Notes on the random walk model

Robert Nau

Fuqua School of Business, Duke University

1. The random walk model
2. The geometric random walk model
3. More reasons for using the random walk model

1. THE RANDOM WALK MODEL

One of the simplest and yet most important models in time series forecasting is the random walk model. This model assumes tha*t in each period the variable takes a random step away from its previous value, and the steps are independently and identically distributed in size ("i.i.d.").* This is equivalent to saying that the *first difference* of the variable is a series to which the *mean* model should be applied. So, if you begin with a time series that wanders all over the map, but you find that its first difference looks like it is an i.i.d. sequence, then a random walk model is a potentially good candidate.

 A random walk model is said to have "drift" or "no drift" according to whether the distribution of step sizes has a non-zero mean or a zero mean. At period n, the k-step-ahead forecast that the random walk model *without* drift gives for the variable Y is:

$$\hat{Y}_{n+k} \ = \ Y_n$$

In others words*, it predicts that all future values will equal the last observed value.* This doesn't really mean you expect them to all be the same, but just that you think they are equally likely to be higher or lower, and you are staying on the fence as far as point predictions are concerned. If you extrapolate forecasts from the random walk model into the distant future, they will go off on a *horizontal line*, just like the forecasts of the mean model. So, qualitatively the long-term point forecasts of the random walk model look similar to those of the mean model, except that they are always "re-anchored" on the last observed value rather than the mean.of the historical data.

For the random-walk-*with*-drift model, the k-step-ahead forecast from period n is:

$$\hat{Y}_{n+k} \ = \ Y_n \ + \ k\hat{d}$$

where $\hat{d}$ is the estimated drift, i.e., the average increase from one period to the next. So, the long-term forecasts from the random-walk-with-drift model look like a trend line with slope $\hat{d}$, but it is always re-anchored on the last observed value. By comparison, in a simple trend line

model fitted by linear regression, the trend line is firmly anchored at the center of mass of the data [2], and it doesn't move very much if it is re-estimated each time a new data point is observed.

In the random-walk-with-drift model, the estimation of the drift term can be tricky. The obvious way to estimate it is to let it be *the average period-to-period change observed in the past,* which is merely the difference between the first and last values in the series divided by n-1:[3]

$$\hat{d} = \frac{Y_n - Y_1}{n-1}$$

This is the slope of a line drawn between the first and last data points, *not* the slope of a trend line fitted to the data by simple regression. It seems like the logical estimate to use when you consider that using the random-walk-with-drift model is equivalent to using the mean model to predict the first difference of the series. However, as will be discussed in section 2, this method of estimating the drift is very sensitive to the choice of how much historical data to use in fitting the model. In fact, if the drift is very small in comparison to the standard deviation of the step size (as it is for stock prices), the behavior of the random walk process is so irregular that it is very hard to estimate the drift from any finite amount of history, particularly in an era of bubbles and busts, as illustrated by the chart on page 9.

The confidence intervals for forecasts from the random walk model (with or without drift) widen very rapidly as the forecast horizon lengthens, while those of the mean model remain the same and those of the simple trend line model widen only very slowly. In fact:

> ☞ *The confidence interval for a k-period-ahead random walk forecast is wider than that of a 1-period-ahead forecast by a factor of square-root-of-k, which is called the "square root of time" rule.*

Thus the confidence interval for a 4-step-ahead forecast is twice as wide as that of a 1-step-ahead forecast. This is the most important (and most measurable!) property of the random walk model.

Now, why the factor of square-root-of-k? The answer is simple. The distance up or down that is traveled in k random steps is the sum of k independently and identically distributed random variables. The *variance* of the distance traveled in k steps is sum of the variances of the individual steps (because variances are added when you add independent random variables), so the variance of the total distance traveled is k times the variance of one step. The *standard deviation* of the distance traveled in k steps, which is your forecast standard error if you are predicting that the series will stay where it is now, is the square root of the variance, so the k-step-ahead standard error is the 1-step-ahead standard error scaled up by the square-root-of-k rather than by k.

---

[2] In a simple trend line model, the trend line always passes through the point $(\overline{X}, \overline{Y})$, the means of both variables. "Trend" and "drift" both mean the same thing in the sense that they are average increases from one period to the next, but I will use the former term in the context of the simple trend line model and the latter one in the context of the random walk model.

[3] To see this, note that if you add up all the period-to-period changes, you get $(Y_2-Y_1)+(Y_3-Y_2)+\ldots+(Y_n-Y_{n-1})$, and after canceling terms you are left with just Y_n-Y_1. This is the total of n-1 differences, so the average difference is $(Y_n-Y_1)/(n-1)$.

So, the really important parameter of the random walk model is the standard error of the 1-step-ahead forecast. Everything else is determined from this one number by the square-root-of-k rule.

For the *random-walk-with-drift model*, the 1-step-ahead forecast standard error is *the standard deviation of the differenced series.* If the differenced series is called Y_DIFF1 as created in RegressIt, then the 1-step forecast standard error is:

$$SE_{fcst(1)} = STDEV(Y\_DIFF1)$$

For the random walk model *without* drift, the 1-step forecast error is slightly different (although the difference is usually too small to matter): it is the *root-mean-square* of the differenced series, i.e., the square root of the mean of the squared values, because the mean value is assumed to be zero rather than being estimated from the sample. In Excel the RMS value of Y_DIFF1 can be calculated like this:

$$SE_{fcst(1)} = SQRT(VARP(Y\_DIFF1) + AVG(Y\_DIFF1)\char94 2)$$

I.e., it is the square root of the-population-variance-plus-the-square-of-the-mean. This looks slightly messy but it is almost the same as STDEV(Y_DIFF1). For both models, the standard error for a k-step-ahead forecast is the 1-step-ahead forecast scaled up by SQRT(k).

Minor technical point: The critical value of the t distribution that should be used to calculate a *confidence interval* based on the forecast and standard error is not quite the same: for the random walk *without* drift model, the critical t-value is one that is based on n-1 degrees of freedom, where n is the sample size, and for the random walk *with* drift model, it is based on n-2 degrees of freedom. For large sample sizes this difference is unimportant, though: a 95% confidence interval is roughly "plus or minus two standard errors" around the point forecast.

As an example, consider the daily exchange rate between the dollar and the euro, from 1/4/1999 to 8/7/2009 (2669 observations):

Plot of daily USDEuroRate vs Date
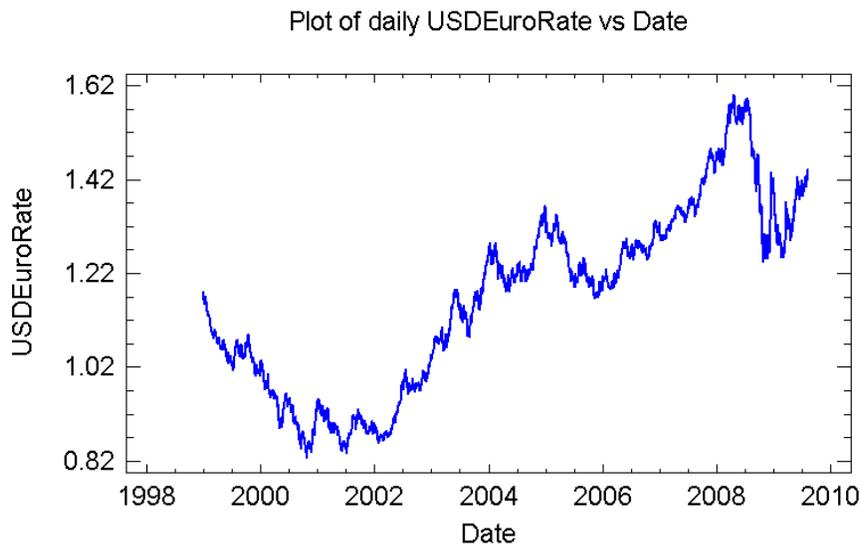
We see lots of short-term volatility, some larger-scale peaks and valleys, and a net upward trend from beginning to end. Obviously we wouldn't fit a simple trend line to this seemingly complicated pattern. What if we look at its first difference (daily changes) instead?


Plot of daily Diff(USDEuroRate) vs Date
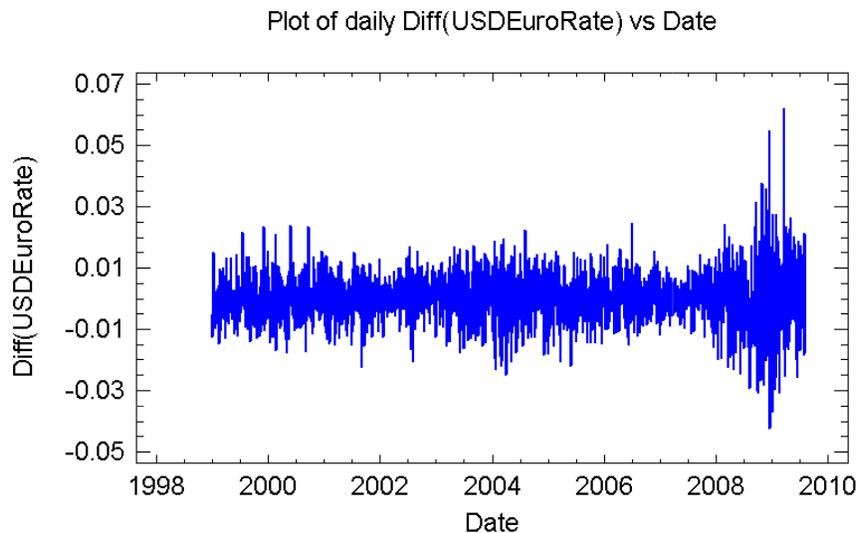
This looks very much like noise, except with intervals of higher and lower daily variance (volatility). The mean of the daily changes is 0.0000887 and the standard deviation is 0.00775 (roughly three-quarters of a cent) over the whole period.

Is the mean of the daily changes significantly different from zero, which is to say, is there a significant day-to-day *drift* over this period? Well, the standard error of the mean is the standard deviation divided by the square root of the sample size, which is $0.00775/SQRT(2669) = 0.000149$. *The t-stat for testing whether the mean is significantly different from zero is the ratio of the mean to its own standard erro*r, which is $0.0000887/0.000149 = 0.595$, which is *not* very significant by the usual t >2 standard.

This simple statistical significance test should not be used to answer the question of whether the exchange rate will continue to drift upward in the future—that's to some extent a policy decision, not a random process—but it does indicate that it won't really matter if we assume the drift to be zero in the short run, so that's what we'll assume for purposes of this example.

Does it appear to be a *random walk* without drift, i.e., are the daily changes *statistically independent* as well having a mean of zero? We can check this by looking at the autocorrelation plot (produced with Statgraphics):

Estimated Autocorrelations for Diff(USDEuroRate)
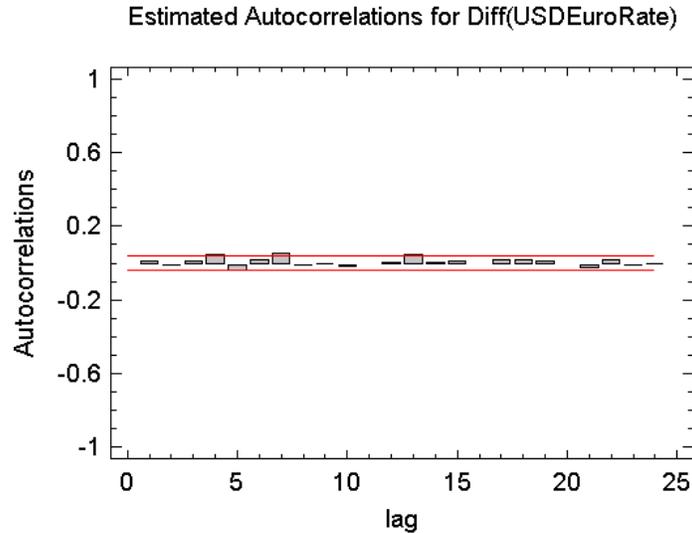
These are *not* significant autocorrelations. The red bands are the 95% limits for testing significance. Nearly all the autocorrelations (particularly the first few) are well within these limits, and there is no systematic pattern. This tells as that, statistically speaking, the USDEuroRate series is an almost perfect random walk without drift, apart from some year-to-year changes in volatility measured in absolute terms. The forecasts and confidence limits for the next 20 business days (4 weeks) look like this:

Random walk forecast for USDEuroRate
(daily volatility based on data from 2008)
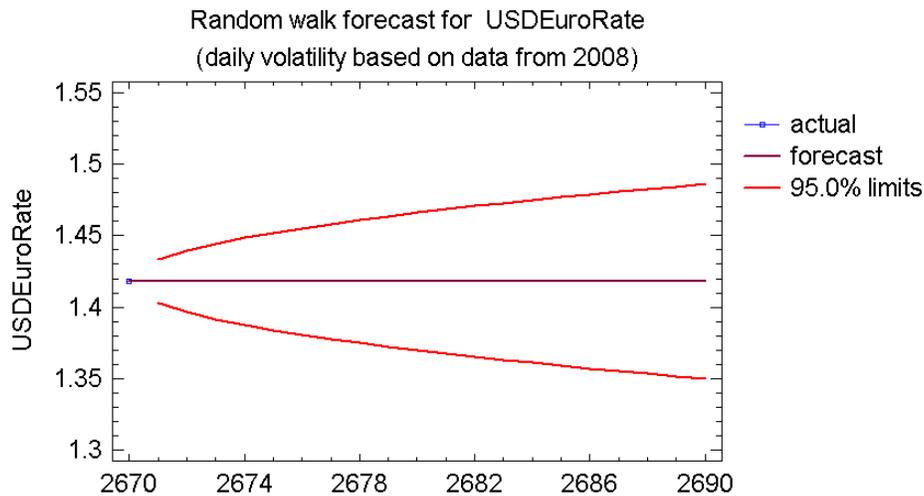
The point forecast remains constant at **1.418**, the last actual value, and the 1-step-ahead forecast error (calculated via the formula two pages back) is **0.00775**, essentially identical to the sample standard deviation of the first difference. The 95% confidence limits for longer-horizon forecasts widen in a "sideways-parabola" fashion as they go farther out, which is the square-root-of-time effect. According to the model, the 95% confidence interval for the rate 4 weeks hence is **[1.350, 1.486]**. Here I would say it is appropriate to report 3 digits after the decimal point, but no more. The 20-step ahead forecast standard error is 0.035, so 0.001 is 1/35 of a standard error. For a 1-step-ahead forecast, whose forecast standard error is 0.0075 as just noted, a precision of 0.001 is about 1/8 of a standard error. The original data is reported to only 3 decimal places of

precision, so it would be strange to claim more precision that that for even a 1-step-ahead forecast.

The daily volatility has been estimated from the whole 10-year sample, which is probably not appropriate in this case. If we restrict the sample to data from 1/1/2008 onward, the estimated volatility is larger, and the confidence limits look like this (on the same scale)

Random walk forecast for USDEuroRate
(daily volatility based on data from 2008)
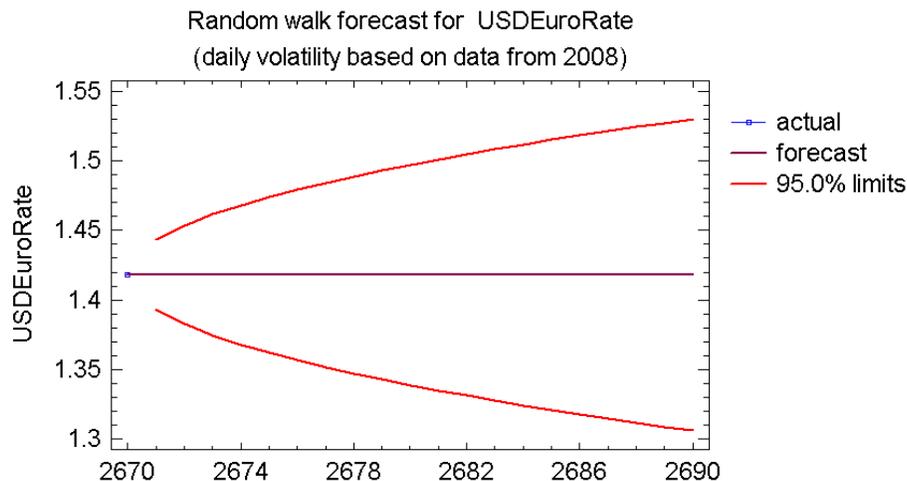
The exact 95% interval for the 4-week-ahead forecast is now **[1.306, 1.530]**, significantly wider, which reflects the greater volatility in the last year-and-a-half. This makes the point that *your results are sensitive to your modeling assumptions, and one of your modeling assumptions is how much past data is considered to be relevant.* It might also be better to measure daily volatility in terms of *percentage* changes rather than absolute changes for this series—we will come back to that later.

What about the interesting "peaks and valleys" we saw in the original data? Aren't these non-random? Well, the patterns did occur "for a reason" *with the benefit of hindsight*, so in a metaphysical sense they are non-random. But the question is whether they occurred for a reason that would have allowed the patterns to be *predicted in advance* based on the historical data. The results of our analysis tell us that they could not be, at least not with a simple statistical model. This is characteristic of financial variables for which speculative markets exist. If short-term ups and downs could be easily predicted with statistical models, they would be arbitraged away by clever traders with big computers. This does happen to some extent on the margin, but not enough to change the big picture unless really huge bets are made by insiders or those with the biggest computers (which has occasionally happened in recent years!).

You will also find that random walks that are *genuinely* random, such as those created by Monte Carlo simulation with tools such as Crystal Ball, almost always have interesting-looking patterns, because your brain tries very hard to find order in chaos. This is an example of a "statistical illusion" such as the hot-hand-in-basketball phenomenon. Players who are perceived to go on hot or cold shooting streaks may be just exhibiting the occasional strings of consecutive hits or misses that you would expect to find in purely random sequences.

2. THE GEOMETRIC RANDOM WALK MODEL

In the *geometric* random walk model, the *natural logarithm* of the variable is assumed to walk a random walk, usually a random walk with drift. That is, the changes in the natural log from one period to the next, which are approximately the *percentage changes,* as we all know by now, are assumed to be independent and identically normally distributed. The k-step-ahead forecasting equation of the geometric random walk model is therefore the same as the one for the random-walk-with-drift model, except that it is applied to LN(Y) rather than Y:

$$LN(\hat{Y}_{n+k}) \ = \ LN(Y_n) + kr$$

where r is the drift measured *in log units*, which is almost the same as the percentage increase per period.[4]   In other words, this is basically the same as predicting that the series will undergo compound growth with a growth factor of (1+r) per period, i.e.,

$$\hat{Y}_{n+k} \ = \ Y_n (1+r)^k$$

For example, if the estimated drift in log units is $\hat{r} = 0.014$, this corresponds almost-exactly to a growth rate of 1.4% per period, which is a per-period compound growth factor of 1.014. [5]   If the first difference of the logged series is called Y_LN_DIFF1 as created in RegressIt, then the 1-step forecast standard error for the geometric random walk model *in log units* is:

$$SE_{fcst(1)} = STDEV(Y\_LN\_DIFF1)$$

The k-step-ahead forecast standard error is obtained by scaling it up by a factor of SQRT(k). Confidence intervals for the forecasts in log units are calculated in the usual way as the point forecasts plus-or-minus an appropriate number of standard errors. Finally, point forecasts and confidence limits for the series in its *original units* are obtained by "unlogging" the logged ones by applying the EXP function to them. Statgraphics will do all these calculations automatically for you, and it is not too hard to do them yourself in Excel.

To illustrate the application of the geometric random walk on a geological time scale, consider the monthly closing values of the S&P500 index January 1950 to July 2009 (714 observations).

---

[4] I'll use the symbol r rather than d to denote drift when it is measured in log units, because that is the case in which it is interpretable as a percentage growth rate.

[5] To see this correspondence, note that if you take the natural log of both sides of the second equation, you get $LN(\hat{Y}_{n+k}) \ = \ LN(Y_n (1+r)^k) \ = \ LN(Y_n) + k\,LN(1+r) \ \approx LN(Y_n) + kr$ because $LN(1+r) \approx r$ as discussed in the last section. The first equation (in terms of natural logs) should thought of as the "true" equation, and the second (in terms of compound percentage growth) should be thought of as a non-geeky approximation.

SP500monthly
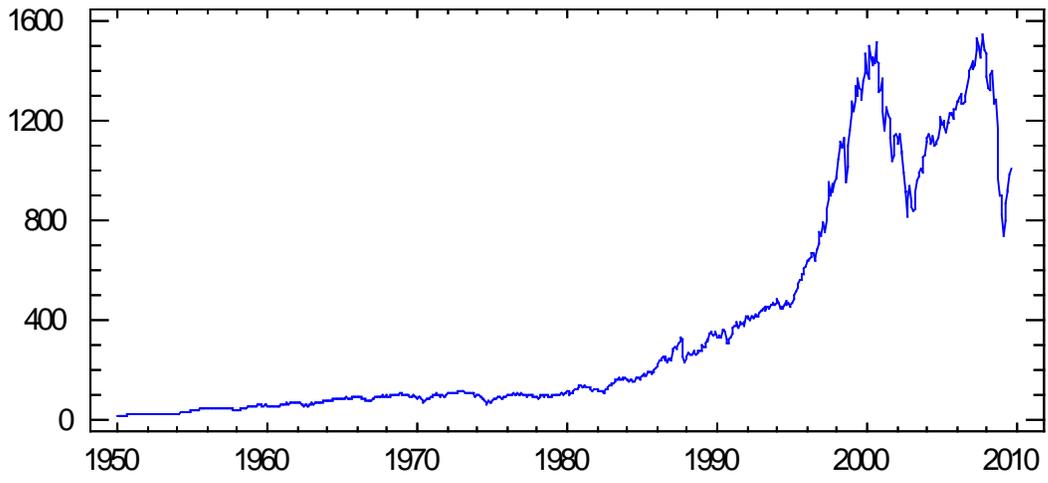
Overall we see what looks like a pattern of general exponential growth, with occasional dips and blips, leading up to the huge dot-com bubble and housing bubble within the last 15 years. It's hard to look at this picture and say that recent years look even remotely similar to the earlier ones. It's also hard to look at this picture and eyeball a long-term growth rate.

If we plot the monthly *change* in the index (i.e., the first difference), we get the following picture:

Diff(SP500monthly)
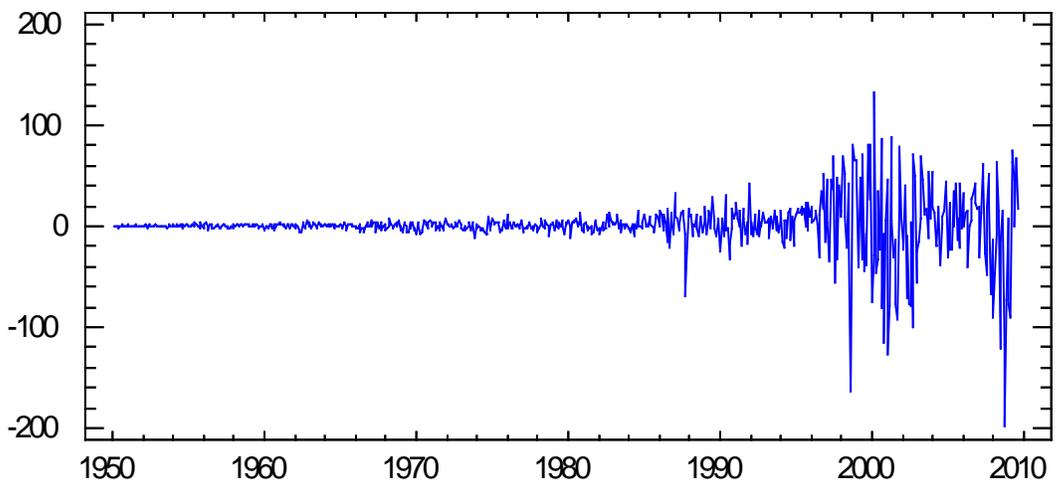
Here too we don't see a pattern that looks stable over time: the monthly changes are obviously a lot bigger toward the end of the series, where the absolute magnitude is also a lot bigger.

Now let's plot the *natural log* of the monthly index:

Log(SP500monthly)
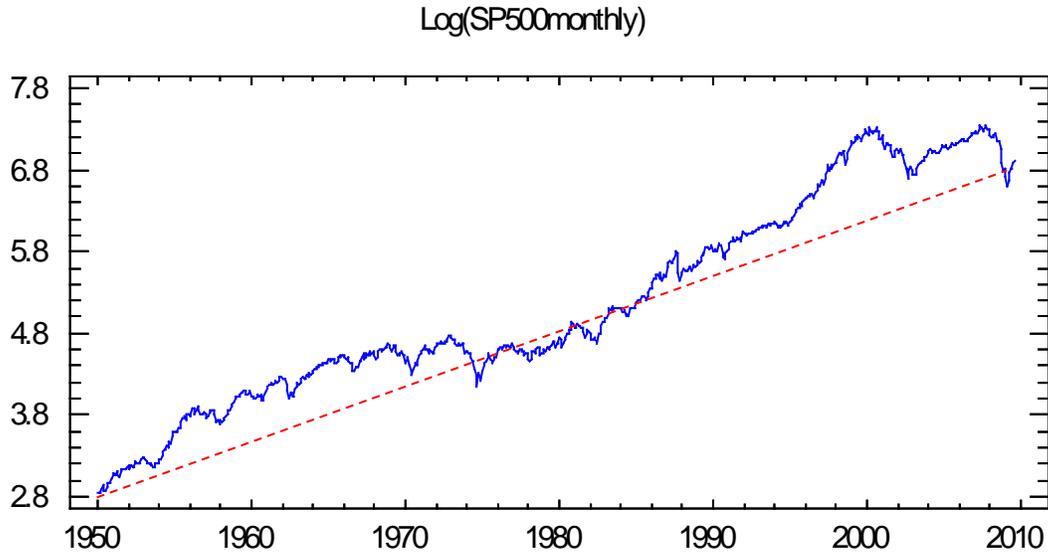
It's a lot easier to make sense of this picture. Recall that the natural log transformation turns an exponential growth pattern into a linear growth pattern, and the slope of a trend line fitted to natural-logged data can be interpreted as a percentage growth rate. Just for fun, I have hand-drawn a dotted line between the first and last value. In very rough terms, this line captures the average linear trend in the logged index. As you can see from the Y-axis scale, its height rises by about 4 natural-log units (from 2.8 to 6.8) over a time span of almost 60 years (714 months). The slope of this line is therefore $4/714 = 0.0056$, which (voila!) can be interpreted as an average monthly increase in percentage terms. That is, the average monthly increase of the S&P 500 index since 1950 has been about 0.56%, which translates into an average annual increase of 6.9%.[6]

You can play this game with any starting and ending points, of course. Pick any two points, draw a line between them, compute the slope, and that tells you the average monthly percentage growth over that interval of time. Clearly there have been good decades and bad decades: the 50's were a high-growth period, the 60's were a normal-growth period, the 70's were a stagnant period, the 80's and early 90's were a period of better-than average growth, and a roller-coaster ride began around 1995. There are also sharp dips and recoveries interspersed throughout the series.

So, there are a lot of interesting stories to go along with this picture, but what can we say about it in statistical terms? Here is a plot of the *first difference of the logged series*:

---

[6] The annual growth rate is computed by compounding the monthly rate for 12 periods: $(1.0056)^{12} = 1.069$.

Diff(Log(SP500monthly))
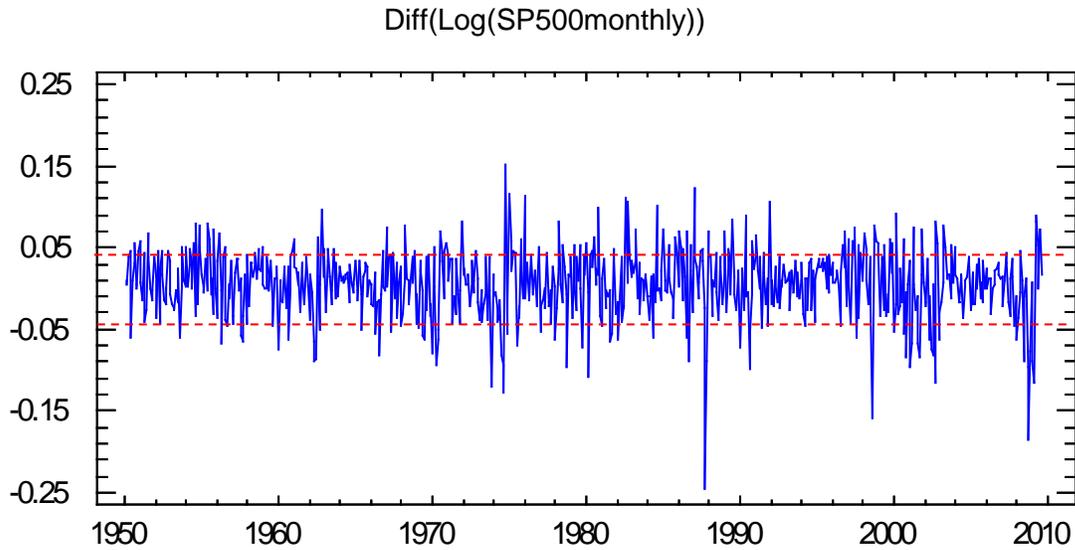
The diff-logs are interpretable as (approximate) percentage changes, so what we see here is a pretty steady stream of monthly changes on the order of +/- 5%. In fact, the mean change is **0.0057** and the standard deviation is **0.042**. (I have hand-drawn some reference lines at roughly +/- 4.2%.)   So, *from one month to the next since 1950, the index has grown on average by 0.57% per month with a standard deviation of 4.2%.* We do see what appear to be periods of somewhat higher and lower volatility, and occasional months of sheer panic, but on the whole the pattern is relatively consistent over this very long stretch of time.

Apart from the modest departures from constant percentage growth and constant monthly volatility noted in the graphs of the logged and diff-logged data, are the monthly percentages *random*? Let's look at their autocorrelations:

Estimated Autocorrelations for Diff(Log(SP500monthly))
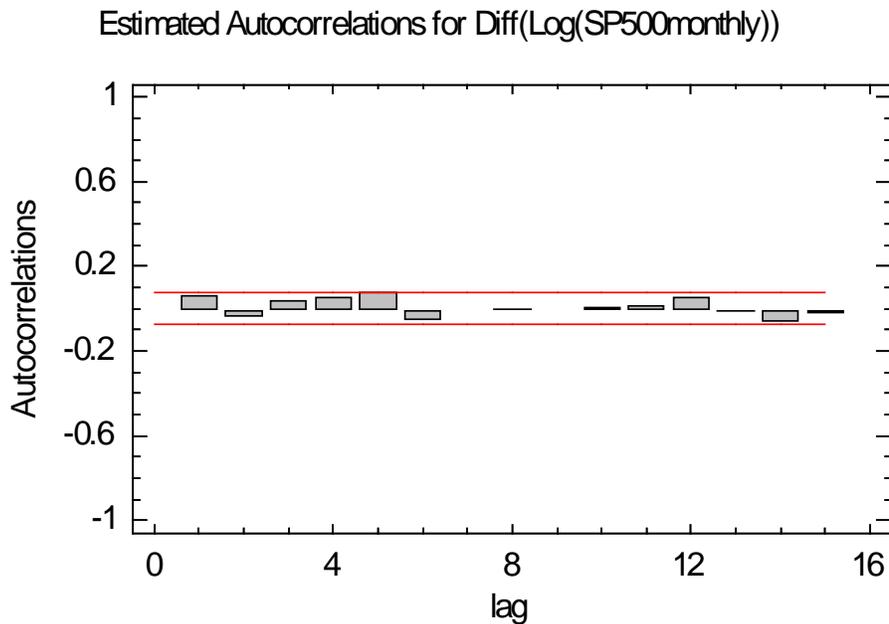
There is no significant pattern here, so the monthly changes appear to be statistically independent.   as well as almost-identically distributed. The S&P 500 monthly closing value therefore looks like an almost perfect random walk.

One more question we might want to ask is whether the monthly percentage changes are *normally distributed* or not. Confidence intervals for forecasts from the random walk model are based on the assumption that the steps are normally distributed as well as i.i.d. We can test the hypothesis of normality by drawing a *normal probability plot* of the diff-logged series:

Normal Probability Plot
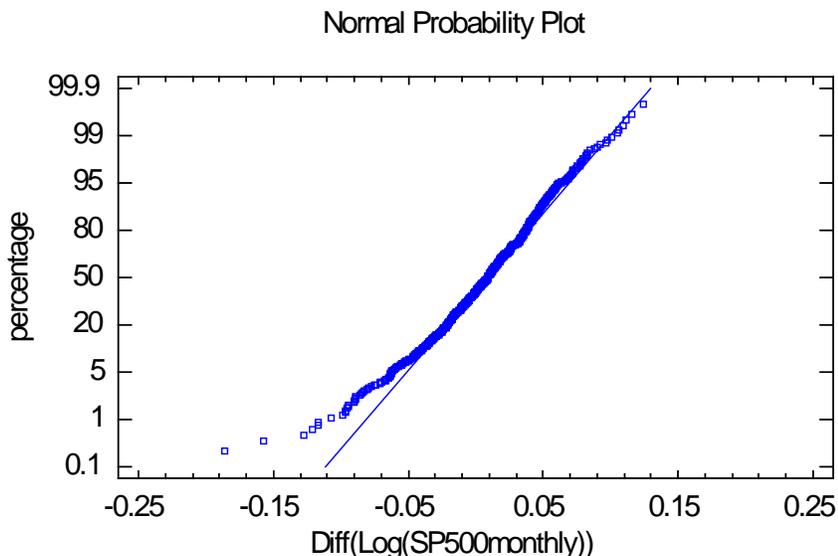
Recall that a normal probability plot is a plot of the values versus the percentiles of a normal distribution having the same mean and standard deviation. If the sample data is normally distributed, the points should lie along the straight line. Here the fact that the plotted points bend to the left at the bottom of the plot means that the distribution is "skewed" to the left, i.e., there are more big values in the lower tail of the distribution than there should be if the distribution is normal. In particular, there was a 24.5% drop in October 1987 and an 18.6% drop in October 2008.[7] The sample standard deviation of the diff-logged values, (which is the forecast standard error of the random-walk-with-drift model) is 4.2% as noted above. So, we have one negative outlier that is more than 5 standard deviations and another that is more than 4 standard deviations. Here are some handy facts:

- A 4-standard-deviation outlier (either positive or negative) occurs less than 1 out of 10,000 times in a normal distribution.

- A 5-standard-deviation outlier occurs less than 1 out of a million times.[8]

Here we have 714 observations, so there is no way we should have observed such large outliers, particularly two on the same side. These two huge outliers by themselves are responsible for

_____

[7] A technical note about normal probability plots: the left and right endpoints (i.e. biggest negative and positive outliers) are not actually shown. They serve to anchor the scale on the left and right, but their points are not plotted. Here, for example, the most negative value plotted is the 18.6% drop of October 2008.

[8] More precisely, a "4-sigma" outlier is a 1-in-16 thousand event, and a 5-sigma outlier is a 1-in-1.7 million event.

most of the left-skewness in the distribution. So apart from a meltdown every 20 years or so, the distribution of monthly percentage changes is not too far away from being normal.

The steps in the analysis shown above follow a logical progression of the sort that you should follow with any sort of time series data. We began by looking a time-series plot of the data as well as its first difference. (The first-difference plot is only relevant for variables that are time series.) We saw a pattern of very non-linear growth and very non-constant volatility in those plots, so we were led to consider the possibility of a transformation of the data. The log transformation is well known to straighten out exponential growth curves and to stabilize the variance of their first differences, and in particular, the log transformation is a standard analytical tool in finance. So we then looked at a time series plot of the logged data and its first difference, and what we saw there was an approximately-linear trend over the very long run and approximately-constant monthly volatility over the very long run. Because we had used the *natural* log transformation, the long-term trend and the monthly changes were directly interpretable in percentage terms.

These observations lead us to the conclusion that it would be reasonable to use a random-walk-with-drift model to predict the *logged* S&P 500 index. With the growth rate and volatility estimated above, we obtain the following picture:
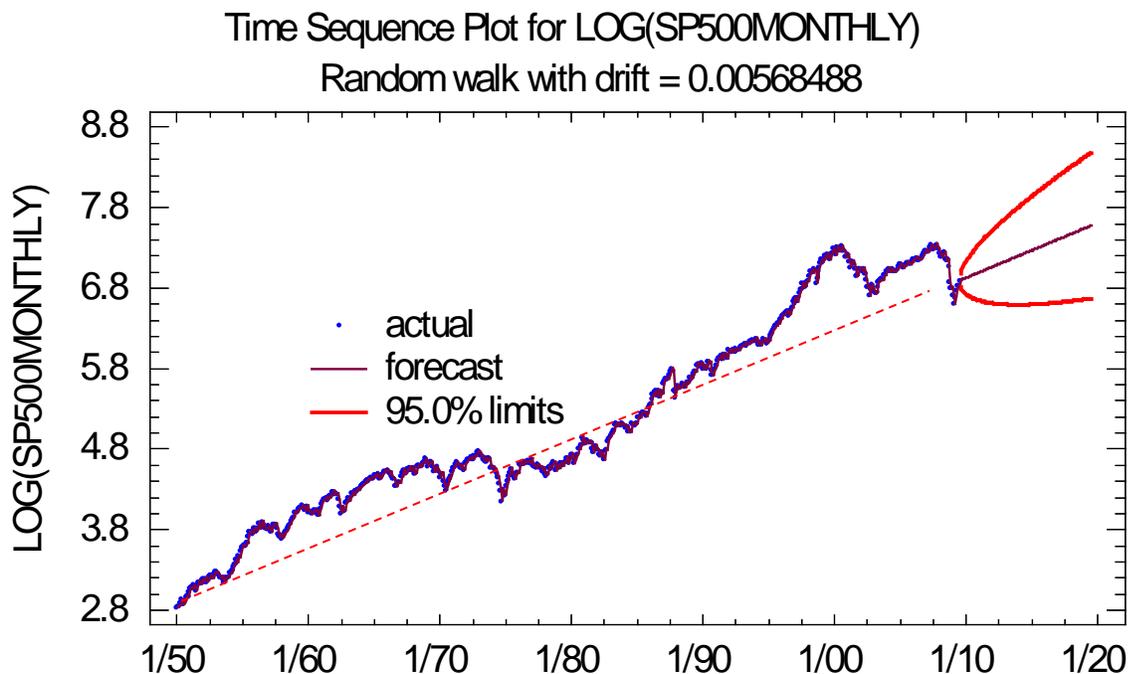
This plot was generated by using the [User-Specified] Forecasting procedure in Statgraphics. The series Log(SP500monthly) was used as the input variable and the random-walk-with drift model was specified, and 120 forecasts (10 years of monthly values) were produced.

The dashed red line was *not* drawn by Statgraphics. I have hand-drawn it myself (the same as in the earlier picture) to make the following point: when the trend in the random-walk-with-drift model is estimated by calculating the average period-to-period increase that was observed in the

historical sample of data (as is the case here), *the point forecasts for the future are actually just an extrapolation of a straight line drawn between the first and last data points!* Knowing this, you can guess how the forecasts for the future would have looked different if we had used a different starting point for the historical sample. For example, if we had used a 10-year sample beginning in late 1999, at the height of the dot-com bubble, then a line extrapolated through the first and last data points would have trended downward, not upward. This shows that it *is very hard to estimate the trend in a random-walk-with-drift model based on the mean growth that was observed in the sample of data unless the sample size is very large.*

When fitting random-walk-with-drift models, it is usually necessary to draw on other sources of information to estimate the trend appropriately. Here it would make no sense to predict that the index will decline in the future, based only on a very simple time series model fitted to a sample of historical data. If everyone believed that, it would already have gone down! Logically the growth rate of a stock index ought to be at least as great as the risk-free rate of interest, whatever that may be at the time. How much greater depends on the "market price of risk" under current economic conditions. Also, *the growth rate that should be predicted for the prices of individual stocks and mutual funds depends on the degree to which their movements are correlated with those of the market index,* a phenomenon explained by the so-called "capital asset pricing model" (CAPM). These are mysteries beyond the scope of this course—ask your finance professor!

The confidence intervals for the long-horizon forecasts widen in the sideways-parabola shape that is characteristic of the random walk model. The width of the parabola depends on the standard error of the 1-month-ahead forecast, which is the estimated monthly volatility. This parameter has also been estimated from the entire data set, under an implicit assumption that volatility is constant in the long run and that volatility in the near future will be "average." This is also too simplistic. Short-term volatility can vary—sometimes very rapidly—and you really ought to draw on other sources of information to estimate it appropriately. Current values of *stock options* can be used to back out the "implied volatility" of the market index as of right now, using the famous Black-Scholes formula. This is also a topic beyond the scope of this course.

Here is what was specified on the input panel of the Forecasting procedure in Statgraphics[9] in order to get the results pictured above:

---

[9] This procedure was designed by your humble instructor many years ago.
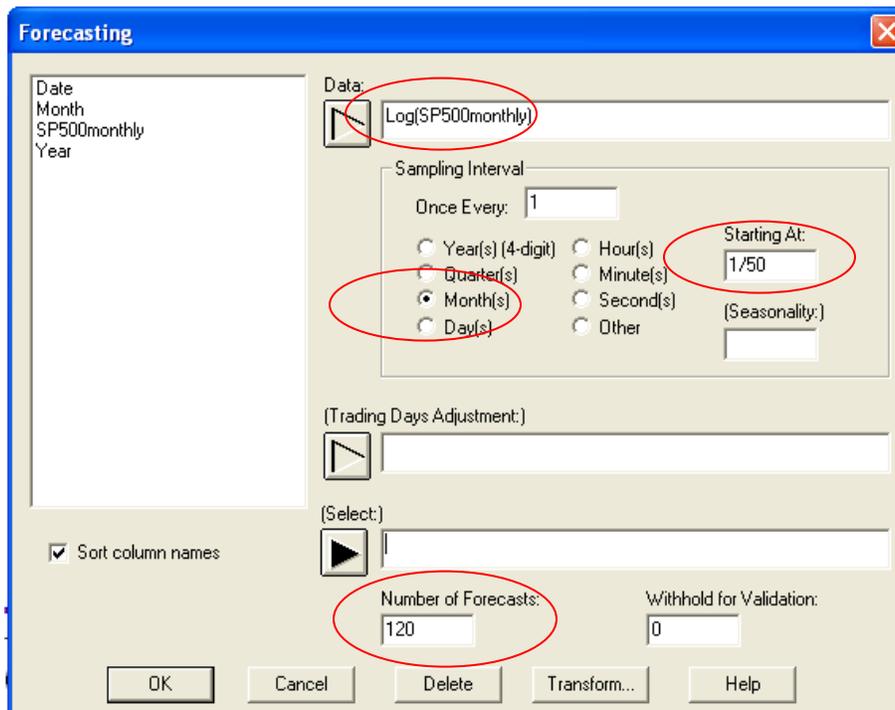
The key things that were done here were: (i) the (natural) log transformation was manually applied to the input variable[10], (ii) the data was specified to be monthly, (iii) the starting date was specified to be 1/50 (this has no bearing on the analysis—it is just used to label the time axis on the plots),. and (iv) 120 forecasts (10 years's worth) were requested.

Here are the right-mouse-button "Analysis Options" that were used inside the procedure in order to specify the model:

---

[10] Recall that in Statgraphics the natural log function is LOG rather than LN. Function names and variable names are *not case-sensitive* in Statgraphics, so it would have been equivalent to type LOG(SP500MONTHLY) here..
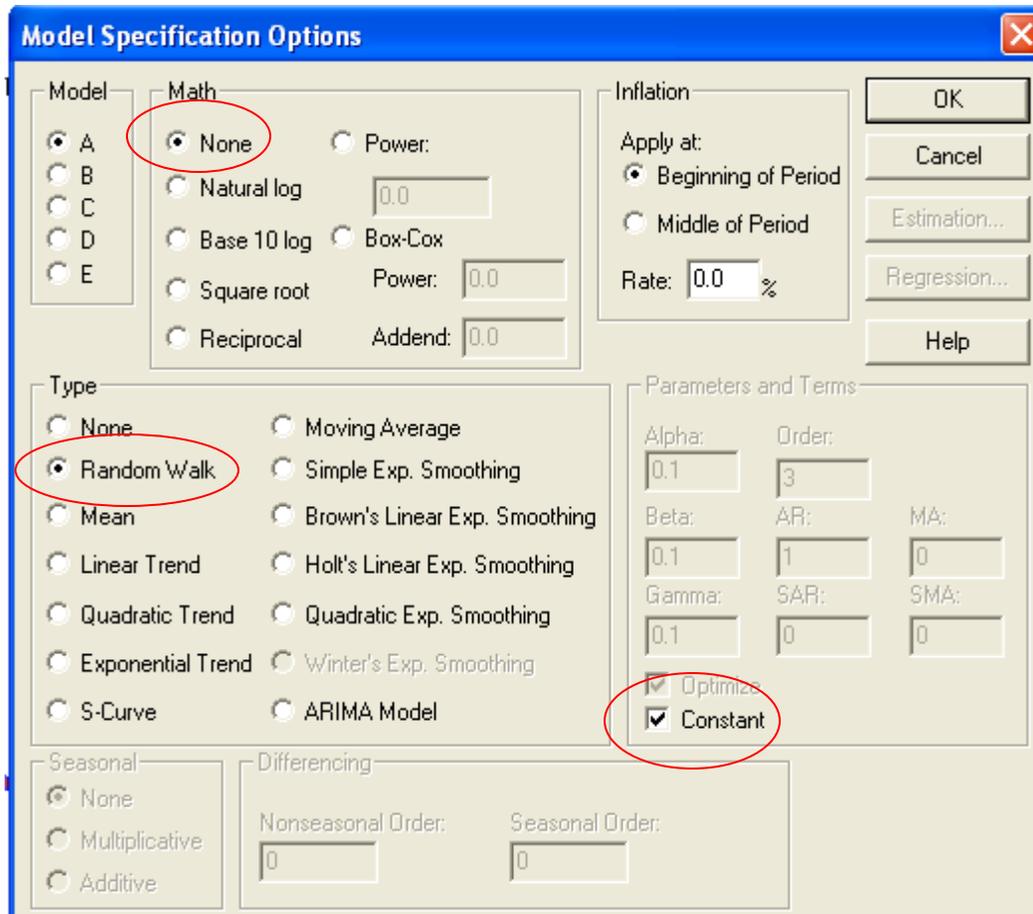
The Random Walk model type was selected with no additional math transformation (because the log transformation was already applied to the input variable), and the "Constant" box was checked to indicate that a constant growth rate should be estimated, i.e., that it is a random walk with *drift.* This model has been specified as "Model A." You can specify up to 5 different models, and the output of this procedure contains a "Model Comparison Report" that shows all all their key stats and forecasts side-by-side.

Fitting a random-walk-with-drift model to the logged series is equivalent to fitting the geometric random walk model to the original series. The fitting of this model that was shown above was carried out in log units, by manually logging the variable on the input panel of the Forecasting procedure. Notice that in log units the plot of the forecasts and their confidence limits looks quite reasonable—almost conservative—given the historical pattern of the data.

Alternatively, you can fit the geometric random walk model by specifying the log transformation as one of the model options *inside* the procedure. If you do it this way, the forecasts and confidence intervals will then be plotted in their original "unlogged" units in the output. Exactly the same model will have been fitted, but the forecasts and confidence intervals will have been unlogged at the very end by applying the EXP function, which yields the following picture in the original units of the S&P 500 index:

Time Sequence Plot for SP500monthly
Random walk with drift = 0.00568488
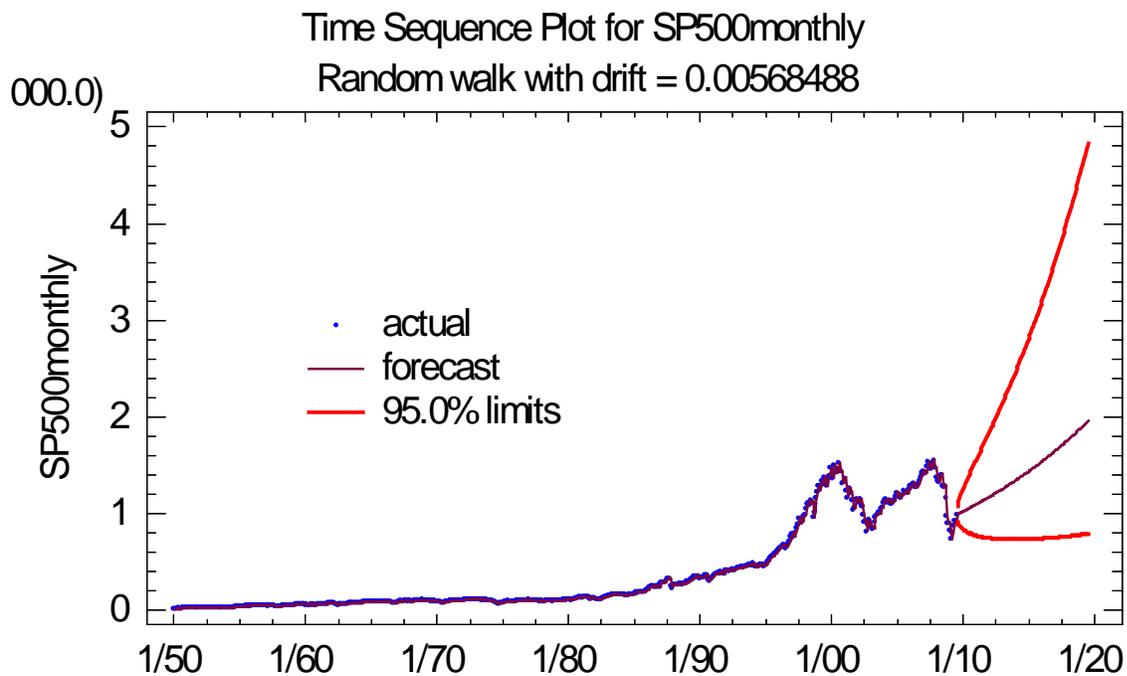
Now *this* looks like a rather dramatic forecast of what might happen over the next 10 years, both on the high and low sides, but it is logically the same as the previous forecast plot. Only the vertical axis scale has been changed. Before it was in log units, here it is in original units.

The settings for this model in Statgraphics are the same as for the previous one except for one small difference. The Log function was *not* applied directly to the input variable, i.e., the input variable was SP500monthly rather than Log(SP500monthly):
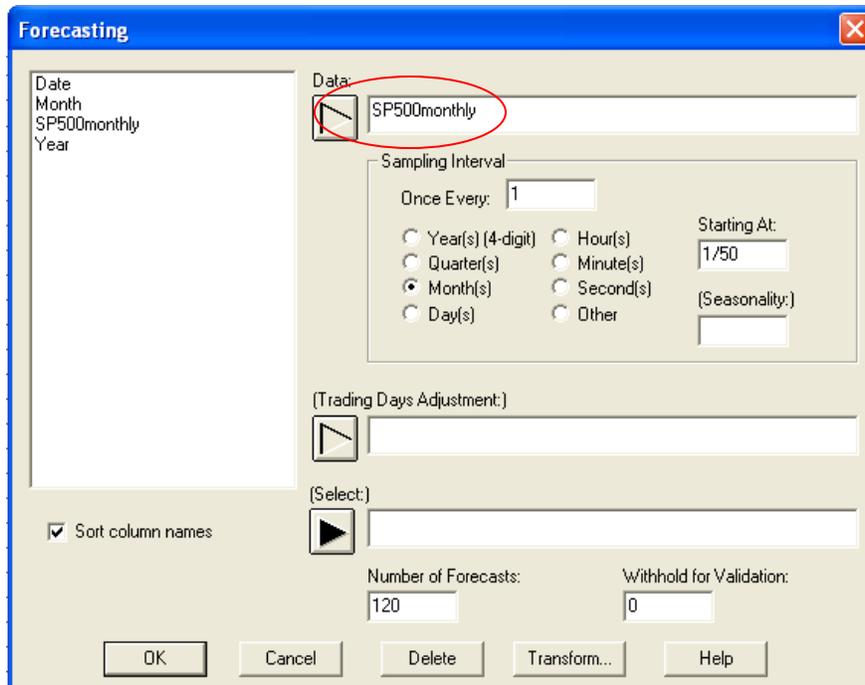
Then, on the right-mouse-button Analysis Options panel, the natural log transformation was specified as a "Math" transformation.
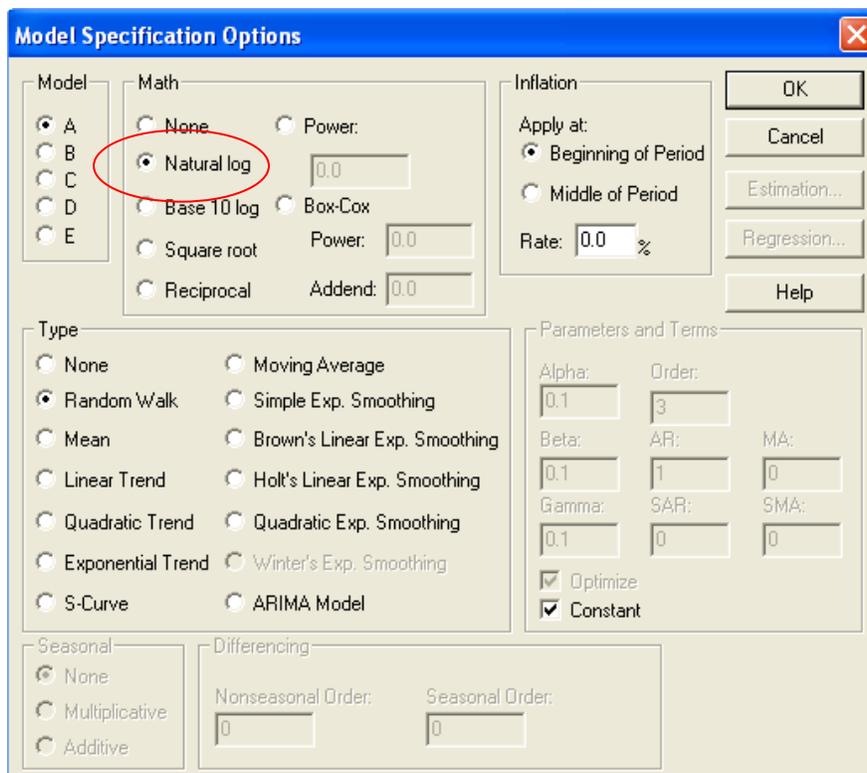
When a transformation is applied inside the procedure like this, what happens is the following: the math transformation is first applied to the input variable, then the forecasting model is

applied to the transformed variable, and at the end the forecasts and confidence intervals are un-transformed to get back to original units before they are plotted.

3. MORE REASONS FOR USING THE RANDOM WALK MODEL

The examples used in this chapter to illustrate the random walk and geometric random walk models were based on data from financial markets.   Finance is a very important application area, because there is a very good theoretical reason for believing that prices of assets for which speculative markets exist ought to behave like random walks.   Namely, if it were easy to predict whether the market would go up or down tomorrow, it should have already gone that way today.   However, the random walk model also crops up in many non-financial applications, and you should be on the lookout for situations in which it is the best you can do, regardless of what you may have hoped or expected.   If you see what looks like pure noise (i.i.d. variations) after performing a 1st-difference or diff-log transformation, then your data is telling you that you that it is a random walk. This isn't very exciting in terms of the point forecasts you should make ("next month will be the same as last month, plus average growth"), but *it has very important implications in terms of how much uncertainty there is in forecasting more than one period ahead.*

Some of the more sophisticated time series that we will study later are merely "fine-tuned" versions of the random walk model, in which lagged values of the dependent variable and/or lagged values of the forecast errors are added to the prediction equation in order to account for the fact that the steps may not be statistically independent—they may have patterns of autocorrelation that can be exploited to improve the forecast.

Another very important use of random walk models is as *a benchmark against which to compare more complicated time series models, particularly regression models*.   There is a tendency in regression analysis to think of R-squared as a measure of the predictive power of the independent variables in the model.   This may not always be the case! R-squared measures the amount by which the error variance of the regression model is lower than that of the *mean* model for purposes of predicting the dependent variable.   But if the mean model is not an appropriate reference point, this is a meaningless statistic. If the time series is nonstationary, a random walk model will often have a much lower error variance than the mean model, and a better question would then be: how much lower is the error variance of the regression model than the error variance of the random walk model? A poorly chosen regression model may actually do worse than a random walk model for predicting a time series, even if it has a big R-squared.

So, the random walk model is an important tool in your toolkit regardless of whether it is the model you end up using for your forecasts at the end of the day.

More wisdom from *The Profit*:

A woman asked,
Which way does the Wind Blow?

And he said:
The Wind is invisible but strong, like vodka...
The Wind is a friend.
The Wind is an enemy.
The Wind is neutral, like Switzerland...
Sometimes the Wind will blow down a tree.
Other times it won't.

There's just no telling with the Wind.

Which way will the Wind blow tomorrow, the woman persisted.

He answered:
The Wind is free and unpredictable, like a bird. It is impossible to guess which way it will blow.
She who would ask such a question would be a fool, and he who would attempt to answer it, a greater fool.

I understand, the woman said, but which way will it blow?

North, he replied.