

9

Short-Term Forecasting with ARIMA Models

All models are wrong, some are useful

GEORGE E. P. BOX (1919 – 2013)



In this chapter, we introduce a class of techniques, called ARIMA (for **A**uto-**R**egressive **I**ntegrated **M**oving **A**verage), which can be used to describe stationary time series and nonstationary time series with changing levels. For seasonal time series, the nonstationary ARIMA process is extended to account for a multiplicative seasonal component. These models form the framework for

- expressing various forms of stationary (level) and nonstationary (mostly trending and seasonal) behavior in time series
- producing optimal forecasts with prediction limits for a time series from its own current and past values
- developing a practical and useful modeling process

The topics covered include:

- what an ARIMA forecasting model is used for
- why stationarity is an important concept for ARIMA processes
- how to select ARIMA models through an iterative three-stage procedure
- how, when, and why we should use the Box Jenkins modeling methodology for ARIMA forecasting models
- its relationship to the modern State Space forecasting methodology

Why Use ARIMA Models for Forecasting?

In chapter 2, we constructed a year-by-month ANOVA table as an exploratory, data-driven technique for examining trend and seasonal variation in a time series. This led us to estimate many coefficients for the monthly and yearly means in the rows and columns of a table. (e.g., In Excel, run Data Analysis Add in > Anova: Two factor without replication). Although the results were not intended for generating projections, the ANOVA table provides a useful preliminary view of the data. Figure 9.1 shows (a) a time plot of the classical Series G Box-Jenkins airline passenger miles data and (b) pie chart representing the contribution in the total variation due to Seasonality (83%), Trend (14%) and Other (3%). These data consists of monthly totals of airline passengers from January 1949 to December 1960.





Figure 9.1 (a) Time plot and (b) pie chart of ANOVA decomposition of monthly totals of airline passengers from January 1949 to December 1960 ($n = 144$). (Seasonality 83%, Trend 14%, Other 3%)



The ARIMA modeling approach offers a *model-driven* technique to time series forecasting by using a theoretical framework developed by George E.P. Box (1919– 2013) and Gwilym M. Jenkins (1932– 1982). The theory was first published in a seminal book titled *Time Series Analysis – Forecasting and Control* (1976),

Many studies have shown that forecasts from simple ARIMA models have frequently outperformed larger, more complex econometric systems for a number of economic series. Although it is possible to construct ARIMA models with only 2 years of monthly historical data, the best results are usually obtained when at least 5 to 10 years of data are available - particularly for seasonal time series.

A significant advantage of univariate ARIMA approach is that useful models can be developed in a relatively short time with automated State Space Forecasting algorithms. Therefore, a practitioner can often deliver significant results with ARIMA modeling early in a project for which adequate historical data exist. Because of the sound theoretical underpinnings, the demand forecaster should always consider ARIMA models as an important forecasting tool whenever these models are relevant to the problem at hand.

A drawback of univariate models is that they have limited explanatory capability. The models are essentially sophisticated extrapolative devices that are of greatest use when it is expected that the underlying factors causing demand for products, services, revenues, and so on, will behave in the future in much the same way as in the past. In the short term, this is often a reasonable expectation, however, because these factors tend to change slowly; data tend to show inertia in the short term. However, there are extensions of the ARIMA approach that incorporate explanatory factors for including information such as price, promotions, strikes, and holiday effects. These models are called transfer function (or **dynamic regression**) models, but are beyond the scope of this book.

Much more time is usually required to obtain and validate historical data than to build the models. Therefore, a practitioner can often deliver significant results with ARIMA modeling early in a project for which adequate historical data exist. The forecaster should always consider ARIMA models as an important option in a forecasting toolbox whenever trend/seasonal models are relevant to the problem at hand.

The ARIMA models have proved to be excellent short-term forecasting models for a wide variety of time series.

The Linear Filter Model as a Black Box

The application of ARIMA models is based on the idea that a time series in which successive values are highly dependent (i.e. having “memory” of the past values) can also be thought of as having come from a process involving a series of independent errors or shocks, ε_t . The general form of a (discrete) linear process is:

$$Z_t = \mu + \varepsilon_t + \psi_1 \varepsilon_{t-1} + \psi_2 \varepsilon_{t-2} + \dots + \psi_n \varepsilon_{t-n} + \dots$$

where μ and all ψ_j are fixed parameters and the $\{\varepsilon_t\}$ is a sequence of identically, independently distributed random errors with zero mean and constant variance.

Why is it called a linear filter? The process is linear because Z_t is represented as a linear combination of current and past shocks. It is often referred to as a black box or filter, because the model relates a random input to an output that is time dependent. The input is filtered, or damped, by the equation so that what comes out of the equation has the characteristics that are wanted.




Figure 9.2 Black-box representation of the linear random process.

A linear process can be visualized as a black box as follows (Figure 9.2). White noise or purely random error $\{\varepsilon_t\}$ is transformed to an observed series $\{Y_t\}$ by the operation of a linear filter; the filtering operation simply takes a weighted sum of previous shocks. For convenience, we henceforth write models in terms of Y_t , which has been mean adjusted, that is, $Y_t = Z_t - \mu$. The weights are known as ψ (psi) coefficients. For Y_t to represent a valid stationary process, it is necessary that the coefficients ψ_j satisfy the condition $\sum \psi_j^2 < \infty$.

A linear process is capable of describing a wide variety of practical forecasting models for time series. It can be visualized as a black-box equation transforming random inputs into the observed data.

It can be shown that any linear process can be written formally as a weighted sum of the current error term plus all past shocks. In many problems, such as those in which it is required that future values of a series be predicted; it is necessary to construct a parametric model for the time series. To be useful, the model should be physically meaningful and involve as few parameters as possible. A powerful parametric model that has been widely used in practice for describing empirical time series is called the mixed autoregressive moving-average (ARMA) process:

$$Y_t = \phi_1 Y_{t-1} + \phi_2 Y_{t-2} + \dots + \phi_p Y_{t-p} + \varepsilon_t - \theta_1 \varepsilon_{t-1} - \theta_2 \varepsilon_{t-2} - \dots - \theta_q \varepsilon_{t-q}$$

where p is the highest lag associated with the data, and q is the highest lag associated with the error term. The ARMA processes are important because they are mathematically tractable and they are capable of providing prediction limits for uncertainty measurement. The inputs to the black box (Figure 9.3) are the shocks ε_t and the output is the observed historical data or time series Z_t .

There are some special versions of the ARMA process that are particularly useful in practice. If that weighted sum has only a finite number of nonzero error terms, then the process is known as a moving average (MA) process. It can be shown that the linear process can also be expressed as a weighted sum of the current shock plus all past observed values. If the number of nonzero terms in this expression is finite, then the process is known as an autoregressive (AR) process. The origin of the AR and MA terminology are described a little later with specific examples.




Figure 9.3 Black-box representation of the ARMA process (with a nonstationary filter).

It turns out that an MA process of finite order can be expressed as an AR process of infinite order and that an AR process of finite order can be expressed as an MA process of infinite order. This duality has led to the *principle of parsimony* in the Box-Jenkins methodology, which recommends that the practitioner employ the smallest possible number of parameters for adequate representation of a model. In practice, it turns out that relatively few parameters are needed to make usable forecasting models with business data.

It may often be possible to describe a stationary time series with a model involving fewer parameters than either the MA or the AR process has by itself. Such a model will possess qualities of both autoregressive and moving average models: it is called an ARMA process. An ARMA (1, 1) process, for example, has one prior observed-value term of lag 1 and one prior error term:

$$Y_t = \phi_1 Y_{t-1} + \varepsilon_t - \theta_1 \varepsilon_{t-1}$$

The general form of an ARMA (p, q) process of autoregressive order p and moving average order q looks like:

$$Y_t = \phi_1 Y_{t-1} + \phi_2 Y_{t-2} + \dots + \phi_p Y_{t-p} + \varepsilon_t - \theta_1 \varepsilon_{t-1} - \theta_2 \varepsilon_{t-2} - \dots - \theta_q \varepsilon_{t-q}$$

In short, the ARMA process is a linear random process. It is linear if Y_t is a linear combination of lagged values of Y_t and ε_t . It is random if the errors (also called shocks or disturbances) are introduced into the system in the form of white noise. The random errors are assumed to be independent of one another and to be identically distributed with a mean of zero and a constant variance σ_ε^2 .

The ARMA process is important because it is mathematically tractable and can be shown to produce a wide variety of useful forecasting profiles for time series.

A Model-Building Strategy

The Box-Jenkins approach for ARIMA modeling provides the demand forecaster with a very powerful and flexible tool. Because of its complexity, it is necessary to establish procedures for coming up with practical models. Its difficulty requires a fair amount of sophistication and judgment in its use. Nevertheless, its proven results in terms of forecasting accuracy and understanding processes generating data and forecast accuracy can be invaluable in the hands of a skilled user.

The Box-Jenkins procedure consists of the following three stages.

1. *Identification* consists of using the data and any other knowledge that will tentatively indicate whether the time series can be described with a **moving average** (MA) model, an **autoregressive** (AR) model, or a mixed autoregressive – moving average (ARMA) model.
2. *Estimation* consists of using the data to make inferences about the parameters that will be needed for the tentatively identified model and to estimate values of them.
3. *Diagnostic checking* involves the examination of residuals from fitted models, which can result in either no indication of model inadequacy or model inadequacy, together with information on how the series may be better described.

The procedure is iterative. Thus, residuals should be examined for any lack of randomness and, if we find that residuals are serially correlated, we use this information to modify a tentative model. The modified model is then fitted and subjected to diagnostic checking again until an adequate model is obtained.

Although a Box-Jenkins methodology is an excellent way for forecasting a time series from its own current and past values, it should not be applied blindly and automatically to all forecasting problems.

Identification: Interpreting Autocorrelation and Partial Autocorrelation Functions

The primary tool for identifying an ARMA process is with autocorrelation functions (ACFs) and partial autocorrelation functions (PACFs). ACFs are quantities used for describing the mutual dependence among values in a time series.

Extreme care should be taken in interpreting ACFs, however; the interpretation can be complex and requires some ongoing experience with real data. Attention should be directed to individual values as well as to the overall pattern of the autocorrelation coefficients. In practice, the autocorrelations of low-order ARMA processes are used to help identify models with the Box-Jenkins methodology.

Autocorrelation analysis can be a powerful tool for deciding whether a process shows pure autoregressive behavior or moving average behavior in ARMA models.

Autocorrelation and Partial Autocorrelation Functions

The ACF and PACF are widely used in identifying ARMA models. The corresponding ordinary and partial correlograms are the sample estimates of the ACF and PACF. They play an important role in the identification phase of the Box-Jenkins methodology for forecasting and control applications. Some examples follow, but to simplify writing model equations, we use a notational device known as the **back-shift** operation.

For example, the ACF of an AR (1) process is depicted in Figure 9.4. There is a decaying pattern in the ACF; the decay is exponential if $0 < \phi_1 < 1$ (Figure 9.4a). For $-1 < \phi_1 < 0$ (Figure 9.4b), the ACF is similar but alternates in sign. The PACF shows a single positive value at lag 1 if $0 < \phi_1 < 1$ and a negative spike at lag 1 if $-1 < \phi_1 < 0$.




Figure 9.4 (a) ACF an AR (1) process ($\phi_1 = 0.70$); (b) ACF of an AR (1) process ($\phi_1 = -0.80$).

The PACF is more complex to describe. It measures the correlation between Y_t and Y_{t-k} adjusted for the intermediate values $Y_{t-1}, Y_{t-2}, \dots, Y_{t-k+1}$ (or the correlation between Y_t and Y_{t-k} not accounted for by $Y_{t-1}, Y_{t-2}, \dots, Y_{t-k+1}$). If we denote by ϕ_{kj} the j th coefficient in an AR(k) model, so that ϕ_{kk} is the last coefficient, then it can be shown that the ϕ_{kj} will be nonzero for $k \leq p$ and zero for $k > p$, where p is the order of the autoregressive process. Another way of saying this is that ϕ_{kk} has a cutoff or truncation after lag p . For example, the PACF of an AR (1) process has one spike at lag 1. It has the value ϕ_1 .

Another basic process that occurs fairly often in practice is the AR (2) process. In this case there are two autoregressive coefficients ϕ_1 and ϕ_2 . Figure 9.5 shows the ACF and PACF of an AR (2) model with $\phi_1 = 0.3$ and $\phi_2 = 0.5$. The PACF shows positive values at lags 1 and 2 only. The PACF is very helpful because it suggests that the process is autoregressive and, more important, that it is second-order autoregressive.

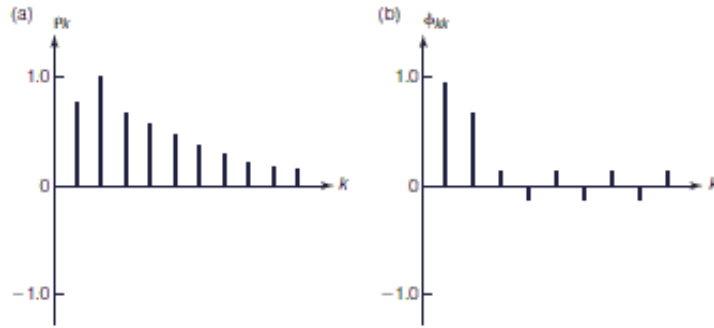


Figure 9.5 (*left*) ACF and (*right*) PACF of an autoregressive AR (2) model with parameters $\phi_1 = 0.3$ and $\phi_2 = 0.5$.




Figure 9.6 (*left*) ACF and (*right*) PACF of an autoregressive AR (2) model with parameters $\phi_1 = 1.2$ and $\phi_2 = -0.64$.

If $\phi_1 = 1.2$ and $\phi_2 = -0.64$, the ACF and PACF have the patterns shown in Figure 9.6. The values in the ACF decay in a sinusoidal pattern; the PACF has a positive value at lag 1 and a negative value at lag 2. There are a number of possible patterns for AR (2) models. A triangular region describes the allowable values for ϕ_1 and ϕ_2 in the stationary case: $\phi_1 + \phi_2 < 1$, $\phi_2 - \phi_1 < 1$, and $-1 < \phi_2 < 1$. If $\phi_1^2 + 4\phi_2 > 0$, the ACF decreases exponentially with increasing lag. If $\phi_1^2 + 4\phi_2 < 0$, the ACF is a damped cosine wave.




Figure 9.7 (a) ACF and (b) PACF of a MA (1) model with positive parameter θ .

The ACF of a MA (q) process is 0, beyond the order q of the process (i.e., it has a cutoff after lag q). For example, the ACF of a MA (1) process has one spike at lag 1, the others are 0. It has the value $\rho_1 = -\theta_1/(1 + \theta_1^2)$ with $|\rho_1| \leq 1/2$.

The PACF of the MA process is complicated, so in Figure 9.7 we display the ACF and PACF of an MA (1) model with positive θ_1 . There is a single negative spike at the lag 1 in the ACF. There is a decaying pattern in the PACF. The ACF of an MA(1) process with negative θ (Figure 9.8) shows a single positive spike, but the PACF shows a decaying pattern with spikes alternating above and below the zero line.




Figure 9.8 (a) ACF and (b) PACF of a MA (1) model with negative parameter θ .

An Important Duality Property

One important consequence of the theory is that the ACF of an AR process behaves like the PACF of an MA process and vice versa. This aspect is known as a duality property of the AR and MA processes. If both the ACF and the PACF attenuate, then a mixed model is called for.

It turns out that the ACF of the pure MA (q) process truncates, becoming 0 after lag q , whereas that for the pure AR (p) process is of infinite extent. MA processes are thus characterized by truncation (spikes ending) of the ACF, whereas AR processes are characterized by attenuation (gradual decay) of the ACF. Derivations of this kind are beyond the scope of this book.

For an AR process, the ACF attenuates and the PACF truncates; conversely, for an MA process, the PACF attenuates and the ACF truncates.

The mixed ARMA (p, q) model contains p AR coefficients ($\phi_1, \phi_2 \dots \phi_p$) and q MA coefficients ($\theta_1, \theta_2, \dots, \theta_q$). This model is useful in that stationary series may often be expressed more parsimoniously (with fewer parameters) in an ARMA model than as a pure AR or pure MA model. In practice, for **mixed ARMA** processes, you should create a catalog of ACF and PACF patterns to establish the orders p and q of the autoregressive and moving average components. The estimated autocorrelation functions, or correlograms, are then matched with the cataloged patterns in order to establish a visual identification of the most useful model for a given situation. Usually, more than one model suggests itself, so that we may tentatively identify several similar models for a particular time series.

Some useful rules for identification are:

- If the **correlogram** cuts off at some point, say $k = q$, then the appropriate model is MA (q).
- If the **partial correlogram** cuts off at some point, say $k = p$, then the appropriate model is AR (p).

- If neither diagram cuts off at some point, but does decay gradually to zero, the appropriate model is ARMA (p', q') for some p', q' .

The ACF and PACF of an ARMA (p, q) model are more complex than either the AR (p) or MA (q) models. The ACF of an ARMA (p, q) process has an irregular pattern at lags 1 through q , then the tail diminishes. The PACF tail also diminishes.

The best way to identify an ARMA process initially is to look for decay or a tail in both the ACFs and PACFs.

The complete chapter can be found in

Change & Chance Embraced

ACHIEVING AGILITY WITH DEMAND

FORECASTING IN THE SUPPLY CHAIN

HANS LEVENBACH, PhD

Copyright © 2017 Hans Levenbach

ISBN-10:0692945989

ISBN-13:978-069294588

Cover art: Rose Tannenbaum, Berkshire TypeGraphic, www.tgo.com

Cartoon art: Don Engel, Isographic LLC

Copy editor: David Coen, Freelance copyeditor

Printed by Createspace, www.createspace.com

ALL RIGHTS RESERVED. No part of this work may be used or reproduced, transmitted, stored or used in any form or by any means graphic, electronic, or mechanical, including but not limited to photocopying, recording, scanning, taping, Web distribution, information networks or information storage and retrieval systems, or in any manner whatsoever without prior written permission.

Professional development: Certified Professional Demand Forecaster (CPDF®), www.cpdftraining.org

Agile Forecasting®, CPDF® are registered trademarks of Delphus, Inc.

Contents



Chapter 1 - Embracing Change & Chance

Inside the Crystal Ball **Error! Bookmark not defined.**

Determinants of Demand	
Demand Forecasting Defined	
Why Demand Forecasting?	
The Role of Demand Forecasting in a Consumer-Driven Supply Chain	4
Is Demand Forecasting Worthwhile?	7
Who Are the End Users of Demand Forecasts in the Supply Chain?	8
Learning from Industry Examples	9
Examples of Product Demand	10
Is a Demand Forecast Just a Number?	11
<i>Creating a Structured Forecasting Process</i>	14
The PEER Methodology: A Structured Demand Forecasting Process	14
<i>Case Example: A Consumer Electronics Company</i>	15
PEER Step 1: Identifying Factors Likely to Affect Changes in Demand	16
The GLOBL Product Lines	17
The Marketplace for GLOBL Products	18
Step 2: Selecting a Forecasting Technique	19
Step 3: Executing and Evaluating Forecasting Models	22
Step 4: Reconciling Final Forecasts	22
<i>Creating Predictive Visualizations</i>	22
<i>Takeaways</i>	26



Chapter 2 - Demand Forecasting Is Mostly about Data:

Improving Data Quality through Data Exploration and Visualization

<i>Demand Forecasting Is Mostly about Data</i>	29
Exploring Data Patterns	29
Learning by Looking at Data Patterns	30
<i>Judging the Quality of Data</i>	30
<i>Data Visualization</i>	35
Time Plots	35
Scatter Diagrams	36
<i>Displaying Data Distributions</i>	37
Overall Behavior of the Data	38

Stem-and-Leaf Displays	39
Box Plots	41
Quantile-Quantile Plots	43
<i>Creating Data Summaries</i>	44
Typical Values	44
The Trimmed Mean	45
Variability	45
Median Absolute Deviation from the Median	45
The Interquartile Difference	46
Detecting Outliers with Resistant Measures	47
<i>The Need for Nonconventional Methods</i>	48
M-Estimators	49
A Numerical Example	49
<i>Why Is Normality So Important?</i>	51
<i>Case Example: GLOBL Product Line B Sales in Region A</i>	52
<i>Takeaways</i>	54




Chapter 3 - Predictive Analytics: Selecting Useful Forecasting Techniques..... 55

<i>All Models Are Wrong. Some Are Useful</i>	56
Qualitative Methods	56
Quantitative Approaches	59
Self-Driven Forecasting Techniques	60
Combining Forecasts is a Useful Method	61
Informed Judgment and Modeling Expertise	62
A Multimethod Approach to Forecasting	64
<i>Some Supplementary Approaches</i>	64
Market Research	64
New Product Introductions	65
Promotions and Special Events	65
Sales Force Composites and Customer Collaboration	65
Neural Nets for Forecasting	66
<i>A Product Life-Cycle Perspective</i>	66
<i>A Prototypical Forecasting Technique: Smoothing Historical Patterns</i>	68
Forecasting with Moving Averages	69
Fit versus Forecast Errors	71
Weighting Based on the Most Current History	73
<i>A Spreadsheet Example: How to Forecast with Weighted Averages</i>	75
Choosing the Smoothing Weight	78
Forecasting with Limited Data	78
Evaluating Forecasting Performance	79
<i>Takeaways</i>	79



Chapter 4 - Taming Uncertainty: What You Need to Know about Measuring Forecast Accuracy80

<i>The Need to Measure Forecast Accuracy</i>	82
Analyzing Forecast Errors	82
Lack of Bias	82
What Is an Acceptable Precision?	83
<i>Ways to Evaluate Accuracy</i>	86
The Fit Period versus the Holdout Period	86
Goodness of Fit versus Forecast Accuracy	87
Item Level versus Aggregate Performance	88
Absolute Errors versus Squared Errors	88
Measures of bias	89
Measures of Precision	90
Comparing with Naive Techniques	93
Relative Error Measures	94
<i>The Myth of the MAPE . . . and How to Avoid It</i>	95
Are There More Reliable Measures Than the MAPE?	96
<i>Predictive Visualization Techniques</i>	96
Ladder Charts	96
Prediction-Realization Diagram	97
<i>Empirical Prediction Intervals for Time Series Models</i>	100
Prediction Interval as a Percentage Miss	101
Prediction Intervals as Early Warning Signals	101
Trigg Tracking Signal	103
<i>Spreadsheet Example: How to Monitor Forecasts</i>	104
<i>Mini Case: Accuracy Measurements of Transportation Forecasts</i>	107
<i>Takeaways</i>	112



Chapter 5 - Characterizing Demand Variability: Seasonality, Trend, and the Uncertainty Factor 114

<i>Visualizing Components in a Time Series</i>	115
Trends and Cycles	116
Seasonality	119
Irregular or Random Fluctuations	122
Weekly Patterns	124
Trading-Day Patterns	124
<i>Exploring Components of Variation</i>	126
Contribution of Trend and Seasonal Effects	127
A Diagnostic Plot and Test for Additivity	130
<i>Unusual Values Need Not Look Big or Be Far Out</i>	132
<i>The Ratio-to-Moving-Average Method</i>	134

- Step 1: Trading-Day Adjustment 135
- Step 2: Calculating a Centered Moving Average 135
- Step 3: Trend-cycle and Seasonal Irregular Ratios 136
- Step 4: Seasonally Adjusted Data 137

GLOBL Case Example: Is the Decomposition Additive or Not? 137

APPENDIX: A Two-Way ANOVA Table Analysis 139

Percent Contribution of Trend and Seasonal Effects 140

Takeaways 140



Chapter 6 - Dealing with Seasonal Fluctuations 141

Seasonal Influences 141

Removing Seasonality by Differencing 143

Seasonal Decomposition 145

Uses of Seasonal Adjustment 146

Multiplicative and Additive Seasonal Decompositions 146

Decomposition of Monthly Data 146

Decomposition of Quarterly Data 151

Seasonal Decomposition of Weekly Point-of-Sale Data 153

Census Seasonal Adjustment Method 156

The Evolution of the X-13ARIMA-SEATS Program 157

Why Use the X-13ARIMA-SEATS Seasonal Adjustment Program? 157

A Forecast Using X-13ARIMA-SEATS 158

Resistant Smoothing 158

Mini Case: A PEER Demand Forecasting Process for Turkey Dinner Cost 162

Takeaways 168



Chapter 7 - Trend-Cycle Forecasting with Turning Points 171

Demand Forecasting with Economic Indicators 171

Origin of Leading Indicators 174

Use of Leading Indicators 174

Composite Indicators 176

Reverse Trend Adjustment of the Leading Indicators 176

Sources of Indicators 178

Selecting Indicators 178

Characterizing Trending Data Patterns 180

Autocorrelation Analysis 180

First Order autocorrelation 182


The Correlogram 183

Trend-Variance Analysis	187
<i>Using Pressures to Analyze Business Cycles</i>	189
<i>Mini Case: Business Cycle Impact on New Orders for Metalworking Machinery</i>	191
1/12 Pressures	192
3/12 Pressures	193
12/12 Pressures	193
<i>Turning Point Forecasting</i>	194
Ten-Step Procedure for a Turning-Point Forecast	195
Alternative Approaches to Turning-Point Forecasting	195
<i>Takeaways</i>	196



Chapter 8 - Big Data: Baseline Forecasting With Exponential Smoothing Models 197

<i>What is Exponential Smoothing?</i>	198
Smoothing Weights	199
The Simple Exponential Smoothing Method	201
<i>Forecast Profiles for Exponential Smoothing Methods</i>	202
Smoothing Levels and Constant Change	204
Damped and Exponential Trends	208
Some Spreadsheet Examples	210
Trend-Seasonal Models with Prediction Limits	216
The Pegels Classification for Trend-Seasonal Models	219
Outlier Adjustment with Prediction Limits	221
Predictive Visualization of Change and Chance – Hotel/Motel Demand	221
<i>Takeaways</i>	225



Chapter 9 - Short-Term Forecasting with ARIMA Models ..226

<i>Why Use ARIMA Models for Forecasting?</i>	226
The Linear Filter Model as a Black Box	227
<i>A Model-Building Strategy</i>	229
Identification: Interpreting Autocorrelation and Partial Autocorrelation Functions	230
Autocorrelation and Partial Autocorrelation Functions	231
An Important Duality Property	233
Seasonal ARMA Process	234
<i>Identifying Nonseasonal ARIMA Models</i>	236
Identification Steps	236
Models for Forecasting Stationary Time Series	236
White Noise and the Autoregressive Moving Average Model	237
One-Period Ahead Forecasts	239
L-Step-Ahead Forecasts	239
Three Kinds of Short-Term Trend Models	241
A Comparison of an ARIMA (0, 1, 0) Model and a Straight-Line Model	241
<i>Seasonal ARIMA Models</i>	244

A Multiplicative Seasonal ARIMA Model	244
Identifying Seasonal ARIMA Models	246
<i>Diagnostic Checking: Validating Model Adequacy</i>	247
<i>Implementing a Trend/Seasonal ARIMA Model for Tourism Demand</i>	249
Preliminary Data Analysis	249
Step 1: Identification	250
Step 2: Estimation	250
Step 3: Diagnostic Checking	251
<i>ARIMA Modeling Checklist</i>	254
<i>Takeaways</i>	255
<i>Postscript</i>	256



Chapter 10 - Demand Forecasting with Regression Models 258

<i>What Are Regression Models?</i>	259
The Regression Curve	260
A Simple Linear Model	260
The Least-Squares Assumption	260
<i>CASE: Sales and Advertising of a Weight Control Product</i>	262
<i>Creating Multiple Linear Regression Models</i>	263
Some Examples	264
<i>CASE: Linear Regression with Two Explanatory Variables</i>	266
<i>Assessing Model Adequacy</i>	268
Transformations and Data Visualization	268
Achieving Linearity	269
Some Perils in Regression Modeling	270
<i>Indicators for Qualitative Variables</i>	273
Use of Indicator Variables	273
Qualitative Factors	274
Dummy Variables for Different Slopes and Intercepts	275
Measuring Discontinuities	275
Adjusting for Seasonal Effects	276
Eliminating the Effects of Outliers	276
<i>How to Forecast with Qualitative Variables</i>	277
Modeling with a Single Qualitative Variable	278
Modeling with Two Qualitative Variables	279
Modeling with Three Qualitative Variables	279
<i>A Multiple Linear Regression Checklist</i>	281
<i>Takeaways</i>	282



Chapter 11 - Gaining Credibility Through Root-Cause

Analysis and Exception Handling 283

***The Diagnostic Checking Process in Forecasting*..... 284**

***The Role of Correlation Analysis in Regression Modeling* 284**

Linear Association and Correlation 285
The Scatter Plot Matrix 286
The Need for Outlier Resistance in Correlation Analysis 287

***Using Elasticities* 288**

Price Elasticity and Revenue Demand Forecasting 290
Cross-Elasticity 291
Other Demand Elasticities 292
Estimating Elasticities 292

***Validating Modeling Assumptions: A Root-Cause Analysis* 293**

A Run Test for Randomness 296
Nonrandom Patterns 297
Graphical Aids 299
Identifying Unusual Patterns 299

***Exception Handling: The Need for Robustness in Regression Modeling* 301**

Why Robust Regression? 301
M-Estimators 301
Calculating M-Estimates 302

***Using Rolling Forecast Simulations* 304**

Choosing the Holdout Period 304
Rolling Origins 305
Measuring Forecast Errors over Lead Time 306

***Mini Case: Estimating Elasticities and Promotion Effects* 306**

Procedure 308
Taming Uncertainty 310

***Multiple Regression Checklist* 311**

***Takeaways* 313**



Chapter 12 - The Final Forecast Numbers: Reconciling Change & Chance

.....316

***Establishing Credibility* 317**

Setting Down Basic Facts: Forecast Data Analysis and Review 317
Establishing Factors Affecting Future Demand 318
Determining Causes of Change and Chance 318
Preparing Forecast Scenarios 318

<i>Analyzing Forecast Errors</i>	319
Taming Uncertainty: A Critical Role for Informed Judgment	320
Forecast Adjustments: Reconciling Sales Force and Management Overrides	321
Combining Forecasts and Methods	322
Verifying Reasonableness	323
Selecting 'Final Forecast' Numbers	324
<i>Gaining Acceptance from Management</i>	325
The Forecast Package	325
Forecast Presentations	326
<i>Case: Creating a Final Forecast for the GLOBL Company</i>	328
Step 1: Developing Factors	329
Impact Change Matrix for the Factors Influencing Product Demand	330
The Impact Association Matrix for the Chosen Factors	331
Exploratory Data Analysis of the Product Line and Factors Influencing Demand	332
Step 2: Creating Univariate and Multivariable Models for Product Lines	334
Handling Exceptions and Forecast Error Analysis	335
Combining Forecasts from Most Useful Models	337
An Unconstrained Baseline Forecast for GLOBL Product Line B, Region A	338
Step 3: Evaluating Model Performance Summaries	341
Step 4: Reconciling Model Projections with Informed Judgment	342
<i>Takeaways</i>	343



Chapter 13 - Creating a Data Framework for Agile Forecasting and Demand Management..... 344

<i>Demand Management in the Supply Chain</i>	345
Data-Driven Demand Management Initiatives	346
Demand Information Flows	347
<i>Creating Planning Hierarchies for Demand Forecasting</i>	349
What Are Planning Hierarchies?	349
Operating Lead Times	350
Distribution Resource Planning (DRP)—A Time-Phased Planned Order Forecast	350
Spreadsheet Example: How to Create a Time-Phased Replenishment Plan	352
<i>A Framework for Agility in Forecast Decision Support Functions</i>	353
The Need for Agile Demand Forecasting	354
Dimensions of Demand	354
A Data-Driven Forecast Decision Support Architecture	355
Dealing with Cross-Functional Forecasting Data Requirements	358
Specifying Customer/Location Segments and Product Hierarchies	358
<i>Automated Statistical Models for Baseline Demand Forecasting</i>	360
Selecting Useful Models Visually	363
Searching for Optimal Smoothing Procedures	367
Error-Minimization Criteria	368
Searching for Optimal Smoothing Weights	368
Starting Values	368
Computational Support for Management Overrides	369
<i>Takeaways</i>	372



Chapter 14 - Blending Agile Forecasting with an Integrated Business Planning Process 373

PEERing into the Future: A Framework for Agile Forecasting in Demand Management 374

The Elephant and the Rider Metaphor	374
Prepare	374
Execute	376
Evaluate	376
Reconcile	381

Creating an Agile Forecasting Implementation Checklist 385

Selecting Overall Goals	385
Obtaining Adequate Resources	386
Defining Data	386
Forecast Data Management	387
Selecting Forecasting Software	387
Forecaster Training	388
Coordinating Modeling Efforts	388
Documenting for Future Reference	388
Presenting Models to Management	389

Engaging Agile Forecasting Decision Support 389

Economic/Demographic Data and Forecasting Services	389
Data and Database Management	390
Modeling Assistance	390
Training Workshops	390

The Forecast Manager's Checklists 391

Forecast Implementation Checklist	391
Software Selection Checklist	392
Large-Volume Demand Forecasting Checklist	393

Takeaways 394