

Machine Learning for Finance

Prepared by Dr. Diaa Salama

FNT633

Forecasting
Time Series Models

Lecture 8



Machine Learning for Finance

Lecture 8

FNT633

Forecasting
Time Series Models

Prepared by
Dr. Diaa Salama

www.midocean.ae

- Forecasting
- Time Series Models

What Kind of Problem do I Need to Solve? How do I Solve it?

The Problem to Solve	The Category of Techniques	Covered in this Course
I want to group items by similarity. I want to find structure (commonalities) in the data	Clustering	K-means clustering
I want to discover relationships between actions or items	Association Rules	Apriori
I want to determine the relationship between the outcome and the input variables	Regression	Linear Regression Logistic Regression
I want to assign (known) labels to objects	Classification	Naïve Bayes Decision Trees
I want to find the structure in a temporal process I want to forecast the behavior of a temporal process	Time Series Analysis	ACF, PACF, ARIMA
I want to analyze my text data	Text Analysis	Regular expressions, Document representation (Bag of Words).

FORECASTING



FORECASTING



Agenda

- Appreciate the importance of forecasting to every organization
- List different types of forecasting method
- Discuss the characteristics of judgmental forecasting
- Use a variety of approaches to judgmental forecasting
- Describe the characteristics of projective forecasting
- Understand the importance of time series
- Calculate errors and a tracking signal for forecasts
- Forecast using simple averages, moving averages, and exponential smoothing
- Forecast time series with seasonality and trend.

- **What is forecasting?**
 - Primary Function is to **Predict the Future** using (time series related or other) **data** we have in hand
- **Why are we interested?**
 - **Affects the decisions we make today**
- **Where is forecasting used ?**
 - **forecast demand for products and services**
 - **forecast availability/need for manpower**
 - **forecast inventory and materiel needs daily**

What makes a good Forecast

- **It should be timely**
- **It should be as accurate** as possible
- **It should be reliable**
- It should be in meaningful units
- **The method should be easy to use and understand in most cases.**

- **Generate forecasts for data with different patterns:
level, trend, seasonality, and cyclical**
- **Describe causal modeling using linear regression**
- **Compute forecast accuracy**
- **Explain how forecasting models should be selected.**

Forecasting steps

- Decide what needs to be forecast
 - Level of detail, units of analysis & time horizon required
- Evaluate and analyze appropriate data
 - Identify needed data & whether it's available
- Select and test the forecasting model
 - Cost, ease of use & accuracy
- Generate the forecast
- Monitor forecast accuracy over time

Reasons for forecasting in organizations

- We have **to calculate a break-even analysis to maximize organization profit** by calculating the number of units of a product that a firm **must sell before it begins to make a profit**
- If sales of a new product are unlikely to reach the break-even point, the organization should not bother making it.
- Unfortunately, it cannot know exactly what future sales will be – so its only option is to forecast likely sales.

Reasons for forecasting in organizations Cont;

- So managers should not base their decisions on present situations, but on conditions as they will be when the decisions become effective.
- These conditions must be forecast, and this suggests that forecasting is a core concern of every organization.
- it becomes clear that virtually every decision made by managers depends on forecasts of future conditions.
- The implication is that forecasting is an essential job in every organization.
- there is no single best way of forecasting, and managers have to choose the method that best suits their needs.

Reasons for forecasting in organizations Cont;

- A lot of the following discussion talks of 'forecasting demand',
 - but this is only a convenient label.
 - In practice, **virtually everything has to be**
- | | |
|---|--|
| <ul style="list-style-type: none">• forecast –demand,• costs,• availability of resources,• weather,• staff turnover,• competitors' actions,• exchange rates, taxes, | <ul style="list-style-type: none">• inflation,• energy consumption,• traffic levels,• customer complaints–• and just about anything else. |
|---|--|

What is a Time Series?

- A collection of data $\{X_t\}$ recorded over a period of time (daily, weekly, monthly, quarterly, yearly, etc.), analyzed to understand the past, in order to predict the future (**forecast**), helping managers and policy makers to make well-informed and sound decisions.
- An important feature of most time series is that **observations close together in time tend to be correlated (serially dependent)**.
- Time could be **discrete**: $t = 1, 2, \dots$, or **continuous**: $t > 0$.

Methods of forecasting

- One classification of forecasting methods concerns the time in the future they cover
 - *Long-term forecasts*; look ahead **several years** – the time typically needed to **build a new factory**
 - *Medium-term forecasts* look ahead **between three months and two years** –the time typically needed to **replace an old product by a new one.**
 - *Short-term forecasts* cover **the next few weeks** – describing the **continuing demand for a product.**

- Another classification of methods draws a distinction between qualitative and quantitative approaches

Types of Forecasting Methods

Quantitative




Qualitative




Methods of forecasting

- Another classification of methods draws a distinction between qualitative and quantitative approaches




- When a company is already making a product, it has records of past sales and knows the factors that affect them. Then it can use a quantitative method
- to forecast future demand. There are two ways of doing this:



Methods of forecasting

- **Causal methods** analyze the effects of outside influences and use these to produce forecasts (linear regression) .
- **Projective methods** examine the pattern of past demand and extend this into the future.
- If demand in the past five weeks has been 10, 20, 30, 40 and 50, it seems reasonable to project this pattern into the future and suggest that demand in the next week will be 60.
- Both of these methods need accurate, quantified data.




Methods of forecasting

- Forecasting methods are classified into two groups:

	Qualitative Methods	Quantitative Methods
1. Characteristics	Based on human judgment, opinions; subjective and nonmathematical.	Based on mathematics; quantitative in nature.
2. Strengths	Can incorporate latest changes in the environment and "inside information."	Consistent and objective; able to consider much information and data at one time.
3. Weaknesses	Can bias the forecast and reduce forecast accuracy.	Often quantifiable data are not available. Only as good as the data on which they are based.


Quantitative Forecasting Methods



Time Series Pattern

- Projective forecasts often work with time series, such as monthly unemployment figures, daily rainfall, weekly sales, quarterly profit, and annual fuel consumption.
- When you have a time series it is always useful to draw a graph, and a simple scatter diagram shows any underlying patterns.
 - constant series (**Level**) ; with observations taking roughly the same value over time, such as annual rainfall
 - series with a **trend** ; with values either rising or falling steadily, such as the gross national product per capita
 - **seasonal** series; which have a **cyclical** component, such as the weekly sales of soft drinks.

Time Series Pattern



Time Series Pattern

- If observations followed such simple patterns, there would be no problems with forecasting.
- Unfortunately, there are nearly always differences between actual observations and the underlying pattern.
- Random noise is applied to the underlying pattern





Figure 10.3 Patterns in time series including noise


Constant series, for example, does not always take exactly the same value, but is somewhere close. Then:
200 205 194 195 208 203 200 193 201 198
is a constant series of 200 with superimposed noise.

Time Series Pattern




The Nature of Time Series Data - Example 1

Quarterly earnings per share for 1960Q1 to 1980Q4 of the U.S. company, Johnson & Johnson, Inc. (Upward trend).




The Nature of Time Series Data - Example 2

A small .1 second (1000 points) sample of recorded speech for the phrase "aaa...hhh". (Regular repetition of small wavelets).




The Nature of Time Series Data - Example 3

Annual numbers of lynx trappings in McKenzie river in Northwest Territories of Canada over the years 1821-1934. (Aperiodic cycles of approximately 10 years).




The Nature of Time Series Data - Example 4

Monthly Airline Passenger Numbers 1949-1960. (Seasonality appears to increase with the general trend).



The Nature of Time Series Data - Example 5

Returns of the New York Stock Exchange (NYSE) from February 2, 1984 to December 31, 1991. (Average return of approximately zero, however, volatility (or variability) of data changes over time).



Components of a Time Series

In general, the time series can be decompose into 4 components: **Secular Trend, Seasonal Variation, Cyclical Variation, and Irregular Variation** that can be modelled deterministically with mathematical functions of time ([see the next slide](#)).

- **General Trend:** The smooth long term direction (upward or downward) of a time series.
- **Seasonal Variation:** Patterns of change in a time series within a year which tend to repeat each year.
- **Cyclical Variation:** The rise and fall of a time series over periods longer than one year.
- **Irregular Variation:** Random and follow no regularity in the occurrence pattern.

- Forecaster looks for data patterns as
 - *Data = historic pattern + random variation*
- Historic pattern to be forecasted:
 - **Level** (long-term average) – data fluctuates around a constant mean
 - **Trend** – data exhibits an increasing or decreasing pattern
 - **Seasonality** – any pattern that regularly repeats itself and is of a constant length
 - **Cycle** – patterns created by economic fluctuations
- Random Variation cannot be predicted

Common Measures of Error

If we define y_t as the actual observation in period t , and F_t as the forecast value, the error is:

$$E_t = y_t - F_t$$

$$\text{mean error} = \frac{\sum E_t}{n} = \frac{\sum (y_t - F_t)}{n}$$

COMMON MEASURES OF ERROR

Mean Absolute Deviation (MAD)

$$\text{MAD} = \frac{\sum |\text{Actual} - \text{Forecast}|}{n}$$

Mean Squared Error (MSE)

$$\text{MSE} = \frac{\sum (\text{Forecast Errors})^2}{n}$$




Figure 10.4 Errors in forecasts

E_t = error in the forecast in period t
= actual observation in period t – forecast value

■ Example

Two forecasting methods give the following results for a time series. Which method is better?

t	1	2	3	4	5
y_t	20	22	26	19	14
F_t with method 1	17	23	24	22	17
F_t with method 2	15	20	22	24	19

Common Measures of Error

- Method 1 gives forecasts that are always nearer to actual demand than method 2, so in this case the decision is easy. We can confirm this by calculating the errors.
- The first method has lower mean absolute error and mean squared error, and is the better choice.
- The second method has slightly less bias, measured

Method 1

t	1	2	3	4	5	Total	Mean
y_t	20	22	26	19	14	101	20.2
F_t with method 1	17	23	24	22	17	103	20.6
E_t	3	-1	2	-3	-3	-2	-0.4
$ E_t $	3	1	2	3	3	12	2.4
$(E_t)^2$	9	1	4	9	9	32	6.4

Method 2

t	1	2	3	4	5	Total	Mean
y_t	20	22	26	19	14	101	20.2
F_t with method 2	15	20	22	24	19	100	20
E_t	5	2	4	-5	-5	1	0.2
$ E_t $	5	2	4	5	5	21	4.2
$(E_t)^2$	25	4	16	25	25	95	19

◎ Naive:

$$F_{t+1} = A_t$$

- The forecast is equal to the actual value observed during the last period - good for level patterns

◎ Simple Mean:

$$F_{t+1} = \sum A_t / n$$

- The average of all available data - good for level patterns

◎ Moving Average:

$$F_{t+1} = \sum A_t / n$$

- The average value over a set time period (e.g.: the last four weeks)
- Each new forecast drops the oldest data point & adds a new observation
- More responsive to a trend but still lags behind actual data

Time Series Model

John Butler runs two dental surgeries, with the following numbers of patients visiting each over the past five weeks. Use simple averages to forecast the numbers of patients visiting in week six. How accurate are the forecasts? What are the forecasts for week 24?

Week	1	2	3	4	5
Surgery 1	98	100	98	104	100
Surgery 2	140	66	152	58	84

Solution

Calculating the simple averages:

- Surgery 1: $F_6 = (\sum y_t)/5 = 500/5 = 100$
- Surgery 2: $F_6 = (\sum y_t)/5 = 500/5 = 100$

Although the forecasts are the same, there is clearly less noise in the figures for Surgery 1 than for Surgery 2, so you should be more confident in the forecast for Surgery 1 and expect smaller errors.

Simple averages assume the underlying pattern is constant, so the forecasts for week 24 are the same as the forecasts for week 6, that is 100.

Simple Averages model

- Simple averages can give good forecasts for constant series
 - Simple averages do not work well when the underlying Pattern changes
 - The problem is that older data tends to swamp the latest figures and the forecast is very slow to follow the changing pattern.

Simple Average model limitations

- Suppose that demand for an item has been constant at 100 units a week for the past two years (104 weeks). Simple averages give a forecast demand for week 105 of 100 units.
- But if the actual demand in week 105 suddenly rises to 200 units, simple averages give a forecast for week 106

$$F_{106} = (104 \times 100 + 200)/105 = 100.95$$

- If demand continues at 200 units a week, following forecasts are:

$$F_{107} = 101.89, F_{108} = 102.80, F_{109} = 103.70, F_{110} = 104.59, \dots \text{etc.}$$

- If demand continues at 200 units a week, following forecasts are:

Simple Average model limitations

- Simple averages only actually work for constant series.
- The forecasts are rising but the response is very slow.
- But very few time series are in fact stable over long periods, so this restriction makes them of limited value.
- The problem with simple averages is that old, out-of-date data tends to swamp newer, more relevant data.
- A way around this is to ignore old data and use only a number of the most recent observations.

- As patterns of demand tend to vary over time
 - **only a certain amount of historical data is relevant** to future forecasts.
 - The problem with simple averages is that **old, out-of-date data tends to swamp newer**, more relevant data
- the principle of moving averages is to **ignore old data and** use only a number of the most recent observations.

$$\begin{aligned} F_{t+1} &= \text{average of } n \text{ most recent observations} \\ &= \frac{\text{latest demand} + \text{next latest} + \dots + n\text{th latest}}{n} \\ &= \frac{y_t + y_{t-1} + \dots + y_{t-n+1}}{n} \end{aligned}$$

Moving Average method


- ◆ MA is a series of arithmetic means
- ◆ Used if little or no trend
- ◆ Used often for smoothing
 - ◆ Provides overall impression of data over time

$$\text{Moving average} = \frac{\Sigma \text{ demand in previous } n \text{ periods}}{n}$$

Moving Average Example

3-Month Moving Average	Actual Shed Sales	Month
10	January	10
12	February	12
13	March	13
16	April	$(10 + 12 + 13)/3 = 11 \frac{2}{3}$
19	May	$(12 + 13 + 16)/3 = 13 \frac{2}{3}$
23	June	$(13 + 16 + 19)/3 = 16$
26	July	$(16 + 19 + 23)/3 = 19 \frac{1}{3}$

Graph of Moving Average



WORKED EXAMPLE 10.3

Epsilon Court Co. has recorded the following numbers of customer complaints each month:

Month, t	1	2	3	4	5	6	7
Complaints, y_t	135	130	125	135	115	80	105

Continuously changing conditions mean that any data over three months old is no longer reliable. Use a moving average to forecast the number of complaints for the future.

Moving Average Quiz Solution

WORKED EXAMPLE 10.3

Epsilon Court Co. has recorded the following numbers of customer complaints each month:

Month, t	1	2	3	4	5	6	7
Complaints, y_t	135	130	125	135	115	80	105

Continuously changing conditions mean that any data over three months old is no longer reliable. Use a moving average to forecast the number of complaints for the future.

Solution

Only data more recent than three months is reliable, so we can use a three-month moving average to

forecast. Consider the situation at the end of month 3, when we can calculate the forecast for month 4 as:

$$F_4 = (y_1 + y_2 + y_3)/3 = (135 + 130 + 125)/3 = 130$$

At the end of month 4 we know the actual number is 135, so the forecast for month 5 is:

$$F_5 = (y_2 + y_3 + y_4)/3 = (130 + 125 + 135)/3 = 130$$

Similarly,

$$F_6 = (y_3 + y_4 + y_5)/3 = (125 + 135 + 115)/3 = 125$$

$$F_7 = (y_4 + y_5 + y_6)/3 = (135 + 115 + 80)/3 = 110$$

$$F_8 = (y_5 + y_6 + y_7)/3 = (115 + 80 + 105)/3 = 100$$

Moving Average example with different value of N


WORKED EXAMPLE 10.4

Column B in Figure 10.5 shows the monthly demand for a product. Use moving averages of three, six and nine months to give forecasts one month ahead.

A	B
Moving Averages	
Month	Demand
1	16
2	13
3	11
4	15
5	17
6	14
7	12
8	15
9	10
10	12
11	18
12	21
13	25
14	30
15	28
16	26
17	34
18	36
19	

Moving Average example with different value of N

Moving Averages				
Month	Demand	Forecasts		
		$n = 3$	$n = 6$	$n = 9$
1	16			
2	13			
3	11			
4	15	13.33		
5	17	13.00		
6	14	14.33		
7	12	15.33	14.33	
8	15	14.33	13.67	
9	10	13.67	14.00	
10	12	12.33	13.83	13.67
11	18	12.33	13.33	13.22
12	21	13.33	13.50	13.78
13	25	17.00	14.67	14.89
14	30	21.33	16.83	16.00
15	28	25.33	19.33	17.44
16	26	27.67	22.33	19.00
17	34	28.00	24.67	20.56
18	36	29.33	27.33	22.67
19		32.00	29.83	25.56



Moving Average model advantages

- You can see from the last example that moving averages give good results for stable patterns, **but they tend to fall behind trends.**
- However, they have a very useful property for data with strong seasonal variations; when you choose n equal to the number of periods in a season

- Although moving averages overcome some of the problems of simple averages, **the method still has defects**, including the following:
 - It **gives all observations the same weight**.
 - It **works well only with constant time series** (as we have seen, it lags behind trends and either removes seasonal **factors** or gets the timing wrong).
 - It needs a lot of historical data to update the forecast.
 - **The choice of n is often arbitrary**

We can overcome the first of these problems by assigning different weights to observations.

Moving Average model for seasonality


WORKED EXAMPLE 10.5

Use a moving average with two, four and six periods to calculate the forecasts one period ahead for the following data.

Quarter	1	2	3	4	5	6	7	8	9	10	11	12
Demand	100	50	20	150	110	55	25	140	95	45	30	145

Moving Average model for seasonality

Moving Averages					
Year	Quarter	Demand	Forecasts		
			$n=2$	$n=4$	$n=6$
1	1	100			
	2	50			
	3	20	75.0		
	4	150	35.0		
2	1	110	85.0	80.00	
	2	55	130.0	82.50	
	3	25	82.5	83.75	80.83
	4	140	40.0	85.00	68.33
3	1	95	82.5	82.50	83.33
	2	45	117.5	78.75	95.83
	3	30	70.0	76.25	78.33
	4	145	37.5	77.75	65.00
4	1		73.3	78.75	80.00



Moving Average model for seasonality

- This data has a clear seasonal pattern, with a peak in the fourth quarter of every year.
- The moving averages with both $n=2$ and $n=6$ have responded to the peaks and troughs of demand,
 - but neither has got the timing right: both forecasts lag behind demand.
 - As you would expect, the two-period moving average is much more responsive than the six-period one
 - the most interesting result is the four-period moving average, which has completely deseasonalised the data.



Thank You