Taking Stock: Here Are Simple Tools For Gauging the Health of the Market

ABREAST OF THE MARKET

By Douglas R. Sease
And Robert Steiner
Staff Reporters of The Wall Street Journal

After Friday's 30-point swoon in the Dow Jones Industrial Average, it's time for a stock market checkup.

The market ran hard and high for six weeks. And there isn't any reason to think there aren't more gains ahead this year and maybe next year, many analysts say.

But just as you don't keeping tooling down the highway in your car without occasionally checking the gas gauge, some investment experts suggest similar occasional checks of the stock market. If there's gas left, relax and enjoy the ride.

Experts can spout enough jargon to make the stock market seem complicated and esoteric to the average investor. Yet if you're not trying to play the daily and weekly shifts in the market – and most individual investors shouldn't be doing that – much of the sound and fury isn't necessary.

For example, Hugh Johnson, chief investment officer at First Albany Corp., says one reason investors can relax is that as bull markets go, this one is still young and vigorous. "The average bull market lasts 43 months and results in a 98.1% gain in stock prices," he says. So far, this bull market, which Mr. Johnson says began in October 1990, is only 15 months old and has provided only a 36% rise.

In using this history to gauge what the future may hold, Mr. Johnson conservatively eliminates the raging bull market of 1982-1990, which lifted the Dow industrials 286%. On that basis, he says, "If this bull market is just an average bull market, it would end in August of 1993 with the Dow industrials at 4045."

He cautions that long-term investors shouldn't try to time the inevitable temporary declines that occur in bull markets. "It's extraordinarily difficult to call short-term swings in the market," he says. "What you want to be able to do is to catch most of the rise in a bull market and avoid most of the losses in a bear market."

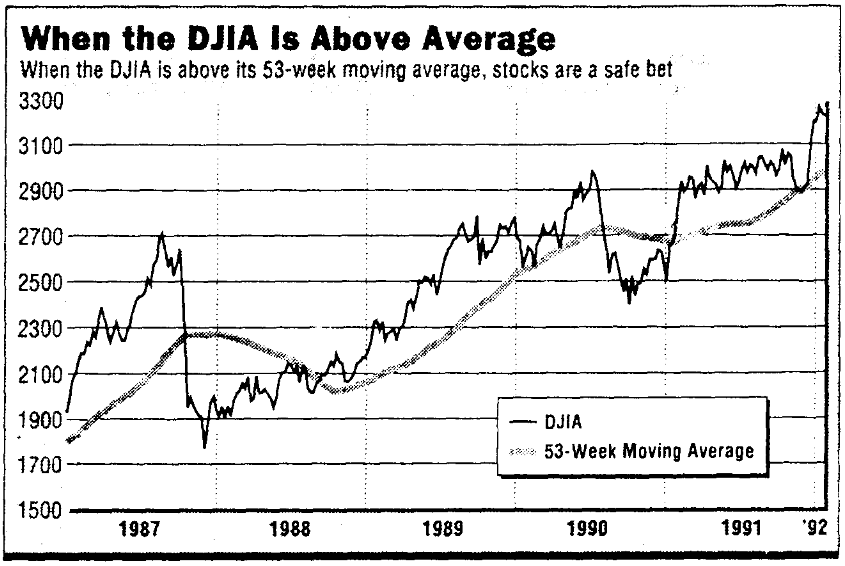

The tool that works best for that, he says, is a simple chart comparing where the market ends each week with the average performance of the market during the preceding 53 weeks. Mr. Johnson says technical analysts have determined that as long as the market finishes every Friday at a level above the 53-week moving average, owning stocks is fine. When it falls below the moving average, it's time to sell.

"That way, you enjoy the ride up, avoid most of the declines, and don't get chewed up by sales commissions on frequent trades," he explains.

Setting up your own chart takes a little initial investment of time and effort, but once it is constructed, maintaining it requires only a few minutes each week. Add the Friday closes of the index you choose for the past 53 weeks (for the Dow Jones industrials the total as of last Friday was 158,557) and divide by 53 to get the average (2992 for the Dow). Behold: Friday's decline to 3225.40 didn't come close to pushing the Dow below the 53-week moving average.

(To keep a moving average, you must each week drop the oldest Friday close and add the latest, again dividing by 53.)

But other analysts contend that it's even easier for typical investors to figure out whether the stock market is the place to be. "The statistics suggest one thing this week and another thing next week, creating a lot of noise," says Abby Cohen, a strategist at Goldman Sachs. "You should try to ignore the noise and just identify the key trends."

She thinks investors should avoid getting bogged down in market and economic minutiae, and instead obtain investment guidance by answering a simple set of questions. Such as: Is the economy going to be better six months from now than it is right now?

"The answer clearly is 'yes,'" she says. "We can't prove yet that things are getting better, but we know they're not getting worse." With the Federal Reserve clearly committed to sparking an economic recovery, she says, there's little reason to worry that another slump lies ahead.

Next question: Is inflation going to be a problem in the next several months? The answer is no. Not only is inflation slowing in the U.S.; it's slowing abroad as well, she says. Raw materials are in abundant supply, there's plenty of idle factory capacity and lots of people are out of work. Indeed, Ms. Cohen suggests that there may be

Please Turn to Page C2, Column 3

Continued From Page C1

more factory capacity and more unemployment, or at least "underemployment," than most analysts think.

"As a nation that imports and exports we should look at factory capacity beyond our own borders," she says, "and if we do we would find a great deal of slack tha[n] government statistics don't pick up."

On the employment front, she adds "we have a lot of part-time workers working fewer hours than they would like to and a lot of discouraged workers who aren't even looking for work. It will take an extended period of economic growth before the labor markets start getting tight."

Finally, ask if there are any attractive investment alternatives to the stock market. "Here it's a question of alternatives that look better in an environment of moderate economic growth," she says. "In that kind of environment what we have found in the past is that stocks are the best place to be."

But it mightn't be all that easy after all. Gail Dudack, chief market strategist for S.G. Warburg, worries that the stock market won't be easy to evaluate until the 1992 presidential election is settled.

"Politics plays a bigger role right now than normal," she says. "The big hurdle between having faith that inflation is under control and fearing that it's out of control has to do with fiscal policy. But we don't know what the Democratic package will be."

She also warns that political uncertainty means investors can't be sure that another cut in interest rates would boost stock prices they way one did in late December.

"Another easing could be seen as infla-

Fund Watch

The Official Guide to High-Performance Mutual Funds

Norman G. Fosback, Editor

Glen King Parker, Publisher

SPECIAL BONUS REPORTS

Part II – INVESTMENT STRATEGIES

Moving Average Trading Systems

Moving-average systems are a popular and easy way to trade mutual funds. They represent a purely "technical" approach to timing, because all trading decisions are the automatic result of price action, with no consideration given to any fundamental factor.

Moving-average systems are designed to keep investors in the market during most price uptrends, while offering some prospect that they will be out of the market during downtrends. All such systems are based upon, and dependent on, the tendency of the market, and the prices of individual funds, to move in trends.

What is a Moving Average?

The heart of these systems, of course, is the moving average itself, which is simply a statistical tool for smoothing a series of price values to eliminate minor fluctuations. For example, a ten-week arithmetic moving average – the dotted line in our chart of *Fidelity Magellan* – consists of successive averages of that fund's ten most recent week-ending prices. Just add up the ten latest Friday prices and divide the total by ten. With each subsequent week, the newest value is incorporated into the average, and the value of eleven weeks previous is dropped, so that the ten most recent weekly prices are always used to calculate the moving average.

Also pictured in Fidelity Magellan's chart is a 40-week moving average, denoted by a dashed line. It is useful to compare the 10-week and 40-week moving averages because they reveal how moving averages perform and how timing systems that are based on moving averages work in practice.

Note that the 10-week moving average is nearly always much closer to the actual price series than is the 40-week average. This is because the former consists only of the ten most recent fund prices whereas the 40-week moving average is also a function of much older prices that may be far removed from the current fund price.

Because they are more sensitive to price, shorter-term

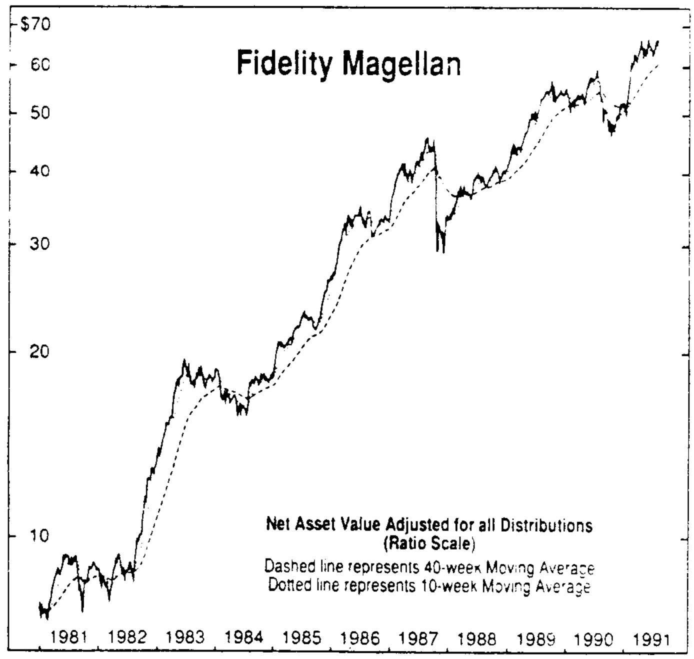

moving averages such as a 10-week, have more twists and turns than longer-term averages. For example, if an investor uses a timing system such that he buys a fund when its current price moves up through a moving average, and sells when the price falls below the moving average, he will obtain many more buy and sell signals with a short-term moving average simply because that moving average is closer to the price series. Longer-term moving averages produce fewer buy and sell signals.

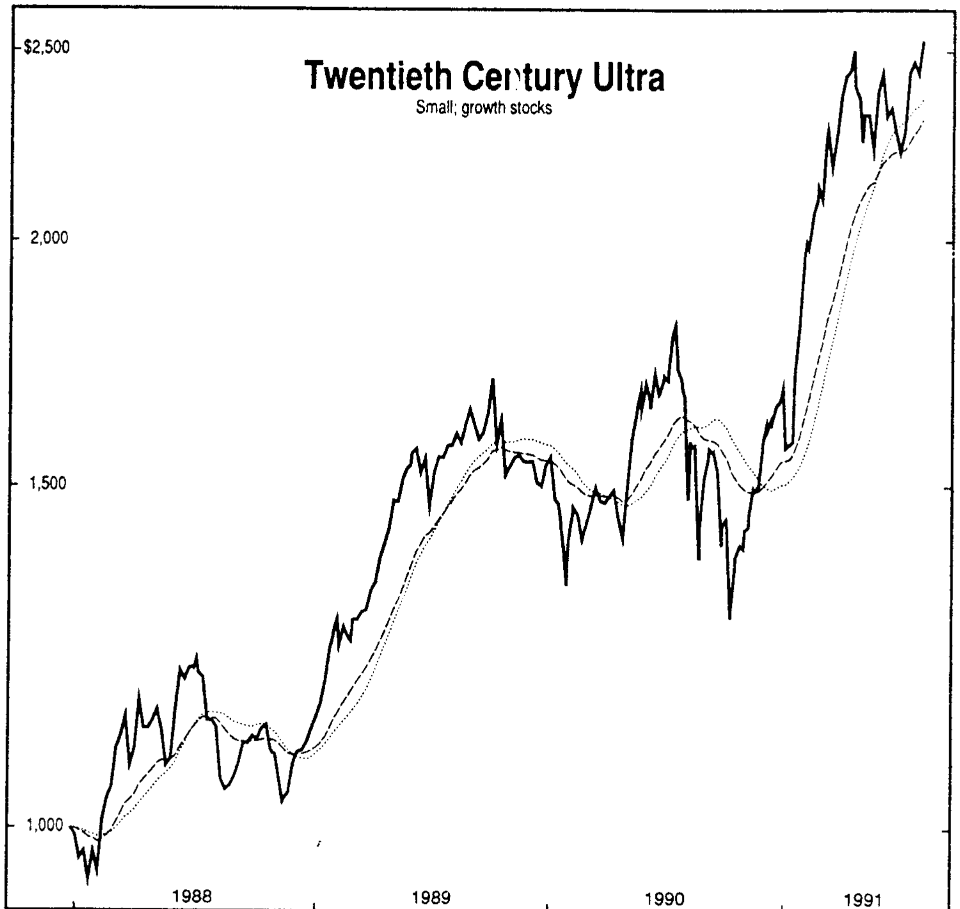

Telephone	Tele.Switching	Sales Load	Red. Fee	Exp.Ratio	Min. Inv.	Str.Rnk.
800-345-2021	1@mo; free	None	None	1.00%+	$0	99

Simple arithmetic moving average are, however, subject to criticism on two counts: First, they assign equal weight to each of the base observations. In a ten-week average, for example, each of the ten values is counted once and has a one-tenth importance (or weight) in the average. But, it would seem logical that more recent observations may, by their very nature, be more important and should receive more weight. Second, as a simple arithmetic average moves through time, its point to point fluctuations are strictly dependent upon two numbers, the one being added and the one being dropped. If the new number is greater than the old one, the average of the values will increase and so, too, will the moving average. If the current number is less than the old one being dropped, the average will decrease. Therefore, even though the *level* of an arithmetic moving average is dependent upon all of the prices being averaged, *fluctuations* in the moving average are dependent solely upon two numbers, the one being added and the one being dropped, and the last of these is older and of questionable relevance.

An alternative type of moving average that overcomes both of these objections, is an *exponential* moving average. An exponential system is based upon the assignment of a fixed weight (say, 10%) to the current price, and all of the remaining weight (in this case, 90%) to the previous value of the moving average itself. The proportional weight assigned to the most recent observation is frequently called a "smoothing constant."

To determine the exponential smoothing constant roughly proportionate to a simple moving average of a given length, use the following easy formula: Divide "2" by one more than the number of terms in the simple moving average you wish to duplicate. For example, to find a smoothing constant to construct an exponential moving average equivalent to a ten-week simple moving average, divide 2 by 11. The result, 0.18, is the smoothing constant. The following table illustrates how this smoothing constant is used to construct an exponential average of a price series.

Example of a Ten Day Exponential Moving Average

Week #	Price	Method of Calculation	Exponential Moving Average
1	$ 10	(to start)	$ 10.00
2	11	(0.18 x 11.00 + 0.82 x 10.00)	10.18
3	10	(0.18 x 10.00 + 0.82 x 10.18)	10.15
4	11	(0.18 x 11.00 + 0.82 x 10.15)	10.30
5	13	(0.18 x 13.00 + 0.82 x 10.30)	10.79
6	14	(0.18 x 14.00 + 0.82 x 10.79)	11.38
7	13	(0.18 x 13.00 + 0.82 x 11.38)	11.67
8	12	(0.18 x 12.00 + 0.82 x 11.67)	11.73
9	13	(0.18 x 13.00 + 0.82 x 11.73)	11.96
10	15	(0.18 x 15.00 + 0.82 x 11.96)	12.51

After arbitrarily extablishing the moving average as equal to the beginning price, the moving average is updated by multiplying the newest price by 0.18 (the smoothing constant) and adding that to the product derived from multiplying the previous exponential moving average value by 0.82 (1.0 minus the smoothing constant). After ten weeks, the exponential moving average value is $12.51. That value is greater than the simple ten week arithmetic average of $12.20 because proportionately greater weight has been assigned to the more recent, and in this case higher, prices. In an exponential moving average, the weight effectively assigned to any given historical value declines as it becomes older, but it is worth noting that *every* historical price always has *some* weight in an exponential moving average. That weight never declines completely to zero but merely trends closer and closer to zero, while never quite reaching it.

The chart on the facing page of *20th Century Ultra Fund* displays its price over a recent three-year period along with two 20-week moving averages – a 20-week arithmetic moving average and a 20-week exponential moving average.

We can quickly see that the two different types of moving averages seldom have the same value. A close inspection reveals that the exponential moving average is usually somewhat more sensitive to price; that is, it tends to move more quickly toward the price series itself. That is because the exponential moving average gives more weight to recent prices while the arithmetic moving average gives an equal amount of weight to all of the different prices (in this case, 20 weeks of prices) contained within it. As a result, an investor using price penetrations of an exponential moving average for buy and sell signals will usually get those signals earlier than an investor using the same length arithmetic moving average. Whether this actually proves to be more or less profitable, of course, depends upon the actual prices at which the buy and sell signals are derived. In all of our studies, we have found that the profitability of the arithmetic and exponential moving averages is relatively similar, with a *slight* edge frequency going to the exponential system.

Although exponential moving averages are typically not as well known to most investors, they are also actually easier to update in real time. Given a marginal superiority in the profitability of trading strategies, they are therefore generally to be preferred.

Arithmetic Moving-Average Trading Strategies (6/30/82 – 6/30/91)

Moving-Average Strategy	Profit on $10,000 Initial Investment	Profit Compared to $28,170 Profit of Buy & Hold	% of Time Out of Funds and in T-bills	Number of Round-Trip Trades
5 Week MA Strategy	$27,040	4% less profit	37% less risk	58
10 Week MA Strategy	30,400	8% more profit	33% less risk	36
15 Week MA Strategy	31,600	12% more profit	31% less risk	26
20 Week MA Strategy	30,770	9% more profit	29% less risk	21
25 Week MA Strategy	29,420	4% more profit	28% less risk	17
30 Week MA Strategy	28,300	1% more profit	27% less risk	15
35 Week MA Strategy	27,880	1% less profit	26% less risk	14
40 Week MA Strategy	27,190	3% less profit	26% less risk	13
45 Week MA Strategy	26,210	7% less profit	25% less risk	12
50 Week MA Strategy	25,490	10% less profit	24% less risk	11

How Profitable Are Moving Average Systems?

The basic theory underlying virtually all moving-average systems is that when the price of a fund is above the moving average, the fund is in an uptrend and should be owned. When the price is below the moving average, the fund is downtrending and should be avoided.

We have created an elaborate computer-based simulation model to test moving-average trading systems. Our initial findings will no doubt be somewhat surprising to technicians, and to the Wall Street establishment. In general, we have discovered that over the last decade, moving-average trading systems have a decidedly mixed track record in terms of producing profit relative to a simple "buy-and-hold" strategy. However, they are invariably excellent at reducing risk.

Our study is based on the performances of all 263 open-end (both load and no-load) and closed-end funds that have been continuously available during the ten-year period June 30, 1981, to June 30, 1991. Obviously, trading strategies cannot be used profitably on funds that levy substantial sales loads. We have included such funds in our simulations merely to broaden the statistical base. Of course, once defined, any trading system involving even moderately high turnover must be applied exclusively to no-load funds.

The first year of the ten-year study period was used to accumulate the values of the moving averages, so the simulation results are based on actual returns for the nine-year period June 30, 1982, through June 30, 1991.

Our initial investigation tests this simple trading rule:

> If, at the end of a week, the current price of a fund is above its moving average, own the fund during the following week. If, at the end of a week, the current price of a fund is below its moving average, own Treasury bills in the following week.

This system is designed to capitalize on trend persistency in the market; to own funds when they are relatively strong (above their moving averages), and to avoid them when they are weak (below their moving averages). We tested two types of moving averages – arithmetic, which gives equal weight to each price, and exponential, which assigns progressively greater weight to recent prices and less weight to older values. We tested every moving average strategy from a 2-week moving average to 52 weeks. The accompanying table presents the average nine-

Moving-Average Trading Systems Versus Buy & Hold Return

Buy when fund price moves above the moving average
Sell when fund price moves below the moving average

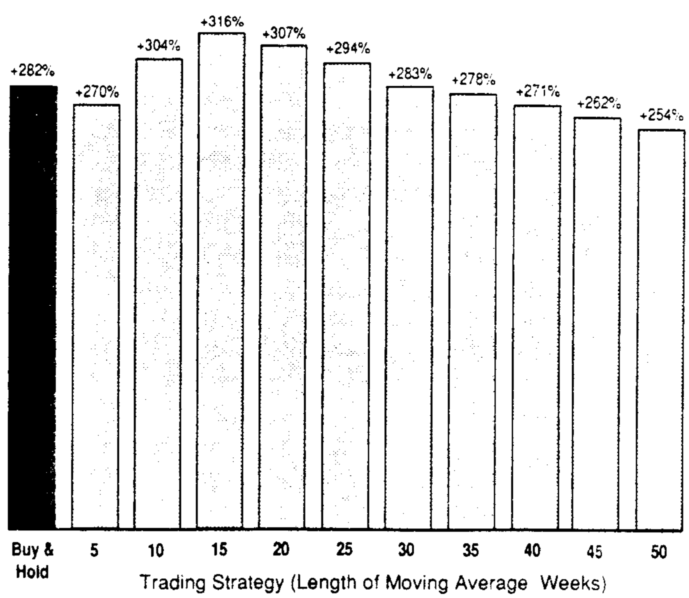

year trading results across all 263 funds for the trading strategies using arithmetic moving averages of 5, 10, 15, 20, 25, 30, 35, 40, 45, and 50 weeks. (The results using exponential averaging are only slightly superior.)

The profits of the various arithmetic moving-average strategies ranged from a low of $25,490 (the 50-week strategy) to a high of $31,600 (the 15-week), based on initial $10,000 investments. Half the strategies produced a lower return than the $28,170 an investor could have earned simply by buying the average fund at the beginning of the nine-year period and holding throughout. The 10, 15, and 20-week strategies most clearly beat buy-and-hold, producing extra profits of 8%, 12%, and 9%, respectively, while the longest-term strategies produced significantly less profit.

Strategies based on moving-average lengths between those shown in the table generally produced proportionate returns. For example, the profits achieved with 21, 22, 23, and 24-week moving-averages were between the amounts earned, and shown in the table, for the 20 and 25-week systems.

Among those moving-average systems now in popular use, the widely-publicized 39-week average approximately matched the 40-week strategy result, producing 3% less profit than a naive buy-and-hold strategy.

In defense of the longer-term moving averages, which showed poorly in this study, it should be noted that the market was unusually strong in the 1980s. The $28,170 buy-and-hold profit is equal to a compounded annual return of 16%, substantially more than the 10% total return realized in the typical year in this century. Such an environment is not amenable to trading systems. The only hope of a moving-average system is to take a trader out of the market during price downtrends. But when the market is in a broad uptrend, it pays to own mutual funds most of the time; in other words, forget trading, and just buy and hold. In the 1980s, market corrections were few and brief.

Short-term moving averages, such as 5, 10, 15, and 20-week, tend to stay much closer to their price series and therefore generate more buy and sell signals. This enables a trader to get in and out of funds faster when prices turn. The downside is that a trader can be whipsawed back and forth when the market is fluctuating in an essentially sideways direction. In contrast, long-term moving-average systems produce fewer buy and sell signals, but those signals tend to occur much later following an actual price turn. If price declines are relatively short, long-term moving-averages simply get a trader out of equities at about the time that it would be better to get back in. Long-term moving-average strategies perform much better in a market environment, such as that in the 1960s and 1970s, that is characterized by occasional long bear markets, which the moving-average systems help avoid.

The long-term moving-average strategies *are* easier to execute because they require fewer trades – an average of only 11 round-trip trades in nine years (about one a year) for the 50-week system in this test, as contrasted with an average of 58 round-trips (about one every two months) for the five-week system. On the other hand, the primary concern of an active trader should be to achieve more profit and incur less risk rather than to minimize the amount of work required to implement the system.

Reverse Strategies

We further validated the efficacy of moving average trading systems, and the underlying theory of trend-following, by applying the strategies to every fund *in reverse* – that is, assuming holding the funds only when their prices were below their moving averages, and owning T-bills when their prices were above the moving averages. Almost invariably, this produced dramatically inferior profits.

Traders Incur Less Risk

A key result of our simulations is that every moving-average system we tested significantly lowered risk. There are many ways to measure risk, of course. The one we use here is based on this easy-to-grasp principle: To the extent that a trader is out of the market and safely ensconced in T-bills, any risk of a price decline has been avoided. For example, the average result for all 263 funds (see table above) for the five-week strategy is that a trader would have been out of funds, and in T-bills, 37% of the time. Such a strategy is therefore 37% less risky than a buy-and-hold strategy that is continuously invested in funds and is always exposed to the risk of price decline.

Since every one of the trading strategies was, from time to time, out of funds and in Treasury bills, *all* of them were less risky than buy-and-hold. In the cases of the 10, 15, 20, 25, and 30-week strategies, traders would have earned more profit than buy-and-hold *and* incurred less risk. Even the strategies that earned *slightly less profit* than buy-and-hold were *much less risky* than buy-and-hold. Thus, on a risk-adjusted basis they, too, were superior.

Naturally, the results vary dramatically among different funds. Some strategies produced enormous profits for some funds and relatively minuscule profits for other funds. For example, the 5, 10, 15, 20, and 25-week arithmetic and exponential moving-average strategies applied to *Wood Struthers Neuwirth* (a no-load fund with unlimited telephone switching) all produced profits from a $10,000 initial investment of between $45,000 and $73,000. That's two to three times more than the $22,000 profit that buy-and-hold investors earned in that fund during the same period.

Case Study: 44 Wall Street Fund

The extent to which moving-average systems can reduce risk is wonderfully illustrated in the case of *44 Wall Street*. The worst-performing fund in our 263-fund population over the ten-year period, 44 Wall Street investors that put $10,000 in the fund and never sold, lost $6,814, more than two-thirds of their capital. However, *every* moving-average system applied to 44 Wall Street made money: The worst made $8,000, the best made $20,000, and the average gain was about $15,000. Of course, several of these returns were less than those achieved by the average trading strategy applied to all 263 funds. That is simply because 44 Wall Street was such a terrible performer during the period that it was usually better to own some other fund *under any circumstances*. The moving-average systems, however, tended to place the 44 Wall Street trader in Treasury bills as often, or more often, than in the fund itself. During the periods of time that the moving-average systems were invested in the fund, the fund generally made money.

For example, the 15-week arithmetic moving-average system produced a $20,010 total profit. With the system, a trader would have owned 44 Wall Street less than half the time, during which he would have achieved a 105% compounded total return. During the balance of the time, the trader was invested in Treasury bills, which produced another 47% compounded profit. During the latter times, the fund incurred an incredible 85% cumulative loss. That loss was suffered by buy-and-hold investors that owned the fund continuously, but was neatly escaped by the moving-average strategy. This illustrates how the intriguing risk-reduction capability of moving-average systems can actually transform a loser like 44 Wall Street into a winner.

Moving Average Interactions

The next step in our exploration of moving average timing systems is to determine the profitability of buying and selling when a short-term moving average of the fund's price moved above or below other moving averages.

Basing trading strategies on the crossing of two moving-average series is a common technique among technicians. For example, they might buy a stock or a mutual fund when the 10-week moving average rises above the 30-week moving average, and sell when the 10-week average falls below the 30-week average. Or, they might buy when the 5-week moving average rises above the 40-week average, and sell when the 5-week average falls below the 40-week. The guiding principle in all cases is to buy when the short-term moving averages rises above a longer-term moving average and to sell when the short-term average falls below the longer-term average.

We tested *every combination* of moving averages, from 2-week to 50-week. However, the profit-loss results from buying and selling funds when the *10-week arithmetic moving average* crossed various longer-term averages proved to be representative. We report these results.

We discovered that nearly every timing strategy using one moving average versus another was *inferior* to the strategies based on current price versus a moving average, and inferior as well to a naive buy and hold strategy. In fact, the typical nine-year total return of this variant of the system (see chart below) was about 100 percentage points less than the original. Both strategies reduced risk by about the same amount versus buy and hold (approximately one-fourth), but with their lower returns, the moving average versus moving average systems are clearly inferior to the price versus moving average approaches.

Moving-Average Trading Systems Versus Buy & Hold Return

Buy when 10-week fund price moves above the moving average
Sell when 10-week fund price moves below the moving average

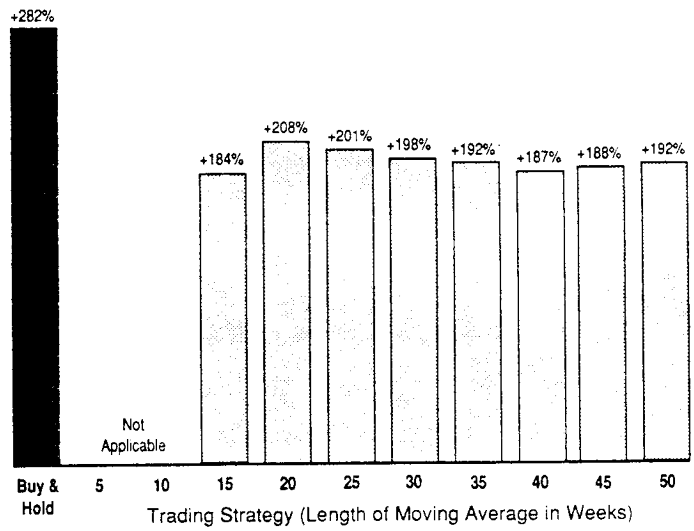

Exponential Smoothing Forecasting Techniques

First-Order (simple) Exponential Smoothing (Assumes no trend):

$$S_t = \alpha X_t + \alpha(1-\alpha) X_{t-1} + \alpha(1-\alpha)^2 X_{t-2} + \dots$$

where: S_t = smoothed value at time t

X_t = observed value at time t

α = a smoothing constant between 0 and 1

Alternatively:

$$S_t = \alpha X_t + (1-\alpha) S_{t-1}$$
$$= S_{t-1} + \alpha (X_t - S_{t-1})$$

Second-Order Exponential Smoothing with Trend Correction:

$$S_t = \alpha X_t + (1-\alpha)(S_{t-1} + T_{t-1})$$

where: T_{t-1} = estimate of trend at time $t-1$

$$T_t = \beta(S_t - S_{t-1}) + (1-\beta) T_{t-1}$$

where: β = a smoothing constant between 0 & 1

After computing S_t and T_t for each time period, compute a forecast at time t for time $t+k$, $F_{t,t+k}$:

$$F_{t,t+k} = S_t + k T_t$$

S&P 500 Daily Closing & 100 Day Mov Ave

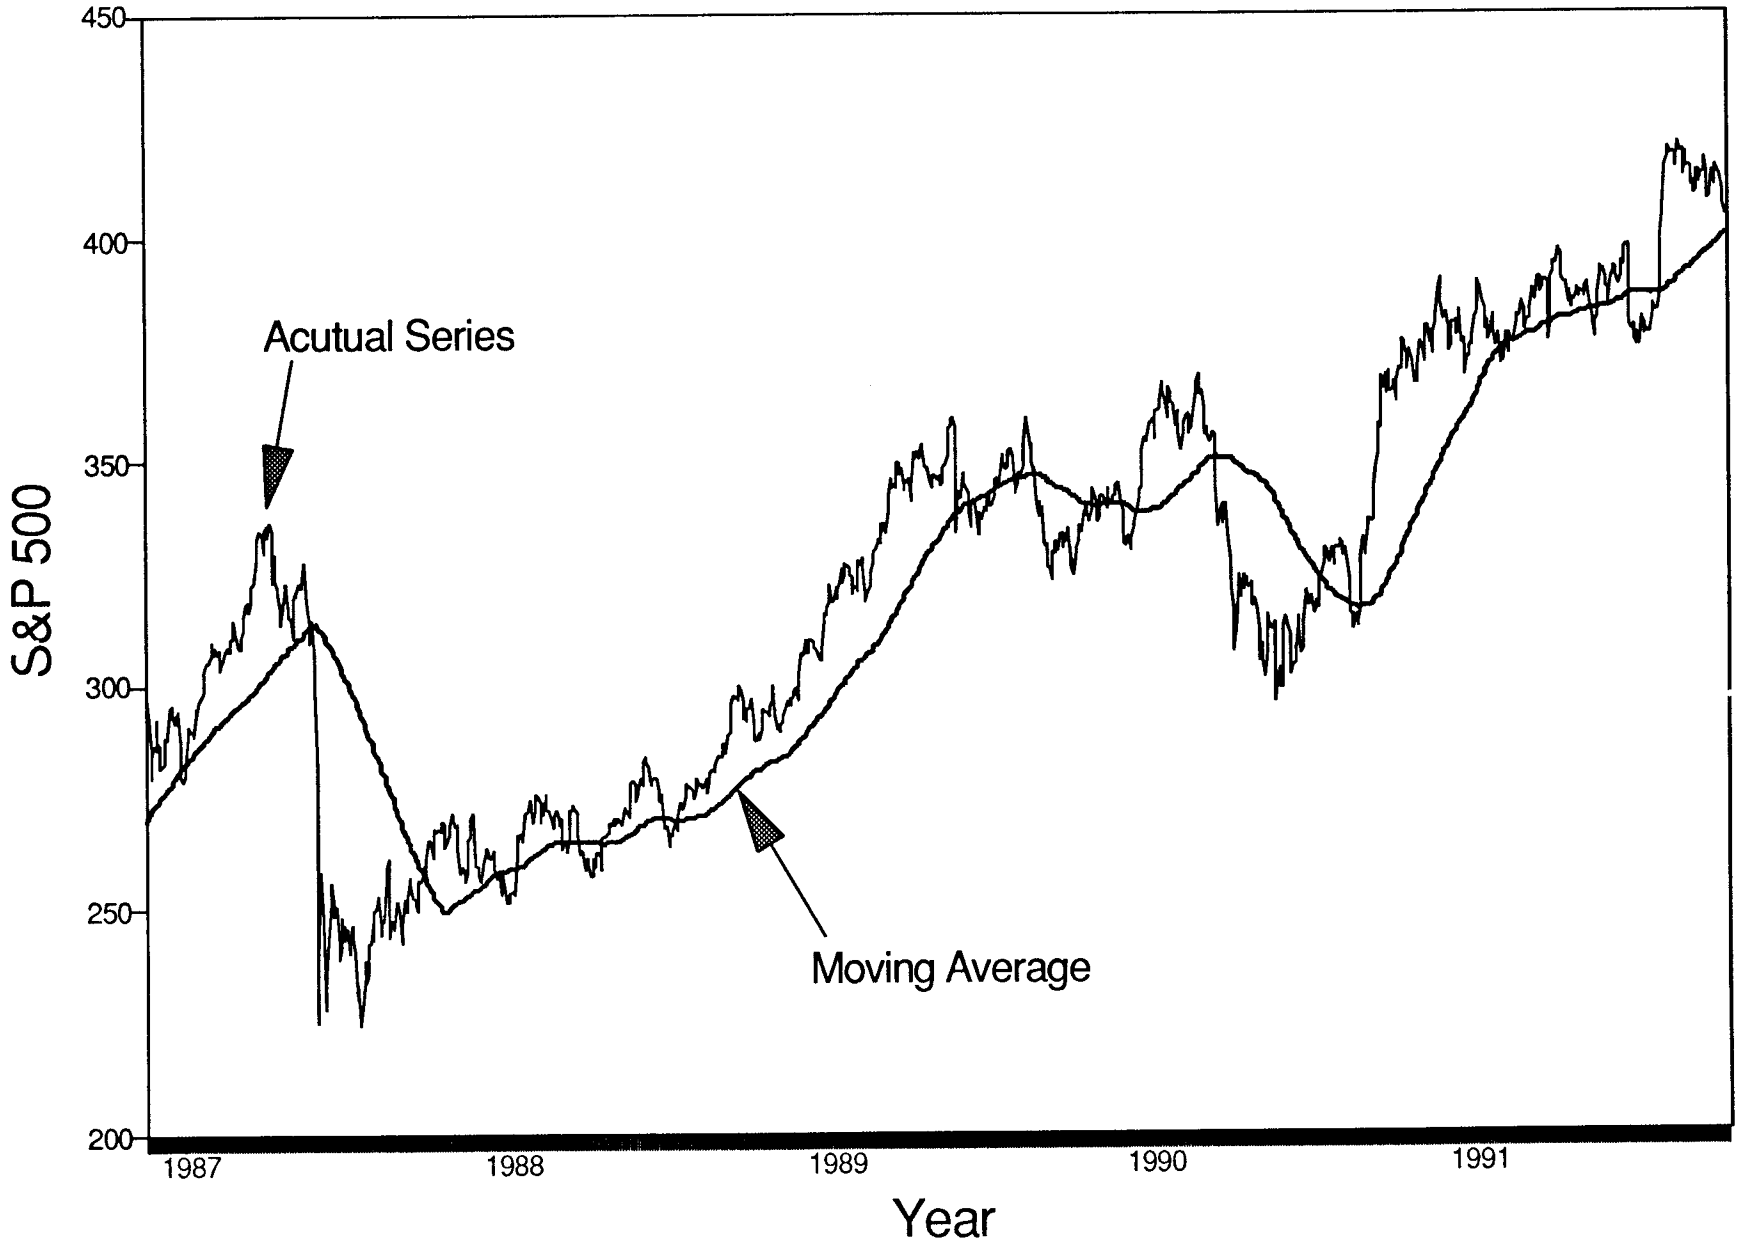

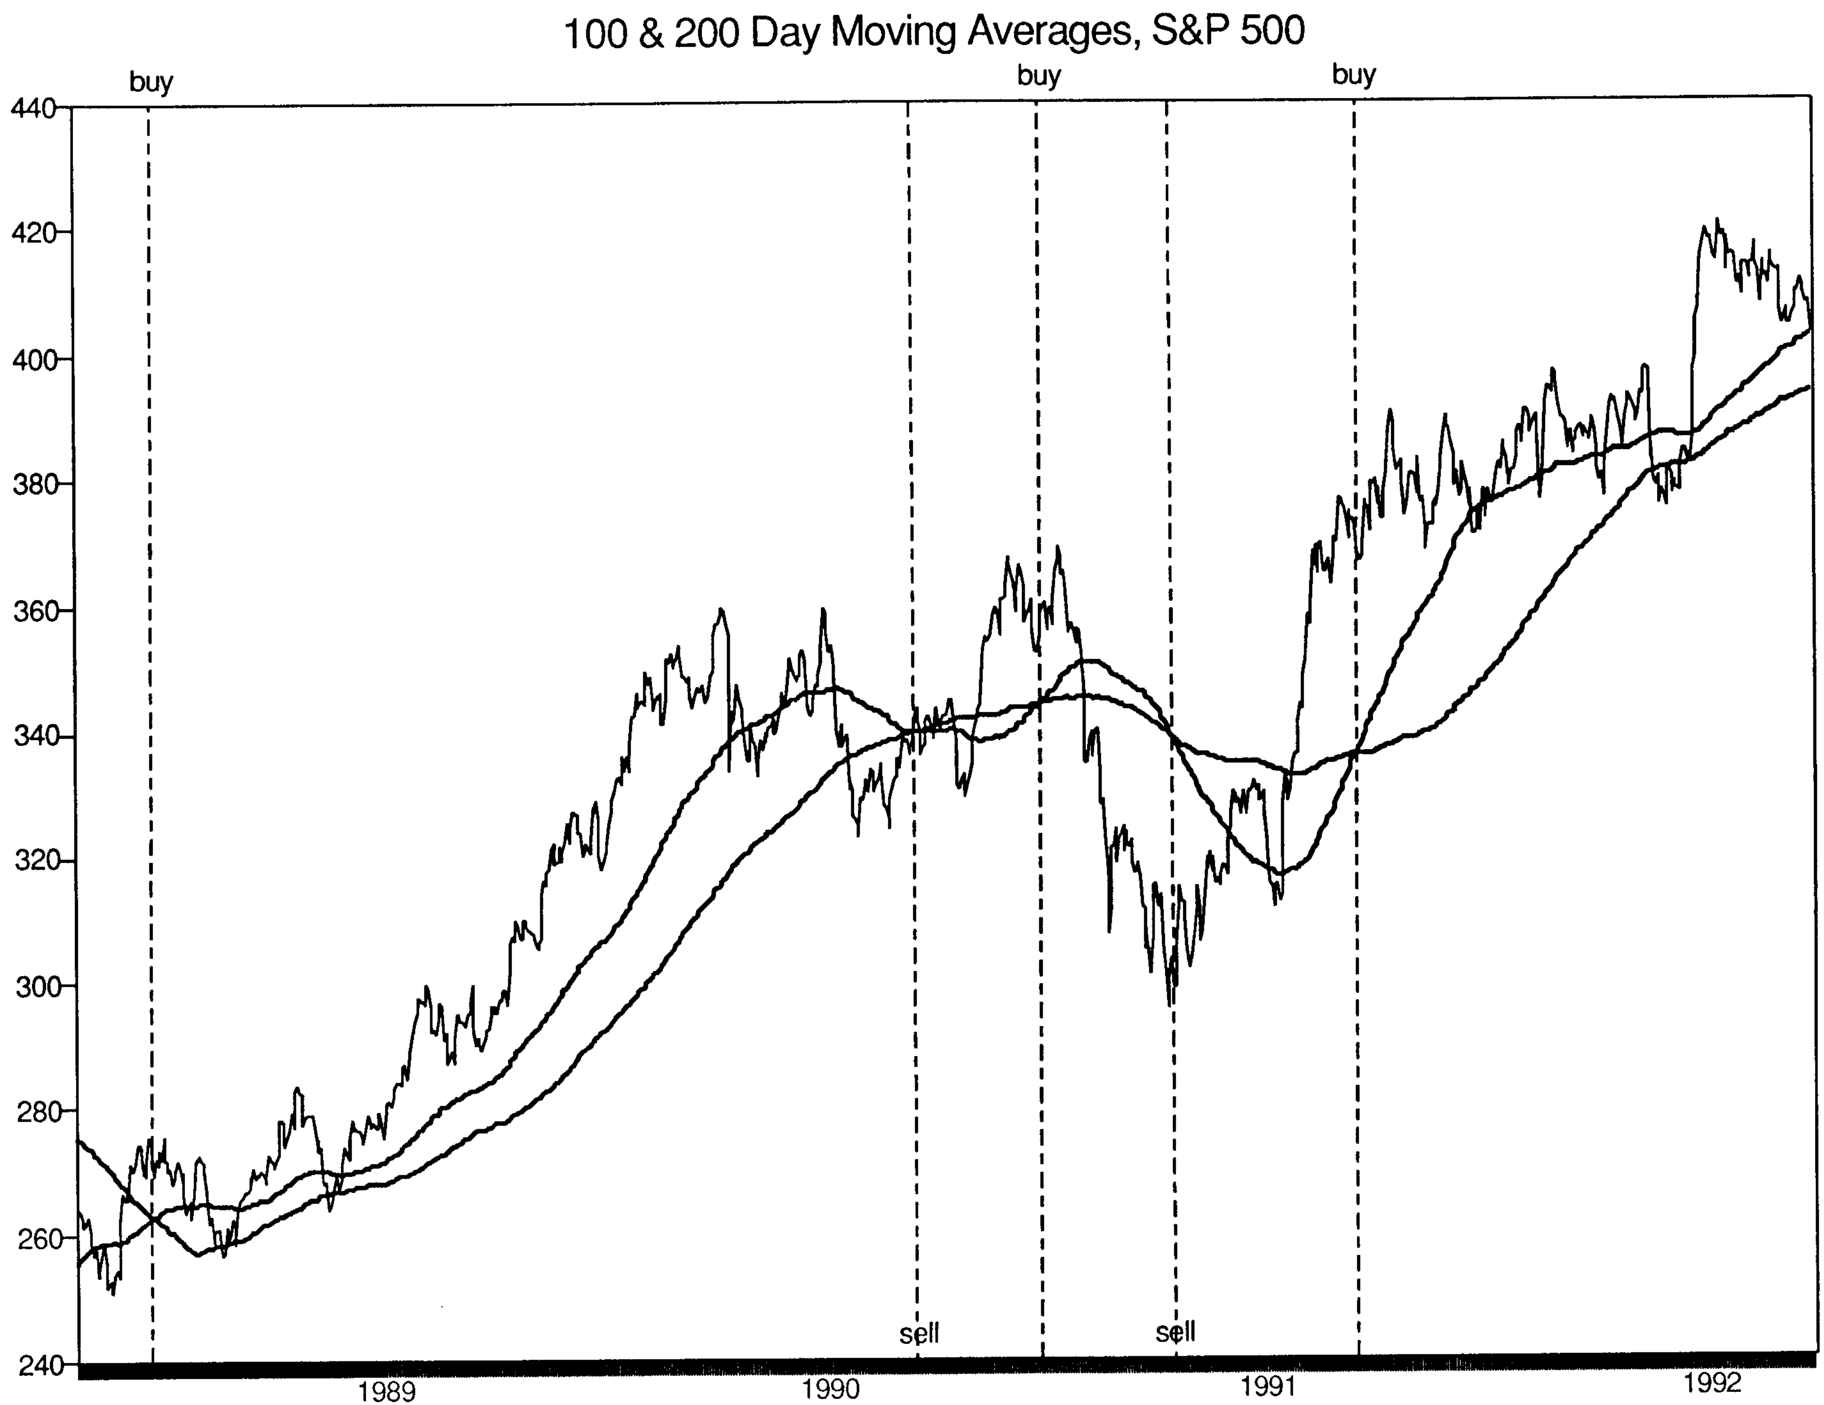

100 & 200 Day Moving Averages, S&P 500
buy
buy
buy
sell
sell
440
420
400
380
360
340
320
300
280
260
240
1989
1990
1991
1992

Professor Stipak

Tricky Short-Cut (TSC) Technique for Setting Up Forecasting Error Analysis in Spreadsheets, Compared to Straight-Forward (SF) Method

A standard way to analyze forecasting error is to compute the mean squared error. We will go over in class a straight-forward method for setting this up in a spreadsheet. Besides the straight-forward (SF) method, there is a "tricky short-cut" (TSC) technique you can use, as shown in the *Lotus Magazine* article, "Short-Range Forecasting." The discussion below tries to clarify these alternatives by contrasting the SF and the TSC methods.

SF Method of Computing Mean Squared Forecast Error

1) Add a column, to the right of the forecasted value, for the forecast error (actual value minus the forecast). 2) Add a column to the right of the forecast error for the squared forecast error. 3) Use the @AVG function to compute the average of this column, yielding the mean squared forecast error.

TSC Method of Computing Mean Squared Forecast Error

1) Add a column, to the right of the forecasted value, for the forecast error (actual value minus the forecast). 2) Calculate the mean squared forecast error directly from the forecast error column using the following formula:

@VAR(forecast error range) + @AVG(forecast error range)^2

The reason this works is because the mean squared error equals the variance plus the square of the mean.

SF Method of Doing Sensitivity Analysis of Which Values of the Smoothing Constants Yield the Best Forecasts for an Exponential Smoothing Model Spreadsheet

Set up your spreadsheet so that you can enter the value for the smoothing constant(s) in an input section near the top of the spreadsheet, and so that you can see the resulting mean squared error, also near the top of the spreadsheet. Try different values of the smoothing constant(s) and see what the result is on the mean squared error.

TSC Method of Doing Sensitivity Analysis of Which Values of the Smoothing Constants Yield the Best Forecasts for an Exponential Smoothing Model Spreadsheet

Use the built-in "what-if" capability of the spreadsheet program to automate trying out different values of the smoothing constant(s). In Lotus use the **/Data Table 1** command, in Quattro Pro the **/Tools What-If 1** command, and in Excel use the **Data Table** command to do what-if analysis for one smoothing constant. To do what-if analysis for both smoothing constants in a second-order exponential smoothing model, use the **/Data Table 2** (Lotus) or **/Tools What-if 2** (Quattro Pro) commands.

Alternative: Use Solver: Change γ, β, Min MSE, s.t. γ β constraint

LOTUS, Vol. 3, No. 2, Feb. 1987

SHORT-RANGE FORECASTING

If prognostication is your vocation, use exponential smoothing the next time you gaze into the future.

BY EVERETTE S. GARDNER JR.

Business decisions depend on good forecasting. Sales forecasting, for example, drives a wide range of decisions: staffing, working-capital needs, inventory levels, production schedules, pricing, advertising, and sales-promotion strategies.

To the spreadsheet user, forecasting is the art of writing formulas that reflect reasonable assumptions about the future. This is the first of two articles that explain how to develop such formulas using a technique called exponential smoothing, which saves you time and usually improves the accuracy of your forecasting.

Exponentially smoothed forecasts are based on weighted averages of past data. The basic assumption of exponential smoothing is that recent data is a better predictor of future performance than older data. Thus, you assign the weights so that recent data gets the greatest weight. Exponential smoothing gets its name because the rate of decline in the weights follows an exponential pattern.

More than 25% of U.S. corporations produce forecasts using some form of exponential smoothing. Why? The computations are relatively simple, certainly simpler than regression modeling, for example. More important, smoothing is quite accurate and compares favorably to highly complex forecasting models.

This article concentrates on short-range forecasting models that look one time period (month, quarter, or year) into the future. Such models are widely used for budgeting and for forecasting sales of mature products. Next month's article will explain a model for projecting trends several time periods into the future.

CRUISE CONTROL

The simplest exponential smoothing model has two equations that are updated each time period:

Error = Data − Current Forecast
Next Forecast = Current Forecast + W * Error

W is a fraction, between Ø and 1, called the smoothing weight.

This model works like an automatic pilot, a thermostat, or a cruise control on an automobile. As the actual level of the data changes, the model continuously adjusts the forecasts over time.

Suppose you are forecasting monthly expenses. Each month you compute the forecast error, which is a positive, negative, or zero value. If the current forecast was too low, the error has a positive sign, and the next forecast is the sum of the current forecast plus a fraction of the error. If the current forecast was high, the error value is negative, and the next forecast is the current forecast minus a fraction of the error. Finally, if you get lucky and the current forecast is perfect, the error is zero, and there is no change in the next forecast.

An example of the model appears in figure 1, a graph of monthly travel expenses and forecasts. The expenses are for the sales force of a pharma-

Everette S. Gardner Jr. is associate professor of decision and information sciences at the University of Houston as well as associate editor of the International Journal of Forecasting.

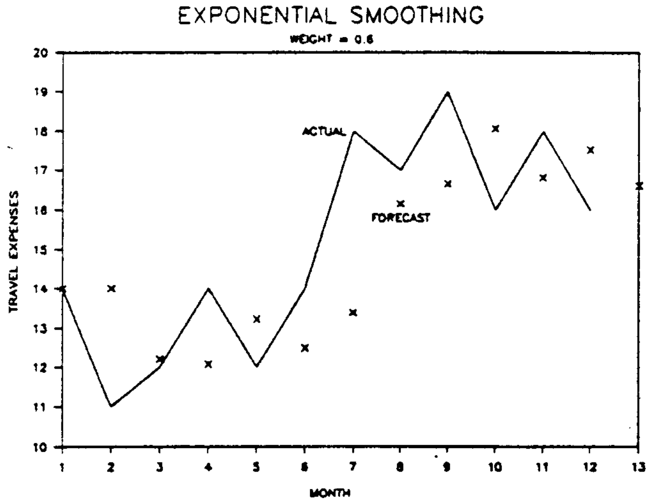

FIGURE 1. This graph of actual travel expenses and forecasts demonstrates the accuracy of exponential smoothing. In July, actual expenses rose to about $18,000 per month, and the July forecast was far too low. In August, however, the forecasts reacted to the change and were close to the new level of expenses thereafter.

	A	B	C	D
1	Exponential Smoothing			
2	Weight =	0.6		
3				
4		Actual	Forecast	
5	Month	Expenses	Expenses	Error
6	1	14	14.0	0.0
7	2	11	14.0	-3.0
8	3	12	12.2	-0.2
9	4	14	12.1	1.9
10	5	12	13.2	-1.2
11	6	14	12.5	1.5
12	7	18	13.4	4.6
13	8	17	16.2	0.8
14	9	19	16.7	2.3
15	10	16	18.1	-2.1
16	11	18	16.8	1.2
17	12	16	17.5	-1.5
18	13		16.6	

FIGURE 2. This simple spreadsheet forecasts the expenses graphed in figure 1. After entering the month numbers, actual expenses, and smoothing weight, you enter two simple formulas that calculate the current month's forecast error and the next month's expenses. As you see, the forecasts change with the errors. If the error for one month is negative, the forecast for the next month goes down, and vice versa.

ceutical manufacturer. The company's controller uses the smoothing model to update cash budget projections on a monthly basis.

Fluctuating around an average expense of about $13,000 per month, expenses for the first six months of the year are relatively stable. In July the company added a new product and increased its sales force. This caused expenses to jump to a new level of about $18,000 per month. As shown in the graph, the July forecast was far too low; however, in August the forecasts reacted to the change and were close to the new level of expenses thereafter.

SPREADSHEET SMOOTHING

Let's look at how to produce these forecasts of expenses. The completed spreadsheet appears in figure 2. Begin by entering the labels in rows 1, 2, 4, and 5. In cell B2 enter the smoothing weight of 0.6. (You'll read more about how to choose this weight in a moment.) Enter the numbers for the months in column A and the actual expense figures in column B. (Note that all expense figures, actual and forecast, represent thousands of dollars.)

Columns C and D contain formulas that calculate the forecasts and errors. To start the smoothing process, you enter the first forecast in cell C6. A good rule of thumb is to set the first forecast equal to the first actual expense figure. Therefore, enter +B6 in cell C6. In cell D6 enter the formula +B6−C6, which calculates the error.

In cell C7 the first updating of the forecast occurs. Enter the formula +C6+$B$2*D6 in cell C7. This formula states that the forecast for month 2 is equal to the forecast for month 1 plus the smoothing weight times the error in the forecast for month 1. Format range C6..D7 by selecting /Range Format Fixed (in *Symphony*, MENU Format Fixed) and specifying one decimal place. Copy the updating formula in cell C7 to range C8..C18. Copy the error-calculating formula in cell D6 to range D7..D17. Cell C18 of the model forecasts that the total expense figure for month 13 will be 16.6. Because the actual expense figure is not yet available for month 13, cell B18 is blank.

Now examine the way the forecasts change with the errors. The month 1 error is zero, so the month 2 forecast is unchanged. The error is negative in month 2, so the month 3 forecast goes down. In month 3, the error is negative again, so the month 4 forecast goes down. In month 4 you get the first positive error, so the month 5 forecast goes up.

As actual expense figures come in each month, updating the spreadsheet is easy. For example, after you enter actual expenses for month 13 in cell B18, two simple copy operations automatically forecast month 14 expenses. Just copy cell D17 to cell D18, then copy cell C18 to cell C19.

WEIGHTY MATTERS

The performance of the smoothing model is controlled by the weight, which ranges from 0 to 1. The higher the value of the weight, the quicker the model will respond to true changes in the data. On the other hand, as the value of the weight increases, the model may become unstable—that is, the forecasts may overreact to purely random fluctuations in the data.

Thus, choosing the best weight is a trade-off. If you are concerned about the possibility of sudden changes in the data, use a high weight. If your data contains a lot of randomness, use a low weight to avoid instability.

	E	F
1	MSE =	4.71
2		
3	Weight	MSE
4		+F1
5	0.0	8.27
6	0.1	7.26
7	0.2	6.25
8	0.3	5.53
9	0.4	5.08
10	0.5	4.82
11	0.6	4.71
12	0.7	4.73
13	0.8	4.86
14	0.9	5.10
15	1.0	5.45

FIGURE 3. Because squaring the errors gives extra weight to large errors, the mean squared error (MSE) of a smoothing weight is a widely used measure of how well a weight performs. Using a data table enables you to select the weight that produces the lowest MSE for use in your forecast.

You can choose the weight intuitively by scanning a graph of your data and picking a value that seems to give a realistic forecast. A more objective approach is to use a *1-2-3* data table or a *Symphony* what-if table that automatically substitutes different weights into the model and gives you a summary of performance.

Figure 3 shows a data table of the mean-squared error (MSE) at different smoothing weights. The MSE is a measure of how well the weight performs. Although you could use the mean error or the mean absolute error as a measure of performance, the MSE is the most widely used measure because squaring the errors gives extra weight to large errors. In most businesses, you can learn to live with small errors, but large errors are extremely disruptive.

As shown in figure 3, the smallest MSE occurs at a weight of 0.6. Therefore, it's safe to assume that 0.6 is the best weight to use in your forecast. To set up this table, enter the labels in rows 1 and 3. In cell F1 enter the following formula and then format the cell for Fixed with two decimal places:

@VAR(D6..D17)+@AVG(D6..D17)^2

This short-cut method of computing the MSE produces the same answer as computing the square of each error and then taking the average of the squares.

Enter the weights in range E5..E15 and then format the range for Fixed with one decimal place. In cell F4 enter +F1. Cell F1 contains the formula evaluated by the data table. It is a good idea to format cell F4 so that it serves as a reminder of which cell is being evaluated. Select /Range Format Text (in *Symphony*, MENU Format Other Literal) and indicate cell F4 as the range to format. Next select /Data Table 1 (in *Symphony*, MENU Range What-If 1-Way). Indicate a Table range of E4..F15. Finally, indicate an Input cell of B2, which contains the weight.

The data-table command automatically substitutes each weight in column E into cell B2, computes the MSEs using the formula in cell F1, and records the results in column F. Format range F5..F15 for Fixed with two decimal places.

	A	B	C	D
1	Analysis of Weights on Past Data			
2	Smoothing Weight =		0.6	
3	Nbr. of Periods =		12	
4				
5		Actual		
6	Month	Expenses	Weight	Data*Weight
7	0	14	0.00002	0.00023
8	1	14	0.00003	0.00035
9	2	11	0.00006	0.00069
10	3	12	0.00016	0.00189
11	4	14	0.00039	0.00551
12	5	12	0.00098	0.01180
13	6	14	0.00246	0.03441
14	7	18	0.00614	0.11059
15	8	17	0.01536	0.26112
16	9	19	0.03840	0.72960
17	10	16	0.09600	1.53600
18	11	18	0.24000	4.32000
19	12	16	0.60000	9.60000
20				
21	Sum		1.00000	16.61219

FIGURE 4. An alternative approach to exponential smoothing is taking the weighted averages of past data. Although this process and smoothing produce the same results, exponential smoothing requires less work.

By simply adjusting the weight values displayed in column E, you can make the data table of MSEs as detailed as you like. However, testing weights in increments of 0.1 is usually good enough.

BEWARE OF EXTREMES

Be careful if the best weight—that is, the one with the lowest MSE—turns out to be zero. A weight of zero causes the model to ignore all errors and never to adjust the first forecast. You can prove this to yourself by plugging a weight of zero into cell B2 of figure 2. Using a minimum weight of 0.1 gives you some protection against changes in the data.

What if the best weight turns out to be 1.0? A weight of 1.0 indicates that each forecast is equal to the last data point: There is no smoothing at all. This could indicate a trend or a regular pattern of growth in your data. One way to cope with a trend is to use a smoothing weight greater than 1. A better approach is to use a model specifically designed for trends. Such a model will appear in next month's article.

A weight of 1.0 does not always mean that a trend exists. There are many data sets in which basic smoothing with a weight of 1.0 is the best forecasting model that statisticians can devise. Stock prices are a good example of this type of data. The best statistical forecasting model for most stock prices is the so-called random-walk model, which is the same as exponential smoothing with a weight of 1.0.

AN ALTERNATIVE APPROACH

Many people compute forecasts by taking weighted averages of past data. This is an intuitively appealing

FORECASTING *continued from page 56*

idea that works well in many applications. If you are currently using a weighted average to forecast, you can save a lot of work by using smoothing instead. It may surprise you that exponential smoothing is actually a short-cut method of taking a weighted average. Let's see the process you'd use to forecast expenses using the weighted-average method.

Figure 4 shows how the smoothing weight determines the individual weights that go into the average. Beginning in an empty spreadsheet, enter the labels shown in rows 1, 2, 3, 5, 6, and 21. Enter the smoothing weight, 0.6, in cell C2, then enter the number 12, which represents the number of periods of available data, in cell C3.

EXPONENTIAL SMOOTHING IS QUITE ACCURATE AND COMPARES FAVORABLY TO HIGHLY COMPLEX FORECASTING MODELS.

Enter the period numbers, which in this case represent months, in range A7..A19, then enter the actual expenses in range B7..B19. The first data observation (month Ø) is the initial forecast. This gets some weight, as you will see in a moment.

The formulas in column C calculate the individual weight on each data point. Enter the formula +(1−C2)^C3 in cell C7. This is the weight on the initial forecast. To obtain the weight on the month 1 data, enter the following formula in cell C8:

+$C$2*(1−$C$2)^($C$3−A8)

Now copy cell C8 to range C9..C19.

The formula in cell C21 is @SUM(C7..C19). In all exponential-smoothing models, the sum of all weights on past data is equal to 1. Notice how the weights decline as the data gets older.

In column D, you multiply each data point times its weight. Enter the formula +B7*C7 in cell D7, then copy cell D7 to range D8..D19. Finally, enter @SUM(D7..D19) in cell D21. The result is 16.6, the forecast for period 13. As you see, this is the same forecast you got for period 13 using the much simpler approach illustrated in figure 2.

Next month's article, which extends the basic smoothing model to handle trends, shows why linear regression is not appropriate for forecasting a trend and offers some tips on applying smoothing models to different types of business data. ■

USING EXPONENTIAL SMOOTHING

Think linear regression is just a line? Smooth your way into the future with this model for producing accurate, realistic forecasts for data exhibiting a trend.

sed by Permission:Everette S. Gardner, Jr. (1987). Using Exponential Smoothing, Lotus Magazine, March 1987, 61-62, 64-64, 6-66. Lotus

BY EVERETTE S. GARDNER JR.

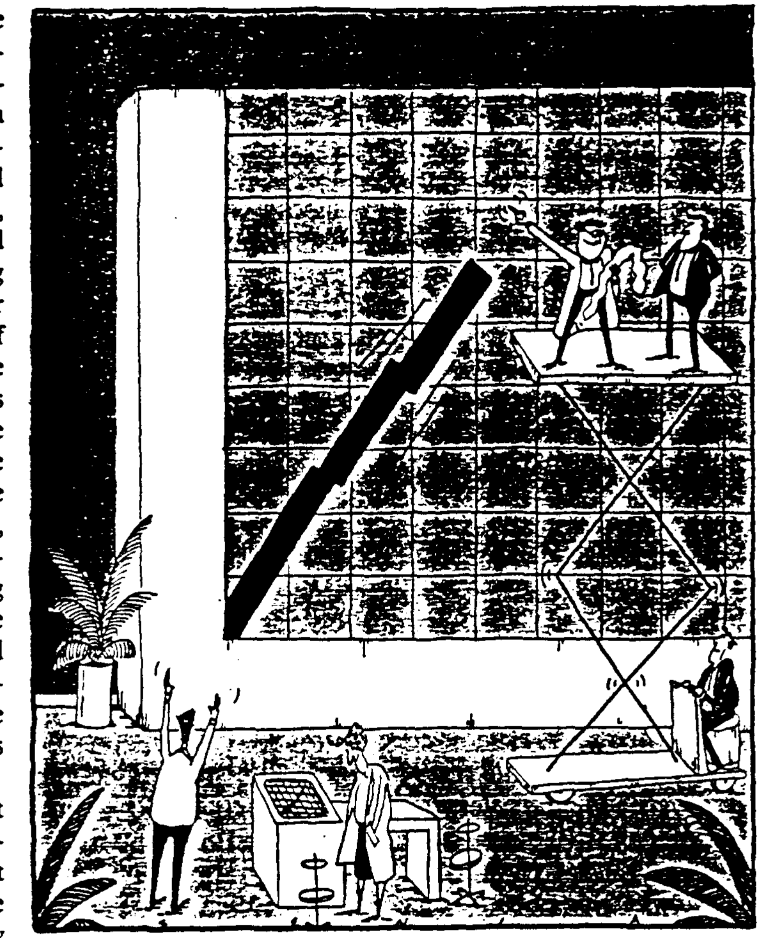

Last month's article discussed the rudimentary exponential-smoothing model, a handy tool for tracking the average level of data such as sales, inventory usage, and expenses. Assuming there is no trend or regular pattern of growth or decline, the basic model gives good forecasts for one time period into the future. However, if the data exhibits a trend, forecasts from the basic model usually lag behind. That is, the forecasts are too small when the trend is going up and too large when the trend is downward.

In practice, the most common way of forecasting a trend is to fit a straight line to the data. This is usually done with linear regression, which computes the best-fitting line. The problem with regression is that, in computing the forecast, it gives equal weight to all past data. Exponential smoothing gives more weight to recent data and usually results in more accurate forecasts.

The difference between straight-line forecasting and smoothing a trend is demonstrated in figure 1, a graph of annual dollar sales and forecasts for a pharmaceutical manufacturer. Actual sales (adjusted for price changes) cover the years from 1975 through 1986. Forecasts extend three years into the future, through 1989.

The regression line was computed using time (the numbers 1 through 12) as the independent variable and sales as the dependent variable. However, since the goal of this article is to convince you not to use linear regression when forecasting a trend, we won't go into the details of computing the regression.

Although the regression line is the best-fitting straight line for this data, it is a poor forecasting model. As shown in the graph, the actual rate of growth in sales changed drastically several times. The regression line runs below actual sales from 1978 to 1982 and above actual sales from 1983 to 1986. You'll probably agree that the regression forecasts for 1987 to 1989 are too optimistic.

The Xs in figure 1 represent the exponentially smoothed forecasts. Notice how these forecasts adapt to changes in sales. From 1977 to 1982, sales grew rapidly, and so did the forecasts. Actual sales fell in 1983. Therefore the 1984 forecast dropped substan-

Everette S. Gardner Jr. is associate professor of decision and information sciences at the University of Houston and associate editor of the International Journal of Forecasting.

MARK KSENAX

FORECASTING *continued from page 61*

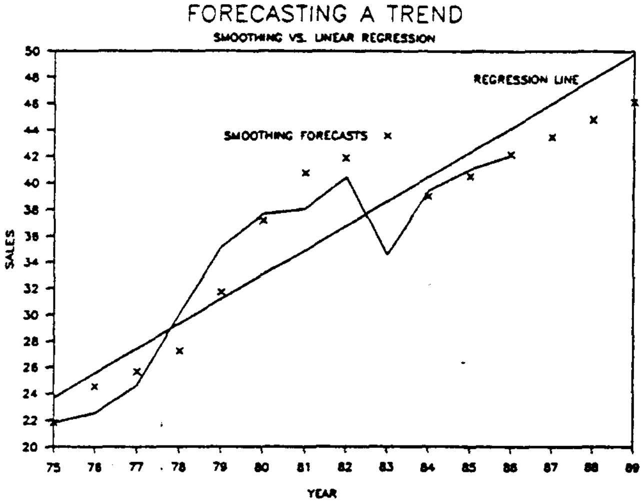

FIGURE 1. The most common way of forecasting a sales trend is to compute a linear-regression line; however, regression forecasting is unrealistic when trends are changing. In this example, regression forecasts of sales for 1986 to 1988 are too large. In contrast, exponential smoothing adapts to changing trends. A smoothing model gives more weight to recent sales in computing the expected trend. After the dip in sales in 1983, smoothing forecasts change drastically. For 1986 to 1988, smoothing forecasts are more realistic than regression.

	A	B	C	D	E	F
1		Exponential Smoothing With Trend				
2						
3		Control Weight		0.40		
4		Level Weight		0.64		
5		Trend Weight		0.16		
6						
7	Year	Sales	Forecast	Error	Level	Trend
8					19.07	2.73
9	75	21.8	21.80	0.00	21.80	2.73
10	76	22.5	24.53	-2.03	23.23	2.41
11	77	24.6	25.64	-1.04	24.97	2.24
12	78	30.0	27.22	2.78	29.00	2.69
13	79	35.1	31.68	3.42	33.87	3.23
14	80	37.6	37.10	0.50	37.42	3.31
15	81	38.0	40.73	-2.73	38.98	2.88
16	82	40.4	41.86	-1.46	40.93	2.64
17	83	34.5	43.57	-9.07	37.76	1.19
18	84	39.4	38.96	0.44	39.24	1.26
19	85	41.0	40.50	0.50	40.82	1.34
20	86	42.1	42.16	-0.06	42.12	1.33
21	87		43.45			
22	88		44.79			
23	89		46.12			

FIGURE 2. Exponential smoothing with trend continuously adjusts the forecasts as sales change. Each forecast is the sum of the level and trend for the previous year. The level is the average value of sales; the trend is the average growth per period. To start the smoothing process, you enter an initial level and trend in cells E8 and F8. Thereafter, formulas adjust level and trend according to the forecast error. The amount of adjustment is determined by the weights in cells D4 and D5. Each level value is adjusted by adding .64 times the forecast error to the previous level. Each trend is adjusted by adding .16 times the forecast error to the current forecast. At the end of 1986, level and trend estimates are 42.12 and 1.33, respectively. The sum (43.45) is the forecast for 1987. The forecast for 1988 is 43.45 plus another increment of the trend, or 44.79.

tially. From 1984 on, the forecasts reflect slower growth in sales.

THE BIG PICTURE

Figure 2 shows the spreadsheet used to develop the trend-smoothing forecasts. Before delving into the formulas that create the forecast, here's an overview of how the spreadsheet works.

The years are listed in column A. Actual sales and forecasts appear in columns B and C, respectively. The forecast error for each year, shown in column D, is calculated by subtracting the forecast from the actual sales figure.

Columns E and F show the component parts of each forecast: level and trend. The level is a smoothed estimate of current sales computed at the end of the year. The trend is a smoothed estimate of the average growth computed at the end of the year. This estimate also represents the growth expected next year.

To get started, you enter initial values for level and trend in cells E8 and F8, respectively. Thereafter, you smooth each year's level and trend values by adding a fraction of the error to each. Cell D3 contains a control weight that governs the model's reaction to changes in the data. Cells D4 and D5 show the individual weights used to smooth the level and trend.

Each forecast is the sum of the latest estimate of level and trend. For example, the forecast for 1975, 21.80, is the sum of the initial level and trend (19.07 + 2.73). The forecast for 1976, 24.53, is the sum of the level and trend at the end of 1975 (21.80 + 2.73).

At the end of 1986 (row 20), you can forecast as many years into the future as you like. The forecast for 1987 is 43.45. This is the level plus the trend at the end of 1986 (42.12 + 1.33). The forecast for 1988, 44.79, is the 1987 forecast plus another increment of trend (43.45 + 1.33).

THE DETAILS

With this background, let's set up the spreadsheet. As shown in figure 2, enter the labels in column B and row 7, the years in column A, and the actual sales figures in column B. Format range B9..B20 by selecting /Range Format Fixed (in *Symphony*, MENU Format Fixed) and specifying one decimal place. Then follow these steps.

1. Choose the control weight. This model has a "control" smoothing weight that determines how fast the forecasts react to changes in the data. You'll use two formulas to convert the control weight into the individual weights for level and trend, shown in cells D4 and D5. For now, enter a control weight of 0.40 in cell D3. (You'll see why this is the best control weight in a moment.)

2. Compute the individual smoothing weights for level and trend. Enter the formula +D3*(2−D3) in cell D4. In cell D5, enter the formula +D3^2. Cell D4 computes a weight used in smoothing the level; cell D5 computes a weight used in smoothing the trend.

The effect of these formulas will be explained later. Format ranges D3..D5 and C8..F23 by selecting /Range Format Fixed (in *Symphony*, MENU Format Fixed) and specifying two decimal places.

3. Enter the initial values for level and trend. You must start the smoothing process with reasonable initial values; otherwise, the forecasts will be distorted. A good rule of thumb is to estimate the initial trend value by calculating the average difference between the first four actual sales figures. To do this, enter the following formula in cell F8:

((B12 − B11) + (B11 − B10) + (B10 − B9))/3

To get the initial level, subtract the initial trend from the first actual sales figure. In cell E8, enter the formula + B9 − F8.

4. Enter the first row of updating formulas, which update the forecast, error, level, and trend for each period. You enter these formulas in the row containing the first actual sales figure (in this case, row 9). The forecast is the sum of the previous year's level and trend. Enter the formula + E8 + F8 in cell C9. You calculate the error by subtracting the current year's forecast from the current year's actual sales figure. Enter the formula + B9 − C9 in cell D9. The level is the current forecast plus the level weight multiplied by the error. Enter the formula + C9 + $D$4*D9 in cell E9. The last formula updates the trend, which is the sum of the previous year's trend and the trend weight multiplied by the error. Enter the formula + F8 + $D$5*D9 in cell F9.

5. Copy the row of updating formulas down the columns for each actual sales figure. Select /Copy (in *Symphony*, MENU Copy) and indicate C9..F9 as the From range and C10..F20 as the To range.

6. Enter the formulas that forecast beyond the last actual sales figure. This final step projects the trend into the future. For 1987, the first year ahead, enter the formula + E20 + F20 in cell C21. As you may notice, this is the same formula used for the previous years' forecasts in column C. However, to forecast 1988, the next year ahead, add another increment of the last calculated trend value to the forecast displayed in cell C21. To do so, enter the formula + C21 + $F$20 in cell C22. Since the reference to cell F20 (which contains the value 1.33) is absolute, you can easily extend the forecasts by copying cell C22 as many times as necessary. For example, to forecast 1989, copy cell C22 to cell C23.

CONTROLLING YOUR WEIGHT

Now let's review how to create a data table that assists you in selecting the best control weight. (For a tutorial on data tables, see the August and September 1986 issues.) You select a control smoothing weight from Ø to 1. A high weight causes the model to react quickly to true changes in the data. However, a high weight may also cause the model to overreact to random changes in the data. A low weight has the opposite effect. A model using a low weight may be slow to react to true changes in the data, but it won't overreact to randomness in the data.

	I	J
1	MSE =	9.76
2		
3	Control	
4	Weight	MSE
5		+J1
6	0.00	26.36
7	0.10	16.55
8	0.20	12.27
9	0.30	10.39
10	0.40	9.76
11	0.50	9.81
12	0.60	10.31
13	0.70	11.28
14	0.80	12.87
15	0.90	15.33
16	1.00	19.08

FIGURE 3. The control weight determines how quickly forecasts respond to changes in sales. This table shows the mean-squared-error (MSE) for the range of control weights in column I. The MSE is the mean of the squared errors in column D of figure 2. The best control weight for this data is 0.40 because it minimizes the MSE.

The most objective way to select a control weight is to calculate the mean-squared-error (MSE) of each weight and then choose the weight with the lowest MSE. Figure 3 shows a data table of weights and MSEs. To build this table, enter the labels shown in rows 1, 3, and 4. Then format range I1..J16 by selecting /Range Format Fixed (in *Symphony*, MENU Format Fixed) and specifying two decimal places. Now enter the weights in range I6..I16. In cell J1 enter the following formula, which calculates the mean-squared-error:

@VAR(D9..D20) + @AVG(D9..D20)^2

Enter +J1 in cell J5. Then to remind yourself of which cell the data table is evaluating, format cell J5 by selecting /Range Format Text (in *Symphony*, MENU Format Other Literal). Now select /Data Table 1 (in *Symphony*, MENU Range What-If 1-Way). Indicate a Table range of I5..J16 and an Input cell of D3. The program substitutes each weight into cell D3 and records the mean-squared-errors in column J. The smallest MSE is 9.76, at a weight of 0.40.

Be careful if the weight producing the lowest MSE turns out to be Ø or 1. If you use a weight of Ø, the model never updates the initial level and trend values. Using a minimum weight of 0.1 gives you some protection against changes in the data.

If you use a weight of 1, the model performs no smoothing at all. The full amount of the error is added to both the level and trend values for each forecasted period. If the best weight is 1, this trend model may be inadequate for your data. This model assumes a linear growth—that is, the amount of increase (or decrease) each period is constant—but true growth in your data may be exponential—that is, the amount of increase (or decrease) increases each period. Another possibility is that your data has been distorted by unusual events.

FORECASTING *continued from page 64*

LEVEL AND TREND WEIGHTS

In figure 2, the formulas in cells D4 and D5 are used to convert the control weight to individual weights for level and trend. Because the mathematical basis of these formulas is beyond the scope of this article, you must take these formulas on faith. However, you should understand the effect of the formulas: The level weight is always higher than the trend weight. To see why, compare the numbers for level and trend in columns E and F of figure 2, which show that the level for any period is much higher than the trend. This is almost always true in business data. Thus the formula that calculates the level is adjusted by a larger fraction of the error than is the trend.

Another way to get level and trend weights is to select them individually—that is, not to use a control weight at all, but to plug weights directly into cells D4 and D5 of figure 2. Selecting level and trend weights individually is a lot more work, but it sometimes improves forecast accuracy. If you want to try this, use the Data Table 2 command to help select the weights.

USING AND MISUSING EXPONENTIAL SMOOTHING

You should keep several caveats in mind when using exponential smoothing. Most important, smoothing models are only as good as the data that you give them. "Scrubbing," or adjusting the data to remove the effects of one-time events, such as a distribution problem or a special promotion, can have a tremendous impact on accuracy. You should also adjust for the number of selling days in each month. Otherwise, your model will get a distorted view of the past.

Second, like the basic model, the trend model assumes that there is no seasonal pattern in your data. A seasonal pattern is a monthly or quarterly cycle caused by weather, holidays, or other factors. If such a pattern exists, adjust for it before you apply exponential smoothing. (For more on seasonal adjustment, see the March 1986 issue.)

Third, understand that the trend model projects linear growth. This means that the amount of increase (or decrease) is constant each time period. This is the most common pattern, but you may see others. For example, the growth for new products is often exponential. The trend model will produce conservative forecasts that lag behind exponential growth.

Smoothing models are available to handle exponential trends and other types of trend patterns. A reference for the beginning forecaster is my article "Smoothing Methods for Short-Term Planning and Control," in *The Handbook of Forecasting: A Manager's Guide* (second edition, edited by Spyros Makridakis and Steven Wheelwright), forthcoming from John Wiley & Sons, New York. Advanced forecasters may find useful my article "Exponential Smoothing: The State of the Art," *Journal of Forecasting*, Volume 4, Number 1, January–March, 1985. ■