

# THE QUANTIFICATION PREDICAMENT

Timothy W. Hayes, CMT

*“This indicator has always produced huge profits! In fact, you would have doubled your money in just six months!”*

Such a claim could be a sales pitch. It could also be an analyst’s enthusiasm about some work just completed. But in either case, such claims appear to be meeting increasing skepticism, perhaps because enough have proven to be based more on fiction than quantifiable fact, perhaps because enough investors have been burned by indicators that have failed to pan out when put to real-time use, perhaps because the combination of ever-strengthening computing power and ever-increasing program complexity have made excessive optimization as easy, and dangerous, as ever.

In any case, the need to quantify accurately and thoroughly is greater than ever. Honest and reliable quantification methods, used in the correct way, are needed for increased research credibility. They are needed to impart objectivity. They are needed for effective analysis and for the sound backing of research findings. The alternative is the purely subjective approach that uses trendlines and chart patterns alone, making no attempt to quantify historical activity. But when the quantification process fails to deliver, instead producing misleading messages, the subjective approach is no worse an alternative – a misguided quantification effort can be worse than none at all. The predicament, then, is how to truly add value through quantification.

## THE CONCERNS

The major reason for quantifying results is to assess the reliability and value of a current or potential indicator, and the major reason we have indicators is to help us interpret the historical data. The more effective the interpretation of historical market activity, the more accurate the projection about a market’s future course. An indicator can be a useful source of input for developing a market outlook if quantitative methods back its reliability.

But for several reasons, quantification must be handled with care. The initial concern is the data used to develop an indicator. If it’s inaccurate, incomplete, or subject to revision, it can do more harm than good, issuing misleading messages about the market that’s under analysis. The data should be clean and contain as much history as possible. When it comes to data, more is better – the greater the data history, the more numerous the like occurrences, and the greater the number of market cycles under study.

This leads to the second quantification concern, and that’s sample size. The data may be extensive and clean, and the analysis may yield an indicator that foretold the market’s direction with 100% accuracy. But if, for example, the record was based on just three cases, the results would lack statistical significance and predictive value. In contrast, there would be fewer questions regarding the statistical validity of results based on more than 30 observations.

The third consideration is the benchmark, or the standard for comparison. The test of an indicator is not whether it would have produced a profit, but whether the profit would have been any better than a random approach, or no approach at all. Without a benchmark, “random walk” suspicions may haunt the results.<sup>1</sup>

The fourth general concern is the indicator’s robustness, or fitness – the consistency of the results of indicators with similar formulas. If, for example, the analysis would lead to an indicator that used a 30-week moving average to produce signals with an excellent hypothetical track record,

how different would the results be using moving averages of 28, 29, 31, or 32 weeks? If the answer was “dramatically worse,” then the indicator’s robustness would be thrown into question, raising the possibility that the historical result was an *exception* to the rule rather than a good *example* of the rule. An indicator can be considered “fit” if various alterations of the formula would produce similar results.

**Figure 1**  
Summary Results From Hypothetical Indicator Tests

These results contain an impressive-looking EXCEPTION to the rule ...

Number of Trades	Moving Average (Periods)	Buy Level	Sell Level	Accuracy Rate (%)	Gain/Annum (%)
40	70	100	110	50	11.2
39	71	99	111	50	11.3
37	72	98	112	65	15.1
37	73	97	113	52	10.1
36	74	96	114	50	9.8

These results would all be good EXAMPLES of the rule ...

50	20	15.6	8.6	55	11.8
49	21	15.8	8.4	56	12.0
48	22	16.0	8.2	56	12.1
47	23	16.2	8.0	57	12.1
46	24	16.4	7.8	56	12.0
<b>Buy-Hold Gain/Annum</b>					6.3

Moreover, the non-robust indicator may be a symptom of the fifth concern, and that’s the optimization process. In recent years, much has been written about the dangers of excessive curvefitting and over-optimization, often the result of unharnessed computing power. As analytical programs have become increasingly complex and able to crunch through an ever-expanding multitude of iterations, it has become easy to over-optimize. The risk is that armed with numerous variables to test with minuscule increments, a program may be able to pick out an impressive result that may in fact be attributable to little more than chance. The accuracy rate and gain per annum columns of Figure 1 compare results that include an impressive-looking indicator that stands in isolation (top) with indicators that look less impressive but have similar formulas (bottom). One could have far more confidence using an indicator from the latter group even though none of them could match the results using the impressive-looking indicator from the top group.

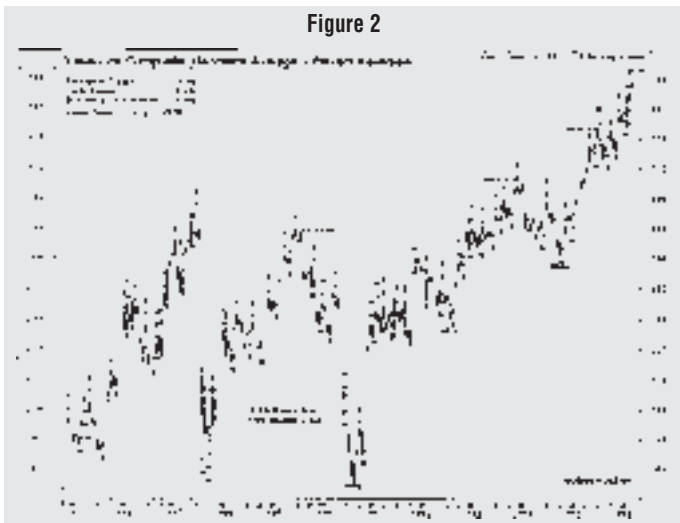
What follows from these five concerns is the final general concern of whether the indicator will hold up on a real-time basis. One approach is to build the indicator and then let it operate for a period of time as a real-time test. At the end of the test period, its effectiveness would be assessed. To increase the chances that it will hold up on a real-time basis, the alternatives include out-of-sample testing and blind simulation. An out-of-sample approach might, for example, require optimization over the first half of the date range and then a real-time simulation over the second half. The results from the two halves would then be compared. A blind-simulation approach might include optimization over one period followed by several tests of the indicator over different periods.

Whatever the approach, real-time results are likely to be less impres-

sive than results during an optimization period. The reality of any indicator developed through optimization is that, as history never repeats itself exactly, it is unlikely that any optimized indicator will do as well in the real-time future. The indicator's creator and user must decide how much deterioration can be lived with, which will help determine whether to keep the indicator or go back to the drawing board.

## TRADE-SIGNAL ANALYSIS

With the general concerns in mind, the various quantification methods can be put to use. The first, and perhaps most widely used, is the approach that relies on buy and sell signals, as shown in Figure 2.<sup>2</sup> When the indicator meets the condition that it deems to be bullish for the market in question, it flashes a buy signal, and that signal remains in effect until the indicator meets the condition that it deems to be bearish. A sell signal is then generated and remains in effect until the next buy signal. Since a buy signal is always followed by a sell signal, and since a sell signal is always followed by a buy signal, the approach lends itself to quantification as though the indicator was a trading system, with a long position assumed on a buy signal and closed out on a sell signal, at which point a short position would be held until the next buy signal.



The method's greatest benefit is that it clearly reveals the indicator's accuracy rate, a statistic that's appealing for its simplicity – all else being equal, an indicator that had generated hypothetical profits on 30 of 40 trades would be more appealing than an indicator that had produced hypothetical profits on 15 of 40 trades. Also, the simulated trading system can be used for comparing a number of other statistics, such as the hypothetical per annum return that would have been produced by using the indicator. The per annum return can then be compared to the gain per annum of the benchmark index.

But the method's greatest benefit may also be its biggest drawback. No single indicator should ever be used as a mechanical trading system – as stated earlier, indicators should instead be used as tools for interpreting market activity. Yet, the hypothetical and actual can be easily confused. Although the signal-based method specifies how a market has done between the periods from one signal to the next, they are not actual records of real-time trading performance. If they were, the results would have to account for the transaction costs per trade, with a negative effect on trading results. Figure 3 summarizes the indicator's hypothetical trade results before and after the inclusion of a quarter-percent transaction cost, illustrating the impact that transaction costs can have on results. The more numerous the signals, the greater the impact.

Also, as noted in the results, another concern is the maximum drawdown, or the maximum loss between any consecutive signals. But again, as long as it is clear that the indicator is for perspective and not for dictating precise trading actions, indicators with trading signals can provide useful input when determining good periods for entering and exiting the market in question.

## ZONE ANALYSIS

In contrast to indicators based on trading signals, indicators based on zone analysis leave little room for doubt about their purpose – they don't even have buy and sell signals. Rather, zone analysis recognizes black, white and one or more shades of gray. It quantifies the market's performance with the indicator in various zones, which can be given such labels as "bullish," "bearish" or "neutral" depending upon the market's per annum performance during all of the periods in each zone. Each period in a zone spans from the first time the indicator enters the zone to the next observation outside of the zone. Unlike the signal-based approach, the indicator can move from a bullish zone to a neutral zone and back to a bullish zone. An intervening move into a bearish zone is not required.

**Figure 3**  
Summary Results For Indicator In Figure 2 — No Transaction Costs

Value Line Geometric		\$ 574,104		1/24/72 – 5/30/96				
Last Signal "Sell"	Profit Current Trade	Number of Trades	Days Per Trade	Gain Per Trade	Batting Average	Model Gain Per Annum	Buy/Hold Gain Per Annum	\$10,000 Investment
5/07/96	-2.9%	240	37	1.9%	50%	18.1%	4.8%	\$574,104
Maximum Drawdown:				-4.68%				

**Summary Results For Indicator In Figure 2 — Including Transaction Costs Of A Quarter Percent Per Trade**

Value Line Geometric		\$173,271		1/24/72 – 5/30/96				
Last Signal "Sell"	Profit Current Trade	Number of Trades	Days Per Trade	Gain Per Trade	Batting Average	Model Gain Per Annum	Buy/Hold Gain Per Annum	\$10,000 Investment
5/07/96	-3.4%	240	37	1.4%	45%	12.4%	4.8%	\$173,271
Maximum Drawdown:				-4.68%				

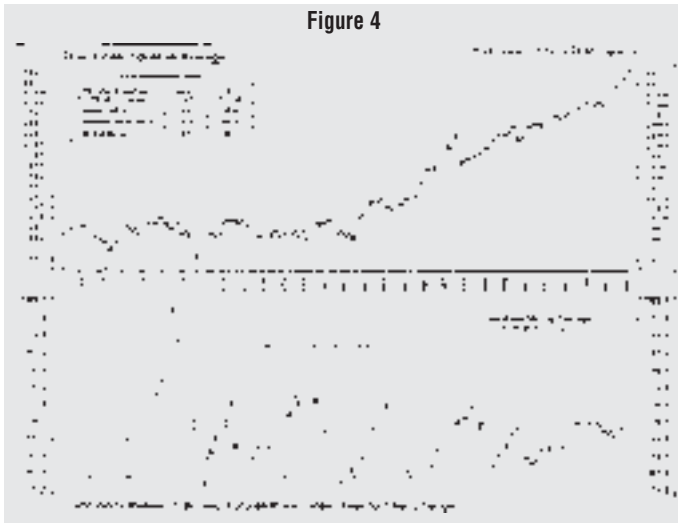
Zone analysis is therefore appealing for its ability to provide useful perspective without a simulated trading system. The results simply indicate how the market has done with the indicator in each zone. But this type of analysis has land mines of its own. In determining the appropriate levels, the most statistically-preferable approach would be to identify the levels that would keep the indicator in each zone for roughly an equal amount of time. In many cases, however, the greatest gains and losses will occur in extreme zones visited for a small percentage of time, which can be problematic for several reasons:

1. if the time spent in the zone is less than a year, the per annum gain can present an inflated picture of performance;
2. if the small amount of time meant that the indicator made only one sortie into the zone, or even a few, the lack of observations would lend suspicion to the indicator's future reliability;
3. the indicator's usefulness must be questioned if it's neutral for the vast majority of time.


A good compromise between optimal hypothetical returns and statistical relevance would be an indicator that spends about 30% of its time in the high and low zones, like the indicator in Figure 4. For an indicator with more than four years of data, that would ensure at least a year's worth of time in the high and low zones and would make a deficiency of observations less likely. In effect, the time-in-zone limit prevents excessive optimization by excluding zone-level possibilities would look the most im-

pressive based on per annum gain alone.

Another consideration is that in some cases, a closer examination of the zone performance reveals that the bullish-zone gains and bearish-zone losses occurred with the indicator moving in particular directions. In those cases, the bullish or bearish messages suggested by the per annum results would be misleading for a good portion of the time, as the market might actually have had a consistent tendency, for example, to fall after the indicator's first move into the bullish zone and to rise after its first move into the bearish zone.




It can therefore be useful to subdivide the zones into rising-in-zone and falling-in-zone, which can have the added benefit of making the information in the neutral zone more useful. This requires definitions for “rising” and “falling.” One way to define those terms is through the indicator’s rate of change. In Figure 5, which applies the approach to the primary stock market model used by Ned Davis Research, the indicator is “rising” in the zone if it’s higher than it was five weeks ago and “falling” if it’s lower. Again, the time spent in the zones and the number of cases are foremost concerns when using this approach.



Alternatively, “rising” and “falling” can be defined using percentage reversals from extremes, in effect using zones and trading signals to confirm one another. In Figure 6, for example, the CRB Index indicator is “rising” and on a sell signal once the indicator has risen from a trough whereas it’s “falling” and on a buy signal after the indicator has declined from a peak. Even though the reversal requirements resulted from optimi-

zation, the indicator includes a few poorly-timed signals and would be risky to use on its own. But the signals could be used to provide confirmation with the indicator in its bullish or bearish zone, in this case the same zones as those used in Figure 4. For example, in late 1972 and early 1973 the indicator would have been rising and in the upper zone, a confirmed bearish message. The indicator would then have peaked *and* started to lose upside momentum, generating a “falling” signal and losing the confirmation. That signal would not be confirmed until the indicator’s subsequent drop into its lower zone.



The chart’s box shows the negative hypothetical returns with the indicator on a sell signal while in the upper zone, and on a buy signal while in the lower zone. In contrast to the rate-of-change approach to subdividing zones, this method fails to address the market action with the indicator in the middle zone. But it does illustrate how zone analysis can be used in conjunction with trade-signal analysis to gauge the strength of an indicator’s message.

### SUBSEQUENT-PERFORMANCE ANALYSIS

In addition to using signals and zones, results can be quantified by gauging market performance over various periods following a specified condition. In contrast to the trade-signal and zone-based quantification methods, a system based on subsequent performance calculates market performance after different specified time periods have elapsed. Once the longest of the time periods passes, the quantification process becomes inactive, remaining dormant until the indicator generates a new signal. In contrast, the other two approaches are always active, calculating market performance with every data update.

The subsequent-performance approach is thus applicable to indicators that are more useful for providing indications about one side of a market, indicating market advances or market declines. And it’s especially useful for indicators with signals that are most effective for a limited amount of time, after which they lose their relevance. The results for a good buy-signal indicator are shown in Figure 7, which lists market performance over several periods following signals produced by a 1.91 ratio of the 10-day advance total to the 10-day decline total.

In its most basic form, the results might list performance over the next five trading days, 10 trading days, etc., summarizing those results with the average gain for each period. However, the results can be misleading if several other questions are not addressed. First of all, how is the average determined? If the mean and the median are close, as they are in Figure 7, then the mean is an acceptable measure. But if the mean is skewed in one

direction by one or a few extreme observations, then the median is usually preferable. In both cases, the more observations the better.

Secondly, what's the benchmark? While the zone approach uses relative performance to quantify results, trade-signal analysis includes a comparison of per annum gains with the buy-hold statistic. Likewise, the subsequent-performance approach can use an all-period gain statistic as a benchmark. In Figure 7, for instance, the average 10-day gain in the Dow Industrials has been 2% following a signal, nearly seven times the 0.3% mean gain for *all* 10-day periods. This indicates that the market has tended to perform better than normal following signals. That could not be said if the 10-day gain was 0.4% following signals.

**Figure 7**

**Percent Change Of Dow Industrials Following 1.91 Ratio Of 10-day Advances To 10-Day Declines**

Signal Date	10-Day A/D	Trading Days Later					
		5	10	22	63	126	252
06/23/47	1.96	-0.1	2.9	5.3	0.3	0.1	3.7
03/29/48	2.05	2.2	3.2	5.8	11.2	4.0	0.6
07/13/49	2.06	1.4	1.9	3.5	7.0	15.2	28.4
11/20/50	2.01	1.5	-1.7	-1.4	10.0	9.8	18.8
01/25/54	2.00	0.5	1.1	0.3	8.3	18.2	36.4
01/24/58	2.00	-0.1	-0.4	-3.1	0.6	10.3	31.4
07/10/62	1.98	-1.4	-2.0	0.9	0.0	14.0	21.5
11/07/62	1.91	2.4	3.5	4.8	10.3	17.3	21.1
01/13/67	1.94	1.4	1.1	2.6	2.9	5.6	6.9
08/31/70	1.91	1.1	-1.8	-0.5	3.9	15.5	17.9
12/03/70	1.95	1.5	1.7	3.6	11.1	14.1	5.0
12/08/71	1.98	1.0	3.5	6.2	10.6	10.4	20.2
01/08/75	1.98	2.8	2.7	12.0	20.9	37.2	41.4
01/06/76	2.05	2.5	6.6	8.3	12.7	11.3	10.9
08/23/82	2.02	0.2	2.6	3.9	14.6	22.6	34.0
10/13/82	2.03	1.9	-0.9	2.4	6.7	13.9	24.6
01/21/85	1.93	1.3	2.3	1.4	0.4	7.6	20.1
01/14/87	2.19	2.9	6.3	7.3	10.7	22.1	-5.4
02/04/91	1.96	4.7	5.8	6.9	6.1	7.8	16.7
01/06/92	1.99	-0.5	1.7	1.8	1.5	4.3	3.4
<b>Median</b>		1.4	2.1	3.6	7.7	12.6	19.4
<b>Mean</b>		3.1	2.0	3.6	7.5	13.1	17.9
<b>Mean – All Periods</b>		0.2	0.3	0.7	2.0	4.0	8.1
<b>% Cases Higher</b>		80	75	85	100	100	100
<b>% Cases Higher All Periods</b>		56	58	60	63	67	70

Signals based on 10-day total of NYSE advances over 10-day total of NYSE declines. Concept courtesy of Dan Sullivan, modified by Ned Davis Research.

A third question is how much risk has there been following a buy-signal system, or reward following a sell-signal system? Using a buy-signal system as an example, one way to address the question would be to list the percentage of cases in which the market was higher over the subsequent period, and to then compare that with the percentage of cases in which the market was higher over any period of the same length. Again using the 10-day span in Figure 7 as an example, the market has been higher after 75% of the signals, yet the market has been up in only 58% of all 10-day periods, supporting the significance of signals. Additional risk information could be provided by determining the average drawdown per signal – i.e., the mean maximum loss from high to low following signals. The mean for the 10-day period, for example, was a maximum loss of 0.7% per signal, sug-

gesting that at some point during the 10-day span, a decline of 0.7% could be considered normal. The opposite approaches could be used with sell-signal indicators, with the results reflecting the chances for the market to follow sell signals by rising, and to what extent.

Along with those questions, the potential for double-counting must be recognized. If, for example, a signal is generated in January and a second signal is generated in February, the four-month performance following the January signal would be the same as the three-month performance following the February signal. This raises the question of whether the three-month return reflects the impact of the first signal or the second one. Moreover, such signal clusters give heavier weight to particular periods of market performance, making the summary statistics more difficult to interpret. Problems related to double-counting can be reduced or eliminated by adding a time requirement. For the signals in Figure 7, for instance, the condition must be met for the first time in 50 days – if the ratio reaches 1.92, drops to 1.90, and then returns to 1.92 two days later, only the first day will have a signal. The time requirement eliminates the potential for double-counting in any of the periods of less than 50 days, though the longer periods still contain some overlap in this example.

**Figure 8**

**Performance Of Dow Industrials Following Initial Index Confirmation (Joint 52-week Highs For The First Time In A Year)**

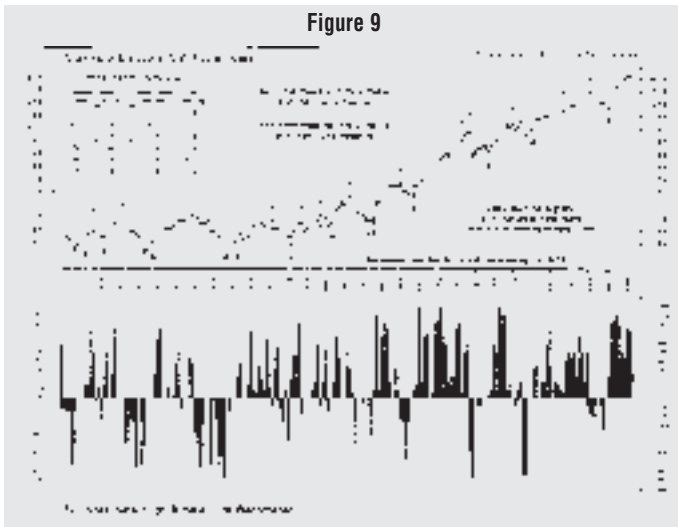
Confirming Index	26 Weeks Later			39 Weeks Later			52 Weeks Later			Latest Close	
	Cases	% Higher	Mean % Gain	% All Periods	% Higher	Mean % Gain	% All Periods	% Higher	Mean % Gain		% All Periods
New York Utilities	7	100	8.79	5.53	100	13.59	7.98	100	16.62	10.35	5/12/95
World Composite	6	100	8.47	5.85	80	9.74	8.88	100	12.91	11.78	9/15/95
Weekly New Highs	9	89	7.79	3.53	78	10.09	5.27	100	14.41	6.98	3/31/95
NYSE Weekly Volume	10	70	6.64	3.53	67	5.16	5.27	89	6.91	6.98	7/14/95
S&P 500 Composite	22	73	5.13	3.53	73	9.92	5.45	82	14.78	7.47	2/10/95
NYSE Composite	20	63	4.02	3.72	68	8.10	5.68	79	13.29	7.60	2/10/95
AMEX Composite	9	67	3.53	5.53	67	8.20	7.98	78	13.62	10.35	3/31/95
OTC Composite	9	56	3.77	5.53	67	7.73	7.98	78	12.38	10.35	3/17/95
Dow Transports	22	77	5.26	3.53	73	8.62	5.45	76	9.99	7.47	4/13/95
S&P High-Grade Index	12	67	5.46	3.53	75	9.73	5.27	75	10.77	6.98	2/17/95
S&P Industrials	12	58	1.66	3.53	58	4.34	5.27	75	10.15	6.98	2/10/95
NYSE Financials	11	45	0.59	3.53	55	4.86	5.24	73	10.06	6.92	4/07/95
Dow Utilities	23	70	6.00	3.27	65	7.63	5.05	73	9.18	6.95	5/05/95
Weekly A/D Line	12	58	2.44	3.53	67	5.30	5.27	73	7.31	6.98	4/13/95
S&P Low-Priced Index	11	55	1.26	3.53	40	2.88	5.27	70	7.31	6.98	7/14/95
Value Line Composite	14	50	1.35	3.53	50	3.74	5.27	69	6.26	6.98	4/13/95

Confirmation occurs, and a case identified, when the DJIA and the index in question both reach 52-week highs, the first such joint occurrence in at least a year. Table is sorted based on percentage of cases in which the index was higher over the subsequent 52-week periods (column shaded). “% All Periods” is the DJIA’s mean gain for all 26, 39, and 52 week periods starting with the beginning of the data series in question. Table updated through 4/04/96.


Another application of subsequent-performance analysis is shown in Figure 8, which is not prone to any double-counting. The signals require that three conditions are met, all for the first time in year – the Dow Industrials must reach its highest level in a year, another index must reach its highest level in a year, and the joint high must be the first in a year. The significance for the various indices can then be compared in conjunction with their benchmarks – i.e., the various all-period gains. Figure 9 uses 12 of those indices to show how subsequent performance analysis for both buy signals and sell signals can be used together in an indicator. For each time span, the chart’s box lists the market’s performance after buy signals, after sell signals, and for all periods.

## REVERSAL-PROBABILITY ANALYSIS

Finally, the subsequent performance approach is useful for assessing the chances of a market reversal. In Figure 10, the "signal" is the market's year-to-year change at the end of the year, with the signals (years) categorized by the amount of change – years with any amount of change, those with gains of more than 5%, etc. In this case, the subsequent-performance analysis is limited to the year after the various one-year gains. But the analysis takes an additional step in assessing the chances for a bull market peak within the one- and two-year periods after the years with market gains, or a bear market bottom within the one- and two-year periods after the years with market declines.



This analysis requires the use of tops and bottoms identified with objective criteria for bull and bear markets in the Dow Industrials. The reversal dates show that starting with 1900, there have been 30 bull market peaks and 30 bear market bottoms, with no more than a single peak and a single trough in any year. This means that for any given year until 1995, there was a 31% chance for the year to contain a bull market peak and a 31% chance for the year to contain a bear market bottom (30 years with reversals / 95 years).



Using this percentage as a benchmark, it can then be determined whether there's been a significant increase in the chances for a peak or trough in the year after a one-year gain or loss of at least a certain amount. The chart's boxes show the peak chances following up years and the trough chances following down years, dividing the number of cases by the number of peaks

or troughs. For example, prior to 1995, there had been 31 years with gains in excess of 15% starting with 1899. After those years, there was a 52% chance for a bull market peak in the subsequent year (16 following-years with peaks / 31 years with gains of more than 15%). The chances for a peak within two years increased to 74%, which can be compared to the benchmark chance for at least one peak in 61% of the two-year periods (since several two-year periods contained more than one top, this is not the exact double of the chances for a peak in any given year).

A major difference in this analysis is that in contrast to signals and zones, which depend upon the action of an indicator, this approach depends entirely on time. Each signal occurs after a fixed amount of time (one year), with the signals classified by what they show (a gain of more than 5%, etc.). Depending upon the classification, the risk of a peak or trough can then be assessed.

## CONCLUSION

Each one of these methods can help in the effort to assess a market's upside and downside potential, with the method selected having a lot to do with the nature of the indicator, the time frame, and the frequency of occurrences. The different analytical methods could be used to confirm one another, the confirmation building as the green lights appeared. An alternative would be a common denominator approach in which several of the approaches would be applied to an indicator using a common parameter (i.e., a buy signal at 100). Although the parameter would most likely be less than optimal for any of the individual methods, excessive optimization would be held in check. But whatever approaches are used, it needs to be stressed that each one of them has its own means of deceiving. By better understanding the potential pitfalls of each approach, indicator development can be enhanced, indicator attributes and drawbacks can be better assessed, and the indicator messages can be better interpreted.

The process of developing a market outlook must be based entirely on research, not sales. The goal of research is to determine *if* something works. The goal of sales is to show that it *does* work. Yet in market analysis, the lines can blur if the analyst decides how the market is supposed to perform, then selling himself on this view by focusing only on the evidence that supports it. What's worse is the potential to sell oneself on the value of an indicator by focusing only on those statistics that support one's view, regardless of their statistical validity. As shown by the various hazards associated with the methods described in this paper, such self-deception is not difficult to do.

Our goals should be objectivity, accuracy, and thoroughness. Using a sound research approach, we can determine the relative value of using any particular indicator in various ways. And we can assess the indicator's value and role relative to all the other indicators analyzed and quantified in a similar way. The indicator spectrum can then provide more useful input toward a research-based market view.

## FOOTNOTES

1. Reference to Burton Malkiel's *A Random Walk Down Wall Street*, which argues that stock prices move randomly and thus cannot be forecasted through technical means.
2. The charts that accompany this paper were produced with the Ned Davis Research computer program.

*Since winning the third Dow Award in 1996, Tim Hayes has expanded upon "The Quantification Predicament" in writing his first book, "The Research Driven Investor," published in November 2000 by McGraw-Hill. A Global Equity Strategist for Ned Davis Research, Tim and his team have developed numerous U.S. and global asset allocation indicators and models in recent years, while also developing global market and sector ranking systems and indicators based on 18 market sectors in 16 countries.*