Time Series Forecasting Methods

Nathaniel Derby

Statis Pro Data Analytics
Seattle, WA, USA

Victoria SAS Users Group, 10/7/2008

Outline

Objectives

- What is time series data?
- What do we want out of a forecast?
  - Long-term or short-term?
  - Broken down into different categories/time units?
  - Do we want *prediction intervals*?
  - Do we want to measure effect of X on Y? (scenario forecasting!)
- What methods are out there to forecast/analyze them?
- How do we decide which method is best?
- How can we use SAS for all this?

What is Time Series Data?

Time Series data = Data with a pattern ("trend") over time.

- Ignore time trend = Get wrong results (tomorrow's talk!)

Valve Orders vs Shipments

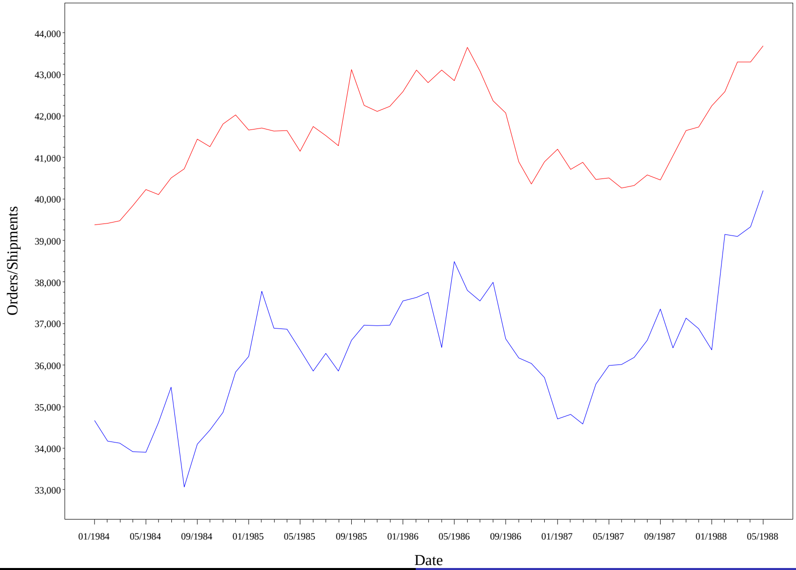

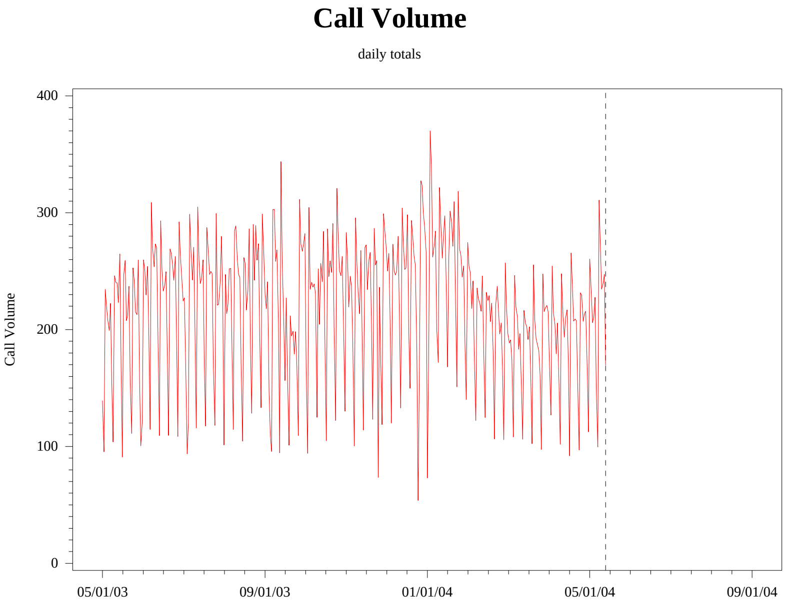
Call Volume
daily totals
Call Volume
400
300
200
100
0
05/01/03
09/01/03
01/01/04
05/01/04
09/01/04

Base Data Set

	date	calls
1	05/01/2003	139
2	05/02/2003	95
3	05/03/2003	235
4	05/04/2003	221
5	05/05/2003	208
6	05/06/2003	199
7	05/07/2003	222
8	05/08/2003	152
9	05/09/2003	104
10	05/10/2003	246
11	05/11/2003	240
12	05/12/2003	240
13	05/13/2003	223

What do we want out of a Forecast?

- Long-term:
  - Involves *many* assumptions! (e.g., global warming)
  - Involves tons of uncertainty.
  - Keynes: “In the long run we are all dead”.
  - We’ll focus on the short term.
- Different categories?
  - Two strategies for forecasting A, B and C:
    1. Forecast their combined total, then break it down by percentages.
    2. Forecast them separately.
  - Idea: Do (1) unless percentages are unstable.

What do we want out of a Forecast?

- Different time units?
  - Two strategies for forecasting at two different time units (e.g., daily and weekly):
    1. Forecast weekly, then break down into days by percentages.
    2. Forecast daily, then aggregate into weeks.
  - Idea: Idea: Do (1) unless percentages are unstable.
- Do we want prediction intervals?
  - Prediction interval = Interval where data point will be with 90/95/99% probability.
  - Yes, we want them!

What do we want out of a Forecast?

- Do we want to measure effect of X on Y?
  - Ex: Marketing campaign $\Rightarrow$ calls to call center.
  - Harder to do, but
  - Allows for scenario forecasting!
  - Idea: Do it, but only with most important Xs.

Remaining Questions: Basis of this talk:

- What methods are out there to forecast/analyze them?
- How do we decide which method is best?
- How can we use SAS for all this?
  - Methods will require `ETS` package.

Strategies

Two stages:

- *Univariate* (one variable) forecasting:
  - Forecasts Y from trend alone.
  - Gives us a basic setup.
- *Multivariate* (many variables) forecasting:
  - Forecasts Y from trend and other variables $X_1, X_2, \ldots$.
  - Allows for “what if” scenario forecasting.
  - May or may not make more accurate forecasts.

Univariate Forecasting - Intro

- Gives us a benchmark for comparing multivariate methods.
- Could give better forecasts than multivariate.
- Some methods can be extended to multivariate.
- Currently three methods:
  - Seasonal moving average (*very* simple)
  - Exponential smoothing (simple)
  - ARIMA (complex)
- More complex methods, for later on (for me):
  - State space (promising)
  - Bayesian (maybe . . .)
  - Wavelets? (forget it!)

Once Again ...

Q: Why not use `PROC REG`?

$$Y_t = \beta_0 + \beta_1 X_t + Z_t$$

- A: We can get misleading results (tomorrow's talk, next year's paper).

Seasonal Moving Average

Simple but sometimes effective!

- Moving Average:

$$\text{Forecast} = \text{Average of last } n \text{ days.}$$

- Seasonal Moving Average:

$$\text{Forecast} = \text{Average of last } n \text{ Mondays.}$$

- After a certain point, forecast the same for each of same weekday.
  - Doesn't allow for a trend.
- Not based on a *model* $\Rightarrow$ No prediction intervals.

SAS Code

Making lags in a `DATA` step (to make the averages) is not fun:

Making 4 lags (Brocklebank and Dickey, p. 45)

```
DATA movingaverage;
  ...
  RETAIN date calls1-calls4;
  OUTPUT;
  calls4=calls3;
  calls3=calls2;
  calls2=calls1;
  calls1=calls;
RUN;
```

SAS Code

Much easier with a trick with `PROC ARIMA`.

Seasonal = "over a weekday", averaging over past 5 weeks:

$$Y_t = \frac{1}{5}\left(Y_{t-7} + Y_{t-14} + Y_{t-21} + Y_{t-28} + Y_{t-35}\right)$$

Forecasting 3 weeks ahead, seasonal moving average

```
PROC ARIMA data=mydata;
  IDENTIFY var=calls noprint;
  ESTIMATE p=( 7, 14, 21, 28, 35 ) q=0 ar=0.2 0.2 0.2 0.2 0.2
    noest noconstant noprint;
  FORECAST lead=21 out=foremave id=date interval=day noprint;
RUN;
QUIT;
```

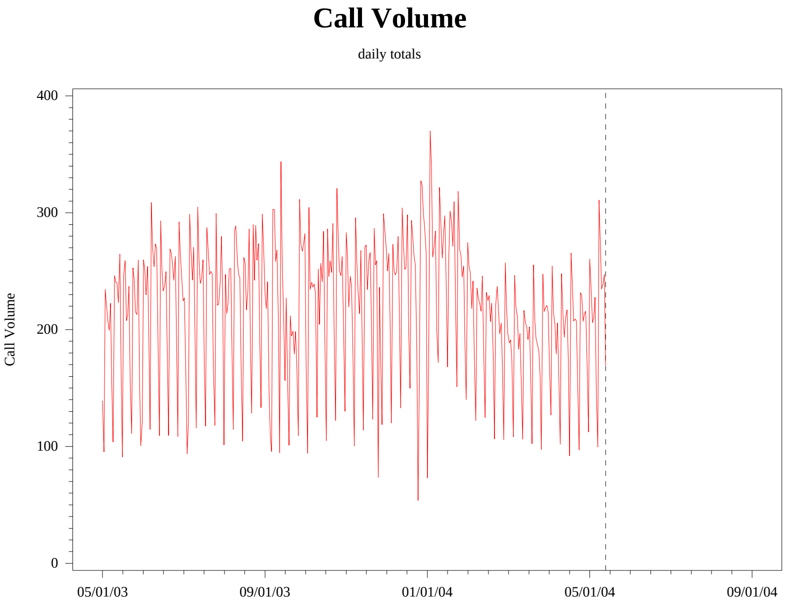
Call Volume
daily totals
Call Volume
400
300
200
100
0
05/01/03
09/01/03
01/01/04
05/01/04
09/01/04

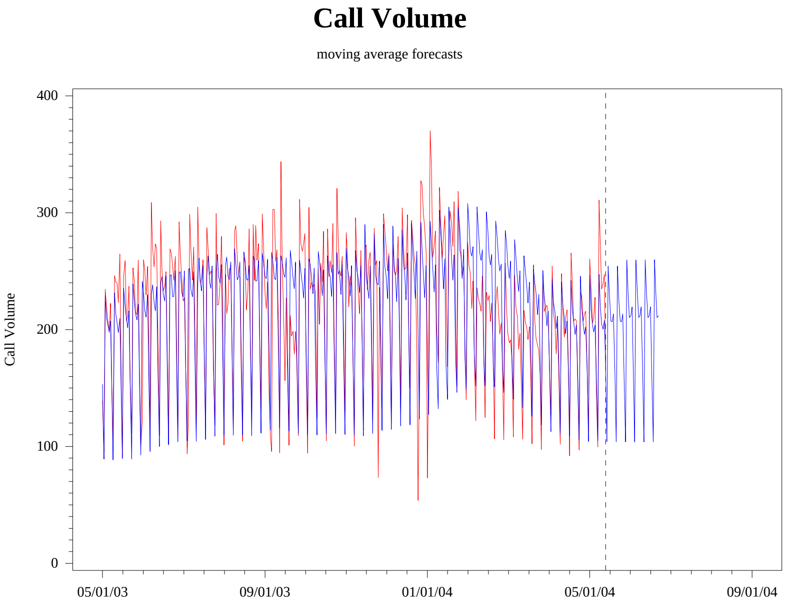
Call Volume
moving average forecasts
Call Volume
400
300
200
100
0
05/01/03
09/01/03
01/01/04
05/01/04
09/01/04

Exponential Smoothing I

Notation: $\hat{y}_t(h)$ = forecast of Y at horizon h, given at time t.

- Idea 1: Predict Y_{t+h} by taking weighted sum of past observations:

$$\hat{y}_t(h) = \lambda_0 y_t + \lambda_1 y_{t-1} + \cdots$$

  Assumes $\hat{y}_t(h)$ is constant for all horizons h.

- Idea 2: Weight recent observations heavier than older ones:

$$\lambda_i = c\alpha^i,\ 0 < \alpha < 1 \ \Rightarrow\ \hat{y}_t(h) = c\left(y_t + \alpha y_{t-1} + \alpha^2 y_{t-2} + \cdots\right)$$

  where c is a constant so that weights sum to 1.

Exponential Smoothing II

$$\hat{y}_t(h) = c\left(y_t + \alpha y_{t-1} + \alpha^2 y_{t-2} + \cdots\right)$$

- Weights are *exponentially decaying* (hence the name).
- Choose α by minimizing squared one-step prediction error.

Overall:

- Just a weighted moving average.
- Can be extended to include trend and seasonality.
- Prediction intervals? Sort of ...

SAS Code

All done with `PROC FORECAST`:

- `method=expo trend=1` for simple.
- `method=expo trend=2` for trend.
- `method=winters seasons=( 7 )` for seasonal.

Forecasting 3 weeks ahead, exponential smoothing

```
PROC FORECAST data=mydata method=xx interval=day lead=21
    out=foreexsm outactual outlimit;
  VAR calls;
  ID date;
RUN;
```

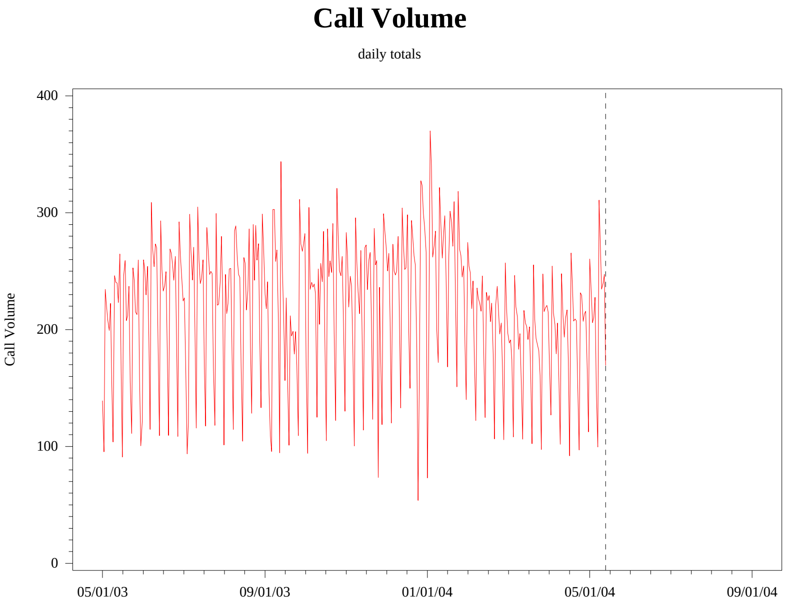
Call Volume
daily totals
Call Volume
400
300
200
100
0
05/01/03
09/01/03
01/01/04
05/01/04
09/01/04

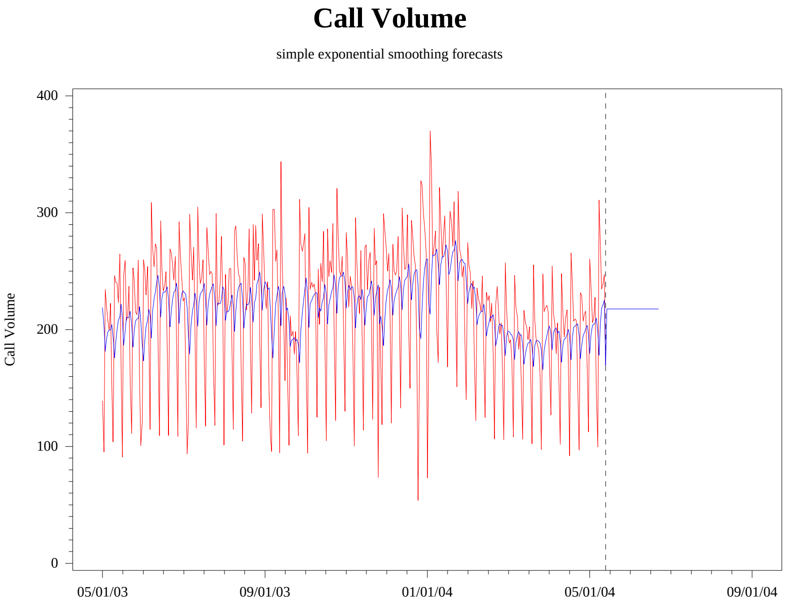
Call Volume
simple exponential smoothing forecasts
Call Volume
400
300
200
100
0
05/01/03
09/01/03
01/01/04
05/01/04
09/01/04

Call Volume

double exponential smoothing forecasts

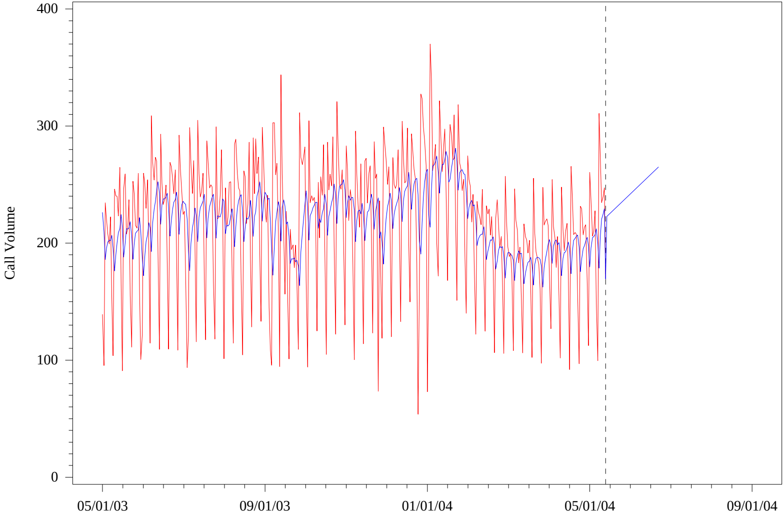

Call Volume

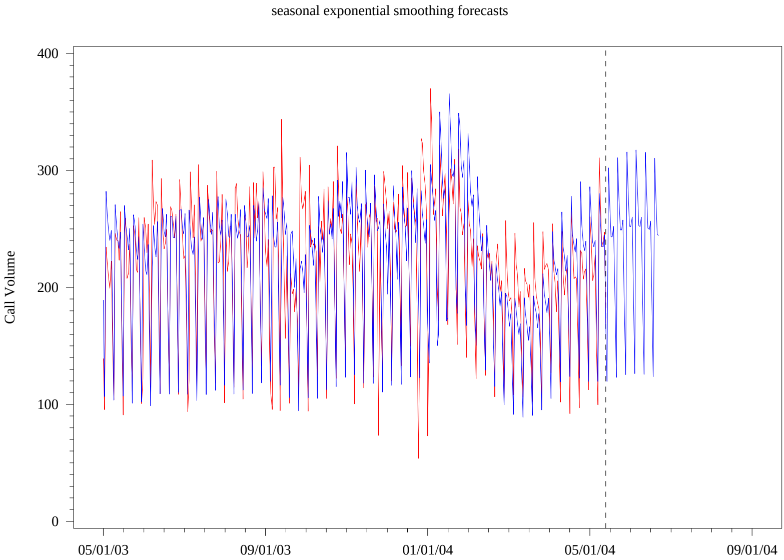
seasonal exponential smoothing forecasts
Call Volume
400
300
200
100
0
05/01/03
09/01/03
01/01/04
05/01/04
09/01/04

Exponential Smoothing VI

Advantages:

- Gives interpretable results (trend + seasonality).
- Gives more weight to recent observations.

Disadvantages:

- Not a model (in the statistical sense).
  - Prediction intervals not (really) possible.
- Can't generalize to multivariate approach.

ARIMA I

- Stands for *AutoRegressive Integrated Moving Average* models.
- Also known as Box-Jenkins models (Box and Jenkins, 1970).
- Advantages:
  - Best fit (minimum mean squared forecast error).
  - Generalizes to multivariate approach.
  - Often used in *statistical* practice.
- Disadvantages:
  - More complex.
  - Not intuitive *at all*.

ARIMA II

Assume nonseasonality for now.

- First, transform, then difference the data $\{Y_t\}$ d times until it is stationary (constant mean, variance), denoted $\{Y_t^*\}$.
- Guesstimate orders p, q through the sample autocorrelation, partial autocorrelation functions.
- Fit an *autoregressive moving average* (ARMA) process, orders p and q:

$$\begin{aligned} Y_t^* - \phi_1 Y_{t-1}^* - \cdots - \phi_p Y_{t-p}^* &= Z_t + \theta_1 Z_{t-1} + \cdots + \theta_q Z_{t-q} \\ \phi(Y_t^*) &= \theta(Z_t) \end{aligned}$$

  where $Z_t \overset{iid}{\sim} N(0, \sigma^2)$, and $\phi_1, \ldots, \phi_p$, $\theta_1, \ldots, \theta_q$ are constants.
- Through trial and error, repeat above 2 steps until errors “look good”.

Above is an ARIMA(p, d, q) model.

Confused Yet?

Q: How do we account for seasonality, period s?

A: We do almost the exact same thing, except for period s:

- Look at $\{Y_t^*, Y_{t+s}^*, Y_{t+2s}^*, \ldots\}$. Are they stationary? If not, difference D times until they are.
- Guesstimate orders P and Q similarly to before.
- Fit "multiplicative ARMA(P, Q)" process, period s:

$$\left(Y_t^* - \Phi_1 Y_{t-s}^* - \cdots - \Phi_P Y_{t-Ps}^*\right) \phi(Y_t^*) = (Z_t + \Theta_1 Z_{t-s} + \cdots + \Theta_Q Z_{t-Qs})\, \theta(Z_t)$$

- Repeat above 2 steps until all "looks good".

Above is an ARIMA(p, d, q)(P, D, Q) process.

SAS Code

If you're still with me ...

Forecasting 3 weeks ahead, ARIMA

```
PROC ARIMA data=mydata;
  IDENTIFY var=calls( 1, 7 ) noprint;
  ESTIMATE p=( 8 ) q=( 1, 5 )( 6, 7 ) noconstant method=ML noprint;
  FORECAST lead=21 out=forearima id=date interval=day noprint;
RUN;
QUIT;
```

Compare with Moving Average

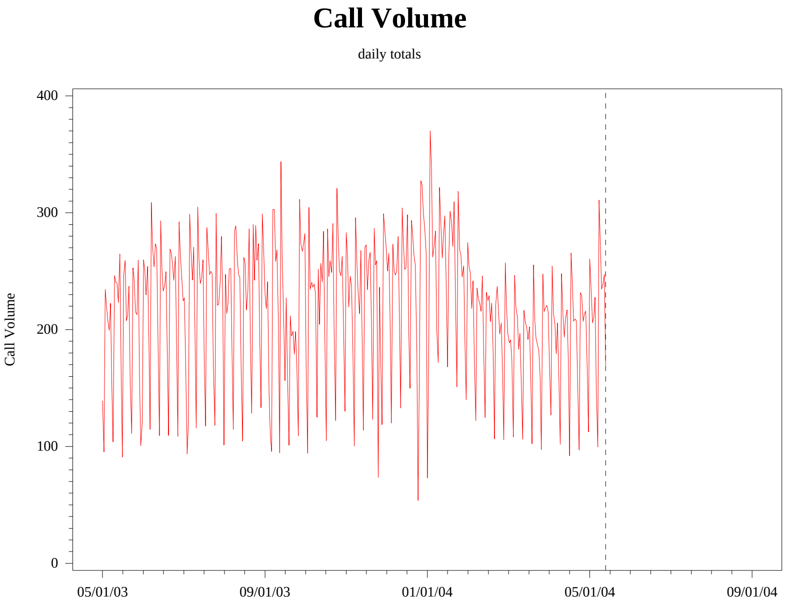
Call Volume
daily totals
Call Volume
400
300
200
100
0
05/01/03
09/01/03
01/01/04
05/01/04
09/01/04

Call Volume


ARIMA forecasts
Call Volume
400
300
200
100
0
05/01/03
09/01/03
01/01/04
05/01/04
09/01/04

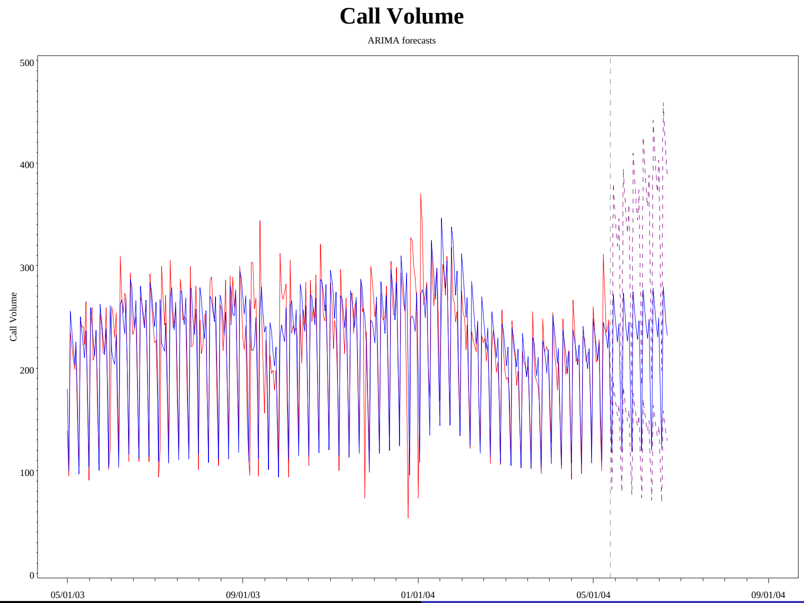
Call Volume
ARIMA forecasts
Call Volume
500
400
300
200
100
0
05/01/03
09/01/03
01/01/04
05/01/04
09/01/04

Which Method Should be Used?

We used three methods, would like to try others later.

Q: Which method should be used?

- Idea: The one that makes the best forecasts!
- For each call category,
  - Make k-week-ahead forecasts for the last n weeks of the data.
    - For $i = 1, \ldots, n$, remove last i weeks of the data, then make forecasts for k weeks in the future.
  - For each method, compare forecasts to actuals.
  - Use forecasts from the method that made the most accurate forecasts.

How Do We Judge Forecasts?

- General standard: Mean Absolute Prediction Error (MAPE):

$$\text{MAPE} = 100 \times \sum_{t=1}^{T} \frac{|\text{forecast}_t - \text{actual}_t|}{\text{actual}_t},$$

Gives average percentage off (zero is best!).

- Our forecasts over 10 weeks before 5/14/04:

Horizon	Mov Average	Exp Smoothing	ARIMA
1-week-ahead	10.98%	**6.36%**	6.51%
3-week-ahead	12.83%	9.53%	**7.77%**

- Sometimes different methods best for different horizons!
- We want the 3-week-ahead criterion.

How Do We Do This with SAS?

Easy way: Purchase SAS High Performance Forecasting!

- Follows (and generalizeds) our framework.
- Implements our methods.
- Allows us to add our own methods.

Harder (but cheaper) way: Program it ourselves.

How Do We Do This with SAS?

SAS Code Excerpt

```
DATA results;
  SET all;     *merged results, sorted by method;
  ape3 = 100*abs( calls - forecast3 )/calls;

PROC MEANS data=results noprint;
  BY method;
  VAR ape3;
  OUTPUT OUT=summ MEAN( ape3 ) = mape3 / noinherit;

DATA skmape;
  SET summ;
  IF method = 'arima' THEN CALL SYMPUT( 'mapearima', mape3 );
  IF method = 'exsm' THEN CALL SYMPUT( 'mapeexp', mape3 );
  IF method = 'mave' THEN CALL SYMPUT( 'mapemave', mape3 );

%LET mapev = &mapearima, &mapeexp, &mapemave;

DATA _null_;
  IF MIN( &mapev ) = &mapearima THEN CALL SYMPUT( 'best', 'arima' );
    ELSE IF MIN( &mapev ) = &mapeexp THEN CALL SYMPUT( 'best', 'exsm' );
    ELSE IF MIN( &mapev ) = &mapemave THEN CALL SYMPUT( 'best', 'mave' );

DATA ftool.bestcalls;
  SET &best.calls;
RUN;
```

Are Our Overall Forecasts Better?

- Better forecasts in training set no guarantee of better forecasts overall!
- Happily, we *do* get better forecasts in general.

1-Week-Ahead MAPEs

Over last 10 weeks up to 5/14/04:

Call Category	Analyst	Forecast Tool	Difference (A - FT)
1	40.61%	5.71%	34.91%
2	16.61%	9.88%	6.73%
3	3.94%	3.54%	0.40%
4	35.84%	13.87%	21.98%
5	14.24%	11.94%	2.30%
6	9.97%	9.68%	0.29%
7	7.78%	6.74%	1.04%
Weighted Average			**24.74%**

3-Week-Ahead MAPEs

Over last 10 weeks up to 5/14/04:

Call Category	Analyst	Forecast Tool	Difference (A - FT)
1	44.36%	7.83%	36.52%
2	19.08%	14.70%	4.39%
3	3.80%	4.22%	-0.42%
4	25.57%	15.06%	10.51%
5	16.78%	14.82%	1.96%
6	13.69%	10.88%	2.81%
7	18.75%	9.81%	8.94%
Weighted Average			**25.44%**

What's Next?

Multivariate Models!

- Takes account of holidays/other irregularities.
- Allows for scenario forecasting!

How will we do this?

How Will We Do This?

One solution: Multivariate ARIMA (transfer models):

$$Y_t = \beta_0 + \sum_{i=0}^{I} \beta_i X_{t-i} + Z_t, \quad Z_t = \text{ARIMA process}$$

- Works all right (using `PROC ARIMA`), but
- Very complicated to use,
- Results not very good/useful!

One big problem: Parameters are fixed over time.

- One outlier (e.g., Sept 11) could screw up entire model.
- If parameters could change over time, model would be (much) more flexible.

How Will We Do This?

Another solution: *State Space* (or *Hidden Markov*) Models

$$Y_t = \beta_{0t} + \sum_{i=0}^{I} \beta_{it} X_{t-i} + Z_t, \quad Z_t = \text{Normal process}$$

- Parameters change (slowly) over time.
  - Modeled by separate equation.
- Complicated, but flexibility makes it worth it.
- Problem: SAS doesn't implement it!
  - `PROC STATESPACE`: Nope! (misleading name)
  - `PROC UCM`: Closer, but still not there.
  - `PROC IML`: Can do it, but a fair bit of work.
  - (Almost) no one else (R, S+, SPSS) does, either.
  - Guess what my PhD project is!

Further Resources

John C. Brocklebank and David A. Dickey.
SAS for Forecasting Time Series.
SAS Institute, 2003.

Chris Chatfield.
Time-Series Foreasting.
Chapman and Hall, 2000.

Nate Derby: `http://nderby.org`

`nderby@sprodata.com`