a picture card on subsequent hands, tilting odds in your favor.

Such a strategy will work as long as the number of decks employed by the dealer is constant. If, however, the dealer intermittently and secretly changes the number of decks in the shoe, the card counting strategy would be imperiled. If the gambler assumed that twelve cards worth 10 or higher were left in the deck because eight had been dealt, the assumption would be faulty if two decks instead of one were being used. By changing the rules for dealing cards, the dealer creates a distribution that is nonstationary.

Clifford Sherry notes the importance of nonstationarity for traders: “If you use these methods and techniques and find that your time series is nonstationary, it is probably best to stop and think carefully about your investment strategy. Nonstationarity implies that the underlying rules that ‘generated’ your time series change from time to time without warning. Therefore, you are dealing with maximum uncertainty about the potential outcome of your investment.” (p. 6). Sherry’s phrase “from time to time” is important. If, say, the card dealer changed the number of decks in the shoe after each and every hand, no card counting strategy would be possible. What makes counting viable is that the cycles are changing, but not constantly changing. A regime—a period in which the market follows a stable set of rules—can last for a while, allowing an alert trader to profit while it is in force.

What should be clear is that a skill essential to trading success is early identification of regime change: those occasions when the cycles are shifting and the distributions of price changes are significantly varying from their recent norms. If a card counter can quickly identify when the number of decks in the shoe have changed, he can avoid betting his old system and take the time to develop a new one. Similarly, once a trader notes that market behavior has shifted, he or she can stand back and identify the new rules that the market is following and position themselves for the new regime.

Amazingly, very few traders bother to look for stationarity and even fewer shift their trading strategies according to the characteristics of recent price change series. This includes technical analysts who employ the same indicators and indicator values across all markets and quantitative traders who fail to properly adjust their lookback periods when testing a relationship between predictors and price change. Would we expect the time series from 2000 to 2002 to provide an accurate database for gauging relationships in the 2003 market? Did the market from 1998 to 2000 provide useful guides over the subsequent two years?

Assuming stationarity when it is not there is one of the cardinal errors of trading. If you are trading a pattern that has been valid in the past and you don’t know if the current distribution of price changes match those from the past, you are flying blind. Successful trading requires that you identify the rules the market is following and base your strategy on those.

Fundamental Uncertainty in Trading

Let's go back to that last sentence. The trader knows that there are ever-changing cycles, but makes a fundamental assumption. That assumption is that the regime that is in place will not change over the next trading interval. The trader assumes that the market's rules will continue to be in force at least one more time. Without that assumption, the trader is either assuming randomness or is assuming regime change in the absence of concrete evidence of such. The only way we know if a regime has changed is by seeing an actual shift in the distribution of price changes. That means that there is a fundamental uncertainty in trading. The next trade may be the one in which the cycles shift. We cannot know for sure. Any trading strategy needs sound money management for this reason. Betting the house on a single trade—or during a single time frame—is courting ruin.

This has some interesting implications. For example, a well-researched trade that loses money may be an important source of trading information. If I have tested a historical period and found stationarity and then test a relationship between predictors and prospective price change over that period, my trade should have a high probability of success—*if* the market is remaining stationary. A losing streak with well-researched trades is often a sign that the markets are changing. Standing aside, waiting for evidence of the new regime, and remodeling the market over the more recent time frame corresponding to the new regime may allow the trader to learn from losses—and recoup them as well!

The fundamental uncertainty of trading is highest in daytrading the stock market—particularly index futures such as the SP/ES and ND/NQ. This is because markets are nonstationary on an intraday basis—almost without fail. Markets are most volatile early in the day's trading, retreat to lowest volatility in early afternoon, and then pick up volatility toward the close (only to plunge in volatility during Globex trading). It is rare indeed that the distribution of price changes from 09:30 – 11:30 AM ET will match those of 11:30 AM – 13:30 PM ET. Using the same indicators and indicator values in morning trading as in early afternoon and Globex sessions is a sure road to the poorhouse. Conversely, identifying regime change and valid relationships with each intraday regime shift requires a nimbleness—and an ability to control losses—that most traders lack.

Interestingly, markets exhibit greater stationarity from day to day and week to week than from hour to hour. That is one of the factors that has sped my transition from intraday trading to swing trading. But if stationarity is as important to trading as Sherry and I believe it to be, then it makes little sense to pigeonhole oneself as a short-term trader, a long-term trader, a daytrader, etc. One should trade the time frames that offer the greatest stationarity. If the market is stationary over a period of weeks and if you can clearly identify the rules the market is following over that period, it makes sense to trade those rules. Later, the market may exhibit stationarity over a shorter time frame, covering a series of days. The rules that capture that regime will provide the basis for trading.

Many times, we hear of the distinction between mechanical and discretionary trading. This is a false dichotomy, because both mechanical and discretionary trading often fail to take ever-changing cycles into account. The real alternative to mechanical trading is flexible trading that searches for regimes and the rules guiding regimes, exploiting these in a rule-governed manner.

From Theory to Practice

Analyzing the market for trades should begin with tests for stationarity. In my new swing trading system, I begin my analysis by identifying the longest swing period in which the markets are exhibiting a stationary series of price changes. (There may be more than one such stationary swing period, permitting diversification of trading by time frame, and—of course—there may be stationarity for certain instruments and not others, permitting diversification by trading vehicles.) My procedure for assessing stationarity is to divide the time series into halves and statistically test to see if the means and standard deviations for the halves are equivalent. For readers interested in the math involved, Sherry's book outlines a practical procedure for testing stationarity. The math is simple; I employ a quick-and-dirty t-test to the data and conduct the test entirely within Excel. What takes time is the repetitive testing of various lookback periods to find the proper window of stationarity.

Once I have that window, I then analyze the market qualitatively. I look at my indicators and observe how they have behaved during the stationary lookback period. The indicators that have consistently traced swing highs and lows over that period are the ones I will use to plan my next trade. I test signals yielded by the indicators (individually and in concert) over the lookback period to examine their entries, exits, and drawdowns. When I have a cadre of indicators that have performed well over the lookback period, I rely on them for my next trade.

But that's curve-fitting, you might protest. Isn't it dangerous to overfit the data with an optimized model?

My response is that optimization is only a problem when you fail to take stationarity into account. If you know you are trading within a stable regime, it makes sense to do your best to capture the rules the market is following over that period. My swing trading methodology might best be described as serial optimization: continually hunting for periods of stationary market behavior and trading optimized models derived from those periods.

Now here's the rub. When markets shift regimes, the window for the new regime is small. In testing the indicators that best follow the new, emerging rules, there aren't enough instances to properly conduct statistical tests. That is where historical tests become important. By identifying past periods of market history where the markets were following the same rules as today, we can see if the indicators and signals that work in the recent lookback period also worked back then. The crucial assumption is that markets that exhibit stationarity and equivalent means and standard deviations in price

changes are following the same rules—regardless of whether those markets were taking place in 2003, 1993, or 1983. If the strategy that we've optimized in the recent, stationary market window also produces profitable trading signals during past, similar regimes, we increase our confidence in the strategy and, indeed, can even test its signals statistically to ensure their departure from randomness.

Perhaps this is why we see so few traders incorporating stationarity into their analyses: It is time-consuming to assess market windows, operative trading rules, and test strategies for exploiting those rules. It is easier—and far more beguiling—to assume that a single system or indicator will produce consistent profits. More than one person has encouraged me to make my writing, research, and trading strategies less complex so that they can be more readily understood and accepted by the bulk of traders who attend seminars, buy trading books, and hire gurus for advice. One seminar organizer even fretted that I might be a threat to the self-esteem of traders, because the majority of traders lack the data and/or statistical background to conduct my kind of trading. I took that, of course, as quite a compliment.

Afterword

If you read the Trading Psychology Weblog with any frequency, you'll notice that many of the charts that I post have a common—and seemingly random—starting date. For instance, as I write this (12/28/03), many of my charts begin with 8/1/03. This is not an accident. The period from August through December represents one of those stationary windows from which we can extract useful swing trading strategies. By posting charts over stationary time frames in the market, the Weblog can assist you in identifying tradable market patterns.

Incorporating stationarity into your market thinking and trading opens the door to innovative trading approaches. For instance, within a longer stationary window (several months), you might identify a smaller window (the past several days) for a short-term trade. By nesting and aligning the short-term trade within a longer-term pattern, you can formulate some high probability trades. Beginning January, 2004, I will be posting real time swing trades to the Weblog that take advantage of rules derived over one or more stationary windows.

Yet another avenue for research is the use of very short-term nonstationarities to identify points of larger regime change. A while ago, when I was exclusively trading the SP on an intraday basis, I noticed how short-term shifts in the NYSE Composite TICK tended to occur at points of trend change in the market. The short-term nonstationarity was a marker for longer-term trend change. I believe the same occurs at all time frames. By monitoring shifts in short-term patterns and indicators, we may be able to hop aboard early phases of regime change.

Brett N. Steenbarger, Ph.D. is an Associate Professor of Psychiatry and Behavioral Sciences at SUNY Upstate Medical University in Syracuse, NY and a frequent trader of

the stock index futures markets. He is the author of The Psychology of Trading (Wiley, 2003) and coeditor of the forthcoming The Art and Science of the Brief Psychotherapies (American Psychiatric Press, 2004). Many of his articles on trading psychology and daily trading strategies are archived at his website, www.brettsteenbarger.com.