

Explore What Did and May Happen with Time Series Questions

By Richard G. Lamb, PE, CPA, ICBB; Analytics4Strategy.com

We make many decisions upon information which are reported sequentially at fixed periods—time series. The possibilities are productivity, KPIs, volumes, costs and many others.

We hope to draw meaning from what happened and, in turn, either reconfirm or reset our expectations for what will happen. It follows that our ability to understand what happened and direct what will happen is fundamental to operational breakthroughs.

But good luck with that! Good luck with that if we cannot see clearly any outliers and volatility in the series. Good luck with that if we cannot see clearly the cycles due to season or calendar. Good luck with that if we cannot see clearly the shape and drivers of the series after removing the cycles. Good luck with that if we cannot see clearly the true spread of the series.

We must build a window into the many serial behaviors along operational processes. The insight is one of the [five core categories of questioning](#)—relationship, difference, time series, duration and apparency. Otherwise, we will not be able to break much from the past.

This article will explain time series analytics in the context of the questions process experts and operatives would ask and answer of a process. Consequently, the article is not intended to be an explanation of statistical method beyond what is necessary to explain the questioning.

The explanation we all want in the end is how to execute. The article, "[DMAIC Done the New-Age Way](#)," explains how the time series questioning of this article, along with the other four questions, are woven into the stages to define, measure, analyze, improve and control processes. Although presented in the context of DMAIC, the explanation is universal to any problem solving sequence.

A Primer on Time Series

A primer on time series must cover three core aspects. They are autocorrelation, components and distortions.

Autocorrelation: The central principle of time series data is what is called autocorrelation. In plain English this means that what happens in one period may be felt significantly by the subsequent one or more periods. In statistics-speak, the observations are not independent of each other.

Figure 1 shows the concept. The upper frame is a time series of the type we always see. In this case a KPI is being reported daily.

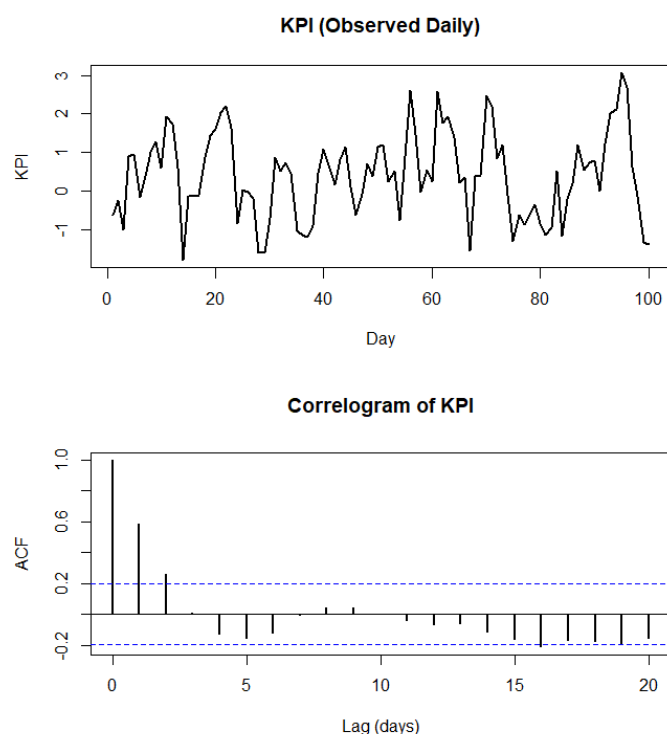


Figure 1: A KPI series with each period influenced by two preceding periods.

The lower frame is what is called a “correlogram” of the series. It is a key tool in series analysis. The acronym “ACF stands for autocorrelation function. The factor at zero lags is always “1,” but placed in the display to give contrasts to all other lags.

With the correlogram, we can see if what happened in one period is felt in subsequent periods. The spikes at lags 1 and 2 are shown as being significantly influenced. This is indicated because they extend beyond the confidence interval of independence (dotted line).

The correlogram shows that a KPI is strongly influenced or explained by the previous period and much less by the second previous period. The degree that a current KPI is explained by the first and second previous periods is the sum of each squared ACF.

If an organization does not recognize this, it will be misled by its reports. They will regard each KPI report as reflecting only what happened on the subject day.

There are other ramifications of autocorrelation. One is that the normally reported variance of a series will be understated (better than we think) when there is autocorrelation in the series. Another is that autocorrelation is applied to evaluate the fit of models to the series. Another is that the mathematics of autocorrelation—called cross-correlation—can be modified to reveal and confirm useful lead-lag relationship between operational variables.

Components: A series has components. We must parse them out as series in their own right. First are the seasonal and calendar cycles. The second, called “trend or level,” is what remains after removing the cycles.

However, beware! We are tempted to think of trend as a fitted straight line. Instead, think of it as what remains after removing the cyclic component from the series.

The cycle and trend components are separated by fitting a model to the series. The same model allows us to understand the characteristics of the series and its component series. We also use the fit to foresee what the future for the subject series may or can look like.

Distortions: The truth of the fit can, unbeknownst to us, be undermined by distortions lurking like trolls in the series. There are two types to ferret out and adjust for. Both are often invisible to the naked eye. It follows that finding them can point to a breakthrough to operational excellence.

First of the distortions are outliers to the fitted model. Second are fluctuations in the range of variance across the series. Both are unearthed by applying the principle of the autocorrelation function to the difference between the fitted and actual values—in statistics-speak, called “residuals.”

If you could ask and answer questions around these principles, try to imagine all you would explore and discover. The subsequent sections will group the questions we would find ourselves asking. That is if we bring the insight of time series into the operational excellence project. They are questions for variable selection, distortion, cycle, trend, sigma and forecast. However, first it is necessary to introduce the analytic tools through which questioning is made possible.

The Analytic Tools

The models to ask and answer time series questions are available from top-tier software. In fact, all of the analytics shown in this article are conducted with the free software “R” (<https://www.r-project.org>).

All models are tested for fit. The test is to compare the modeled series to the actual series and then evaluate the resulting residuals for autocorrelation. If lags are the case, as seen in a correlogram such as Figure 1, we know that a more powerful model is needed to acceptably fit the cycles and trend of the series.

The most straightforward models are the decompose and the Holt-Winters models. If the decompose and Holt-Winter models do not pass the test, we may go to a linear regression model. The elements of the regression models are coefficients for the trend and the cycle at each of its periods. Another method is to fit the cycles with a sine wave, called harmonic models.

However, linear regression is only applicable if we are working with cycles and trend that have drivers. In statistics-speak the series is “deterministic.”

A series may not have drivers. In statistics-speak, the series is “stochastic.” Random series, random walk series and drifting random walk series are such cases.

Actually, any series will be one of a possible permutation of with and without cycles, deterministic and stochastic components, and periods of lag and moving average for each component. As the series is more complex, we move to an autoregressive moving average (ARMA) model or beyond to an autoregressive integrated moving average (ARIMA) model.

“AR” stands for autoregressive and “MA” stands for moving average. The “I” deals with the stochastic and deterministic aspects of the series. However, the article will not go into a discourse to explain them.

Amongst the options for the decompose, Holt-Winter, linear regression, ARMA and ARIMA models, we find the best fit by trial and error. In other words, we are using the comparative fit of models to determine the underlying characteristics permutation of the subject series. This may be the most important implication of times series questioning—knowing what we are dealing with.

Before accepting a final model, we should search for distortions that are significant to the fit and the truth of the model. As mentioned earlier, there are two types—outliers and stability of variance.

Analytical software provides tools to ferret out five types of outliers to be presented later in the article. They are unearthed by testing all records in a series as outliers to the fitted model. This is compared to outliers to dataset to the model. The [tsoutlier function](#) of “R” can be applied to conduct the search and assessment of outliers. However, all top-tier software include equivalent functions. Ultimately, we make decisions to delete, ignore or adjust each.

A series may be unstable with respect to the variance. When fit to a model, the generated measures of autocorrelation may show the residuals to be random with time, thus, a good fit. However, when the autocorrelation is squared, a picture of instability may emerge. When this appears, we move to the generalized autoregressive conditional heteroskedastic (GARCH) model. The model allows for changes in variance along the series.

Which Variables to Explore?

The first stage of questioning is directed at identifying the variables of interest. Many variables of a process can be extracted in datasets that include period-type variables for each record.

Consequently, many of the other variables in in the dataset can be explored through the questioning of time series.

The questions are asked in the context of subject and surrounding operational processes. From which elements of the process will insight into the past tell a story for the operational process as a whole? For which must we have insight because they are the proverbial canary or foretell future outcomes of the process? What will we do with the insight into the components and distortions to them? These are just a few of the universe of possibilities.

Are There Distortions and Why?

The next range of questions orbit around whether there are distortions in the data and why. More specifically are there outliers to the fitted model and shifts in the variance of residuals? If so, are they significant? Do they reveal immediate and longer-term actionable opportunities to change the future?

The first question of distortion is answered by unearthing any of five types outliers to a seemingly well-fitting model and then rerunning the model without those that are significant. The second question of distortion is dealt with by moving to a different class of model (GARCH) if testing shows instability of variance.

Figure 2 is a representative graphical output from the analysis for outliers using the `tsoutlier` function of “R.” They have the jargon-based names of additive, innovative, temporary, level shift and seasonal level shift outliers.




Figure 2: The five types of outliers to a series model.

The article will not get into defining each beyond the largely self-explanatory graphical representations. The point is that we cannot be comfortable that we would spot them just from inspecting the traditional line or column chart of the series. The analytics of packages, such as `tsoutliers` in R, will tell us where the cases exist and if they are significant to the fit and truth of the model.

Once revealed, we have a host of questions. Are any of the outliers due to error in data collection—fat finger error? Which are due to noncompliance with the rules of conduct for the process? Of those that remain, what are we finding in the cases that we did not know before about the process? What is the opportunity of the discoveries with respect to the improving the subject and surrounding operational processes?

Thence, we have questions for whether or not to remove the distorting outliers from the series for modeling. The answers will rest on other questions. Which outliers have a significant effect on the fit and truth of the model? If due to noncompliance and other process issues, can we fix them? Are they, instead, a reality that must remain reflected in the final model?

A fundamental underlying principal to time series analytics is that the variance is constant without regard to time period. In line with the principle, can you see statistically significant change in the vertical range of Figure 3?




Figure 3: Can you see changing ranges in the variance of the series?

Without analytical help, a team could debate forever if instability is present and still not reach a consensus. In fact, using analytic tests, we would discover that there is volatility in the series.

Now we have new questions. Why the shifts? Do the shifts undermine the process? Can we or should we find a way to prevent or reduce the shifts? Should we redesign the process to realign itself at each shift—if we cannot beat them, join them?

Are There Cycles?

From the questioning of the previous stage and with cleansed data we now go to the questions we can ask once we parse the components by modeling.

Figure 4 is an example of parsing a series using the Holt-Winters method. The top frame is the series. It is also what organizations have traditionally been limited to working with and being forced to ponder upon intuition rather than insight.

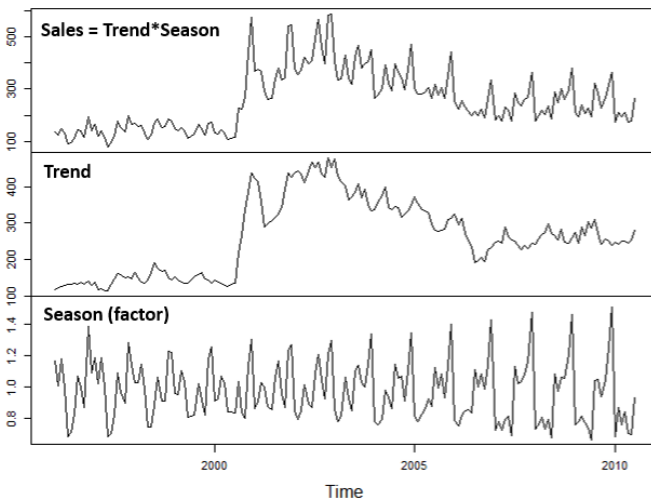


Figure 4: A series and its constituent components.

Of course, the first question is whether or not the series actually includes a cycle. If as in the lower frame of Figure 4 there is a cycle, is it seasonal, calendar or both? What is the pattern of the cycle? Is it deterministic or stochastic?

Some of the questions are answered in the bottom frame—season. Others may require the ARIMA model if the cycle is stochastic.

We can also see evidence of an affirmative answer to another fundamental question for which the answer is not apparent in the top frame. With the passage of time, are the cycles growing in range or holding steady?

What Type is the Trend?

As shown in Figure 4, we have extracted the cycle from the series and are left with the trend component. Now there is another set of fundamental questions.

Now the first and most fundamental question is what kind of trend is it—deterministic or stochastic? Figure 5 is a good example of why the question is so central to the project team.

We can immediately rule out a purely random series (sometimes called white noise). Therefore, is it deterministic? Or is it a stochastic random walk caused by positive autocorrelation? Or is it a drifting random walk?




Figure 5: What is the trend and can we effect it?

It is safe to bet that most of us would interpret the trend as deterministic and draw all sorts of conclusions. We would likely design improvements to the subject process on the conclusions. However, the methods of time series analysis would reveal it to be a random walk, but not a drifting random walk.

Additional questions would arise if we had found the trend to be deterministic. All would be directed at the shape of the plot. Can we see long-term trends such as the influence of business cycles? Can we see change beginning at the time of some operational strategy or event? Can we see signs of process decay?

What is the Process True Sigma?

An essential insight to an operational process is the spread of some percent of its observations around the mean (average). That is the fundamental point of “six sigma.” It is also an important concept because as the saying goes, “We can drown in an average of one foot of water.”

Time for a touch of statistic-speak. Sigma is associated with the residuals around the mean of a process element. Called the “standard error of the mean,” one sigma is the standard deviation of the residuals divided by the square root of the observations reduced by one. The spread around mean is found by multiplying the standard error by the number of sigmas we have chosen to operate by.

The standard computation of variance assumes that all observations in a dataset are independent of each other. However, as already mentioned, independence can never be safely assumed for a time series.

When independence is not the case, the width of the confidence interval around the series at a specified sigma is understated—a little or a lot. This is because the standard error is understated by the autocorrelation. Consequently, the process of the series does not work as

tight and sharp as we think it does. Our six sigma process may be only five sigma—a very different reality, especially if we are not aware.

Therefore, what is the true range to the series? Is the process represented by the series actually as capable as we need it to be or is it actually less? What upstream elements of the process effect the range of the time series element?

To answer the questions, we can find true sigma through the residuals to the fitted model. If the model is a good fit, its residuals will be a series of independent points. Thence, we compute the true sigma with the residuals and construct the confidence interval with it.

What is the Future to Look Like?

The preceding stages of questioning are the feedstock to forecasting. As seen in Figure 6, what was found by modeling is projected into the future. That is if there is a deterministic pattern in the series. If the series is a random walk, any forecasting can only be a very short-term moving average.

The first question is what will we do with forecasts and scenarios? How far out do we want to project? Given what we now know, what did we find in the series' history that will be the assumptions we want to project into the future.? How will improvements to the process change the assumptions and then the picture of the operational future?




Figure 6: Forecasting the past into the future.

There is another intriguing way to forecast from series analytics. We can forecast by identifying and taking advantage of variables that have a lead-lag relationship. Two variables may have the same pattern after extracting the cycles, but have a lead-lag between the patterns. We can also relate two variables through intermediate variables when the two of interest do not have a matching pattern.

The method is called cross-correlation. It is essentially the autocorrelation method represented by Figure 1. The difference is to correlate upon the lags between the residuals of variable after each is fit to a model.

Figure 7 shows the case of building applications related to construction activity. The insert cross-correlogram shows that the level of building activity in a given period is most foretold by building applications submitted one quarter in advance.

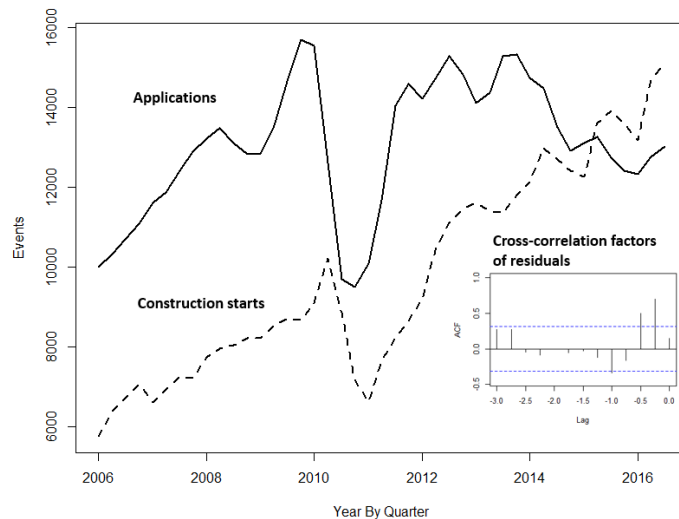


Figure 7: Cross-correlation to find lead-lag series.

Asking and answering questions to lead-lag forecasting is opportunistic. Which sets of variables will increase the readiness of the process to meet its changing challenges? How can we capitalize on the relationships? Can we reach further back for lead indicators by using a chain of variables?

Finally, time series information is the most common of all types consumed in enterprise operations. As goes the enterprise’s ability to look behind the curtain of each series of interest, so goes the enterprise’s ability to reach for operational excellence.

Sources for self-directed learning: *Discovering Statistics Using R*, Field and Miles, 2012 | *Multilevel Modeling Using R*, Holmes, 2014 | *Machine Learning with R*, Lantz, 2015 | *ggplot2, Elegant Graphics for Data Analysis*, Wickham, 2016 | **Introductory Time Series with R**, Cowpertwait and Metcalfe, 2009 | *Event History Analytics with R*, Bostrom, 2012 | **Package “tsoutliers,”** Javier López-de-Lacalle, 2017.

Question-specific articles: [Find What Matters With Relationship Questions](#) | [Know That Improvements Work By Asking Difference Questions](#) | [Explore What Did And May Happen With Time Series Questions](#) | [Find The Time That Is Money By Asking Duration Questions](#) | [Dive Below The Surface Of Process Functioning With Apparency Questions](#)

Richard G. Lamb: [Professional Mission and Bio](#)

Educational website: analytics4strategy.com

