Automatic Statistical Forecasting of Univariate Time Series

Use in Supply Chain Management (SCM) and Comparison of

Three Implementations

Marcel Baumgartner
marcel.baumgartner@nestle.com

Outline

- Demand Planning / Statistical Forecasting in the Food Industry
- ForecastPRO
- State-Space Model of Exponential Smoothing
- SAP *Advanced Planning and Optimization* (APO)
- Comparison / Conclusions

Make-to-Stock

The Food Industry uses the *Make-to-Stock* approach: due to industrial constraints, we need to manufacture our products in advance, and we cannot wait for the order of the customer.

Therefore, foreseeing future orders is paramount. We need to make sure that we have the

- right product,
- at the right location,
- at the right moment in time,
- with the right amount.

Demand Planning Process

At Nestlé, this planning process is known as *Consensus Demand Planning*.

The outcome of this process is a quantity per product, location and week (short term) / months (mid-term, up to 18 months), agreed upon by Sales, Marketing, Supply Chain and Finance functions, within the context of:

- high number or products
- high innovation/renovation rate
- need to forecast customer orders, and not real consumer demand
- promotion driven business (need to capture consequences of internal trade and marketing activities)
- many categories depend on weather situation

SAP APO Screen

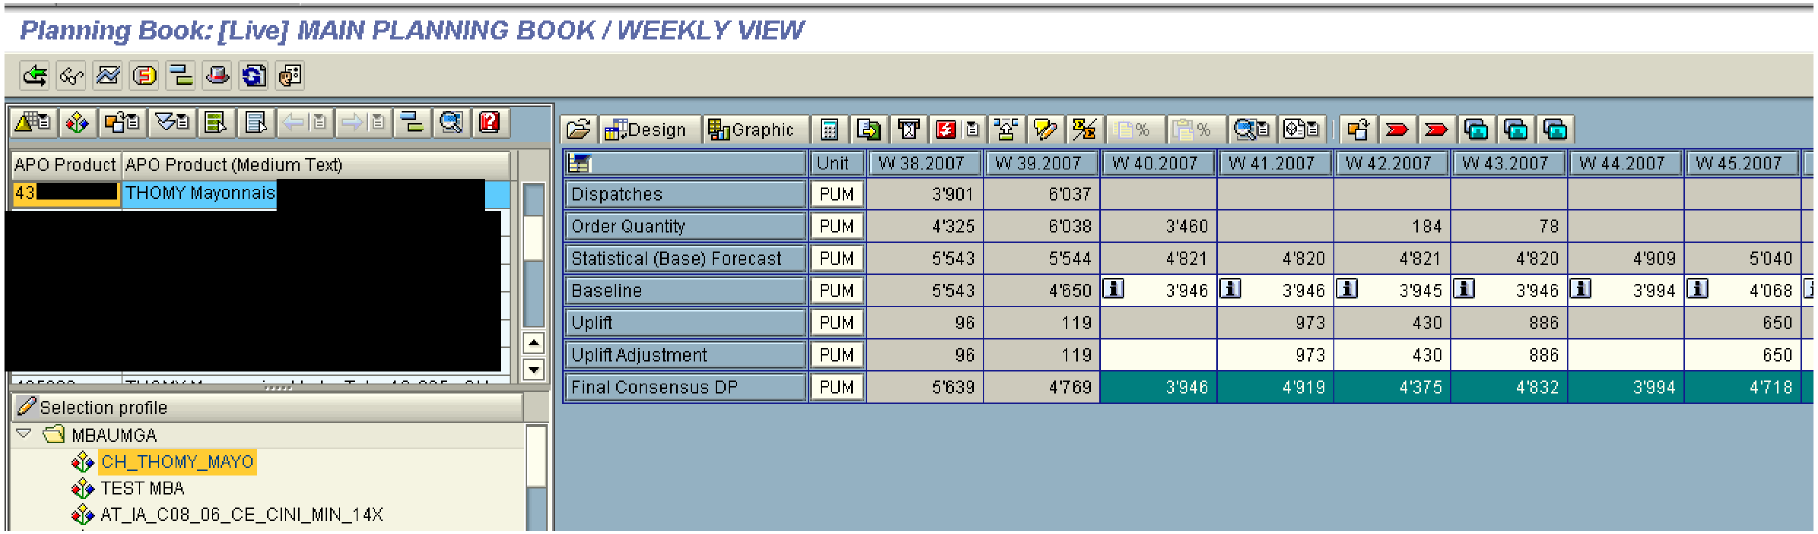

This is part of the screen used by a Demand Planner. He/She sees historical data, and can input/calculate future planned orders, per product, location and/or customer.

The Planner can also use *Statistical Forecasting* algorithms available in this SAP Module.

Examples of Time Series

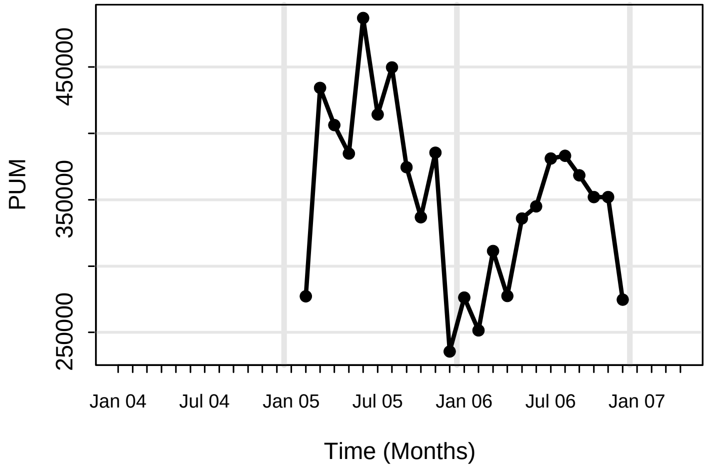

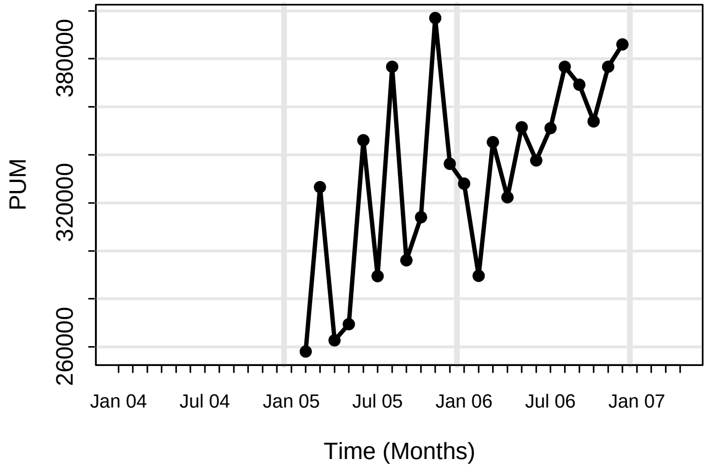

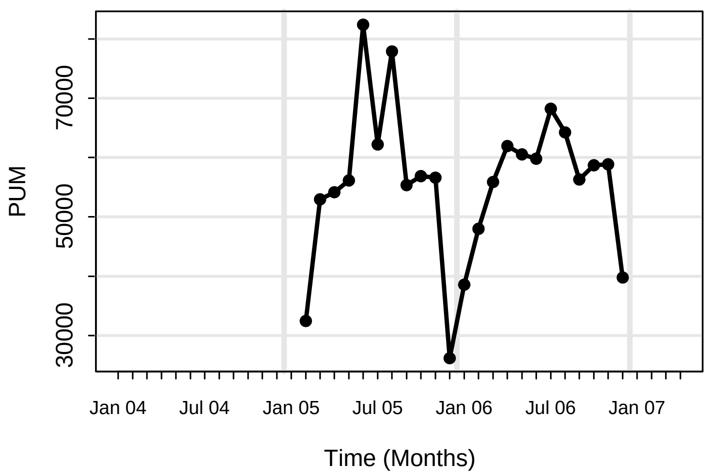

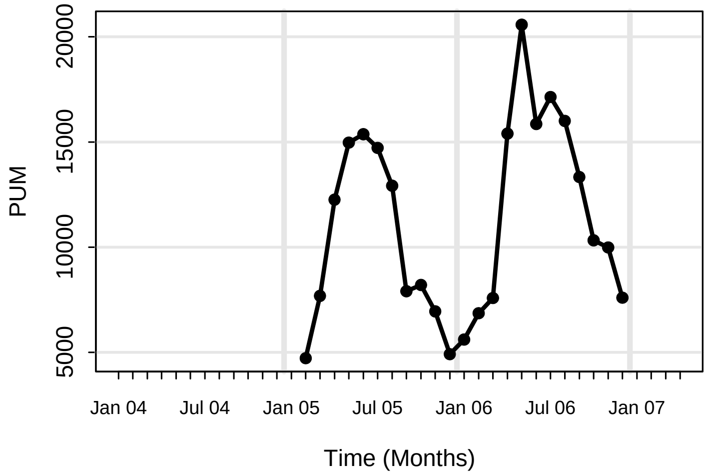

Examples of Time Series

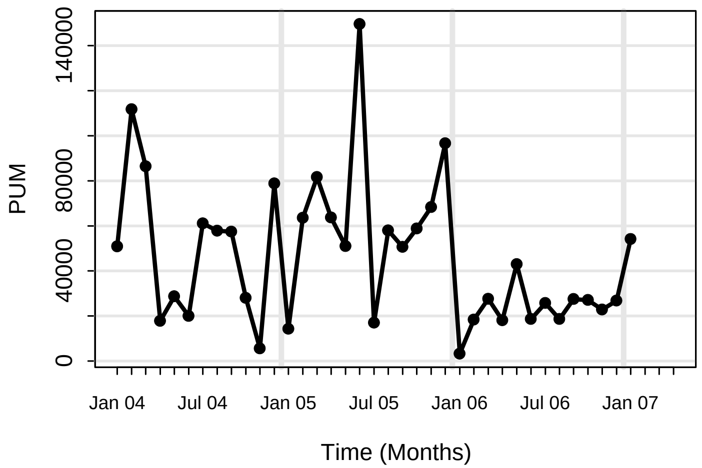

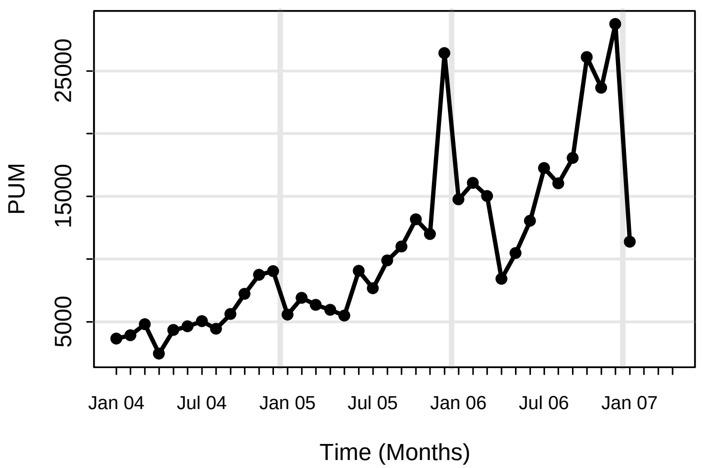

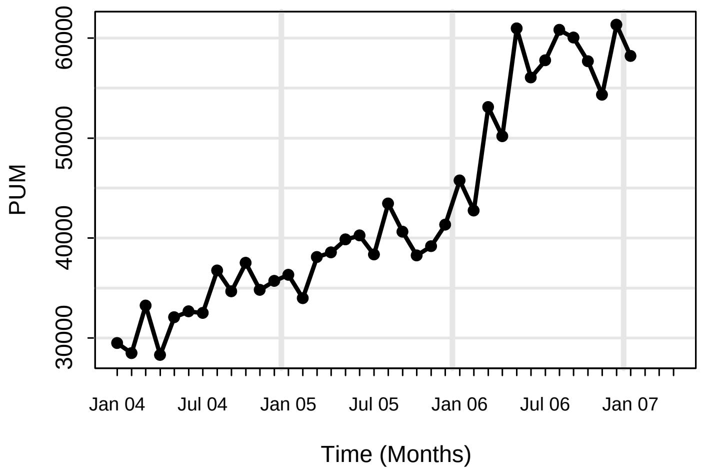

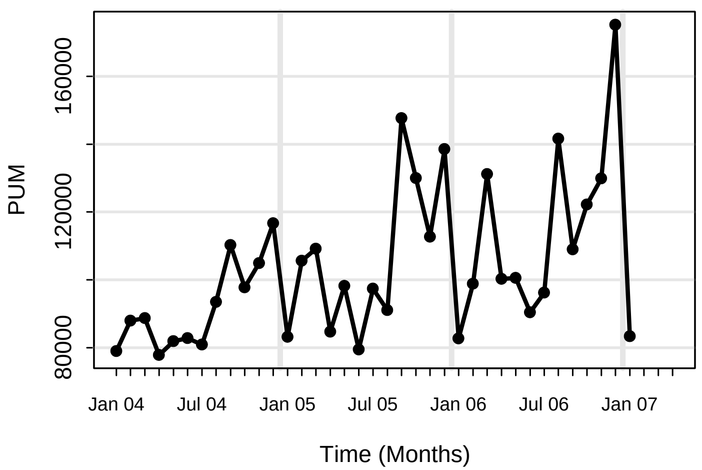

ForecastPRO

ForecastPRO http://www.forecastpro.com/

ForecastPRO is a renowed, best-of-breed statistical forecasting software. It has a proven track record, notably in the famous "M3 Forecasting Competition".

Its *Expert Selection* forecasting methods uses *Exponential Smoothing* and ARIMA type algorithms. It can be run as a full black-box type method.

Advanced users can also choose themselves the family of method and the smoothing parameters.

Layout of ForecastPRO

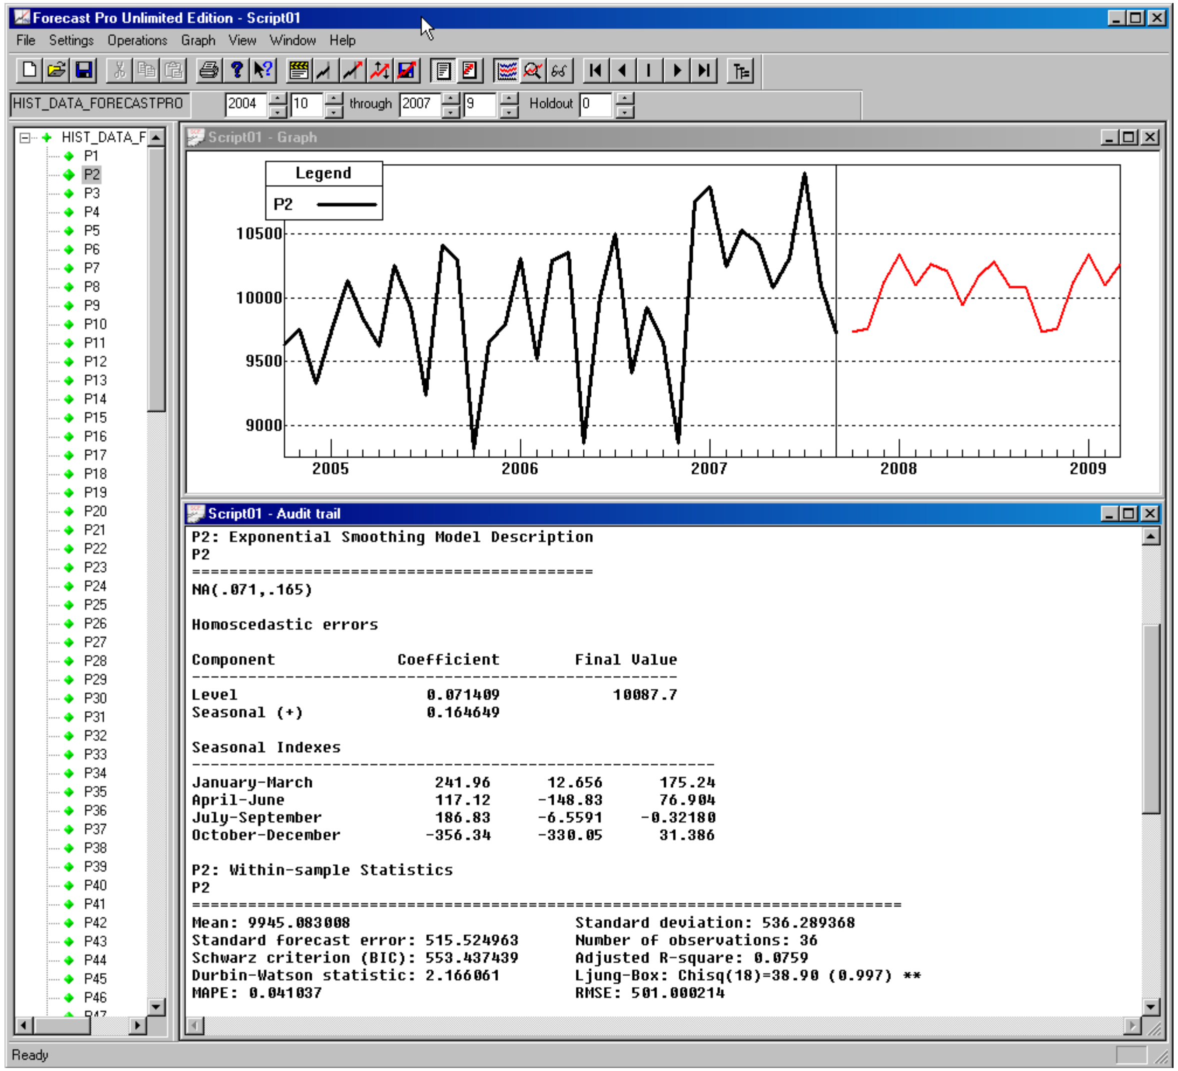

State Space Model of Exponential Smoothing

www.RobHyndman.info

This brand new approach is developed by Rob Hyndman, Monash University, Melbourne.

Rob Hyndman is co-author of "Business Forecasting: Methods and Applications" by Makridakis et al., and the lead author of a forthcoming book on Exponential Smoothing.

Hyndman has developed a new framework for *Exponential Smoothing* based on *State-Space Models*. This allowed him to develop a much more robust approach to parameter fitting, and using Akaike type error measurements to estimate the "best" models.

His algorithm can also be run fully automatically. It is available in the package 'forecast' for R.

ETS() in R

```
> x.ts
        Jan   Feb   Mar   Apr   May   Jun   Jul   Aug   Sep   Oct   Nov   Dec
2004                                                            11338  9668  9337
2005   8010 10016  9736  9656  9124  9389  9128 10064  9854  9436  9121  8990
2006  10089  9236  9660  9628  8659  8503  9417  9994  8729  8746  9333  8649
2007   9148  8287  9317  9213  9262  8728  8654  8849  9539
> x.ets <- ets(x.ts)
> x.ets
ETS(M,M,N)

Call:
 ets(y = x.ts)

  Smoothing parameters:
           alpha = 0.01
           beta  = 0.01

         Initial states:
           l = 9863.2379
           b = 0.9965

         sigma:  0.0606
         AIC:    591.9277
         AICc:   593.218
         BIC:    598.2618
> forecast(x.ets,h=18)
         Point Forecast    Lo 80    Hi 80    Lo 95    Hi 95
Oct 2007       8749.078 8076.056 9468.787 7697.493 9828.300
Nov 2007       8736.517 8080.797 9406.716 7698.910 9748.150
Dec 2007       8723.973 8044.222 9410.581 7690.806 9795.340
Jan 2008       8711.448 8014.593 9389.144 7694.054 9719.676
Feb 2008       8698.941 8015.549 9377.469 7643.466 9698.905
Mar 2008       8686.452 8000.297 9342.348 7609.226 9717.750
Apr 2008       8673.981 8006.143 9357.990 7671.909 9699.004
May 2008       8661.528 7979.596 9338.303 7599.730 9669.771
Jun 2008       8649.092 7961.689 9328.927 7639.438 9705.380
Jul 2008       8636.675 7966.255 9331.372 7594.205 9698.840
Aug 2008       8624.275 7929.004 9324.538 7562.499 9691.458
Sep 2008       8611.893 7941.915 9318.171 7575.280 9678.965
Oct 2008       8599.529 7904.413 9310.827 7533.724 9657.084
Nov 2008       8587.182 7866.665 9272.525 7539.509 9642.586
Dec 2008       8574.854 7878.827 9264.928 7527.169 9671.179
Jan 2009       8562.543 7853.625 9272.497 7484.743 9657.017
Feb 2009       8550.250 7833.811 9262.454 7469.180 9665.408
Mar 2009       8537.974 7799.030 9287.598 7439.766 9679.611
>
```

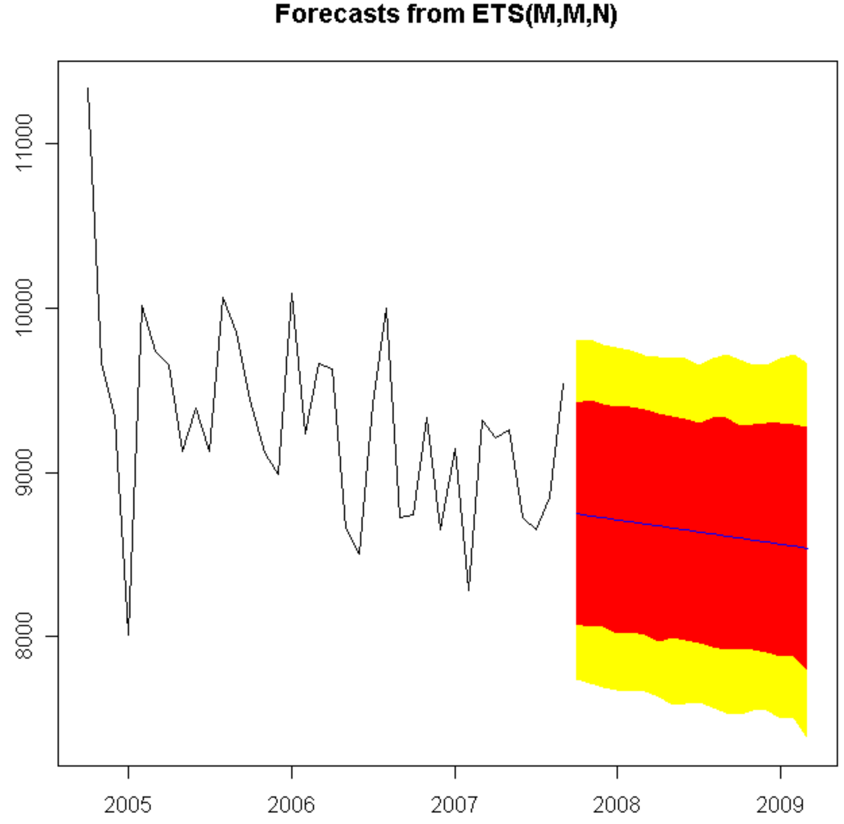

SAP APO DP

SAP is one of the biggest business software provider. Nestlé supports its processes and best practices with SAP modules, in all kinds of areas, like Finance, Sales, Human Resources, Manufacturing, but also *Supply Chain Management* (SCM).

The main module for SCM is called APO (Advanced Planning and Optimization). A sub-module is related to Demand Planning. It contains statistical forecasting algorithms, essentially based on *Exponential Smoothing* and *Linear Regression*.

APO DP has a fully automated method as well, but its architecture is quite different compared to Hyndman's approach or the one from ForecastPRO.

Logic of APO's Automatic Method

APO uses three statistical tests to define the methods to be estimated.

The trend test is based on the *Confidence Interval* of the estimated slope of a linear regression. The seasonal test is based on the autocorrelation coefficient for the period of the data (e.g. 12 month seasonality). The white noise test is the *Box-Ljung Portmanteu* test, using the χ^2 distribution.

Trend Test	Seasonal Test	White Noise Test	Constant	Trend	Season, No Trend	Season and Trend	Linear Regression	Seasonal Linear Regression
-	-	-	X					
+	-	-	X	X			X	
-	+	-	X		X			
+	+	-	X	X		X	X	X
-	-	+	X					
+	-	+	X					
-	+	+	X					
+	+	+	X					

Overview of Methods

ForecastPRO	www.RobHyndman.info	SAP
Best-of-breed statistical forecasting software. Internally developed expert selection algorithm, without any detailed documentation. Proven track record in forecast competitions. Easy to use. Extremely fast !	Brand new approach, "only" available in R. Based on a solid statistical foundation, using modern model selection criteria. Algorithm documented in details. Seems to compare well in the classic M3 competition. Execution is slow, but easy to use and to parametrize.	Simple and straightforward implementation of exponential smoothing, with an automatic method based on statistical tests. Well documented, not as fast as ForecastPRO, but faster than the state-space model. Has not been tested in forecasting competitions. Well integrated in a business software for Demand Planning.

Simulation Study

In order to study the quality of the forecasts of these three methods, and also how differently they behave, we have simulated time series.

We vary 3 types of seasonality and 5 types of trend. This results in an experimental design of 15 combinations. For each combination, we generated 10 samples (different random numbers).

Examples of Time Series

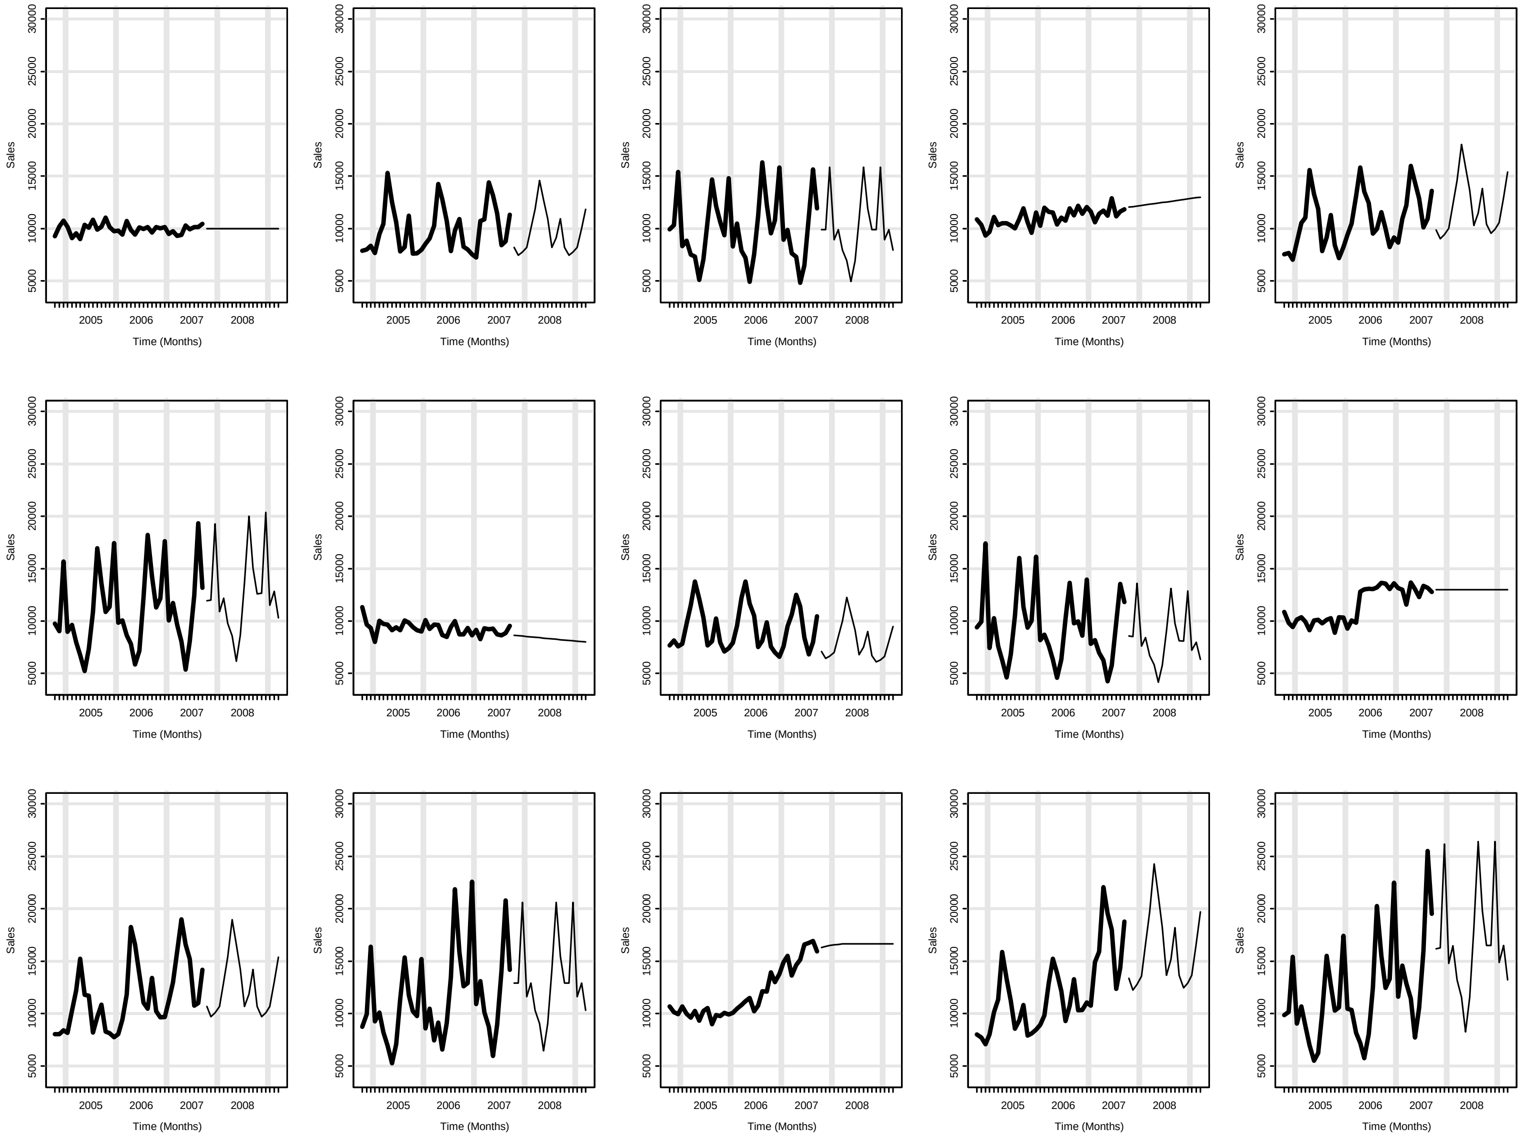

Simulation Results (I)

To measure the quality of the forecasts, we compare the estimates of the 3 methods with the "true" future. We then compute an error rate:

$$Error = 100 \cdot \frac{|Forecast - Truth|}{Truth}$$

We simply sum the next 3 months. Here are the average scores for each method:

APO	FPRO	Hyndman
13.983	13.053	10.508

Hyndman's method seems to slightly outperform the two other methods.

Simulation Results (II)

, , **No Season**

	No Trend	Upwards	Downwards	Shift	Exponential
APO	5.5	7.4	10.0	13.1	15.9
FPRO	4.0	5.6	12.2	9.5	13.7
Hyndman	2.6	5.1	7.6	7.9	13.1

, , **Summer-Winter**

	No Trend	Upwards	Downwards	Shift	Exponential
APO	9.9	14.2	13.7	16.1	31.5
FPRO	8.1	13.6	13.4	10.6	28.8
Hyndman	8.2	10.5	11.4	17.2	14.6

, , **Easter-Christmas**

	No Trend	Upwards	Downwards	Shift	Exponential
APO	10.6	8.7	12.7	10.8	29.6
FPRO	10.9	9.7	16.1	16.0	23.6
Hyndman	8.3	7.8	9.4	10.2	23.6

Interpretation

, , **No Season**

	No Trend	Upwards	Downwards	Shift	Exponential
APO	5.5	7.4	10.0	13.1	15.9
FPRO	4.0	5.6	12.2	9.5	13.7
Hyndman	2.6	5.1	7.6	7.9	13.1

, , **Summer-Winter**

	No Trend	Upwards	Downwards	Shift	Exponential
APO	9.9	14.2	13.7	16.1	31.5
FPRO	8.1	13.6	13.4	10.6	28.8
Hyndman	8.2	10.5	11.4	17.2	14.6

, , **Easter-Christmas**

	No Trend	Upwards	Downwards	Shift	Exponential
APO	10.6	8.7	12.7	10.8	29.6
FPRO	10.9	9.7	16.1	16.0	23.6
Hyndman	8.3	7.8	9.4	10.2	23.6

Hyndman's method seems to perform consistently better, particularly for time series with trend components.

APO and ForecastPRO do have similar results on average, but their peformance varies with different trend patterns.

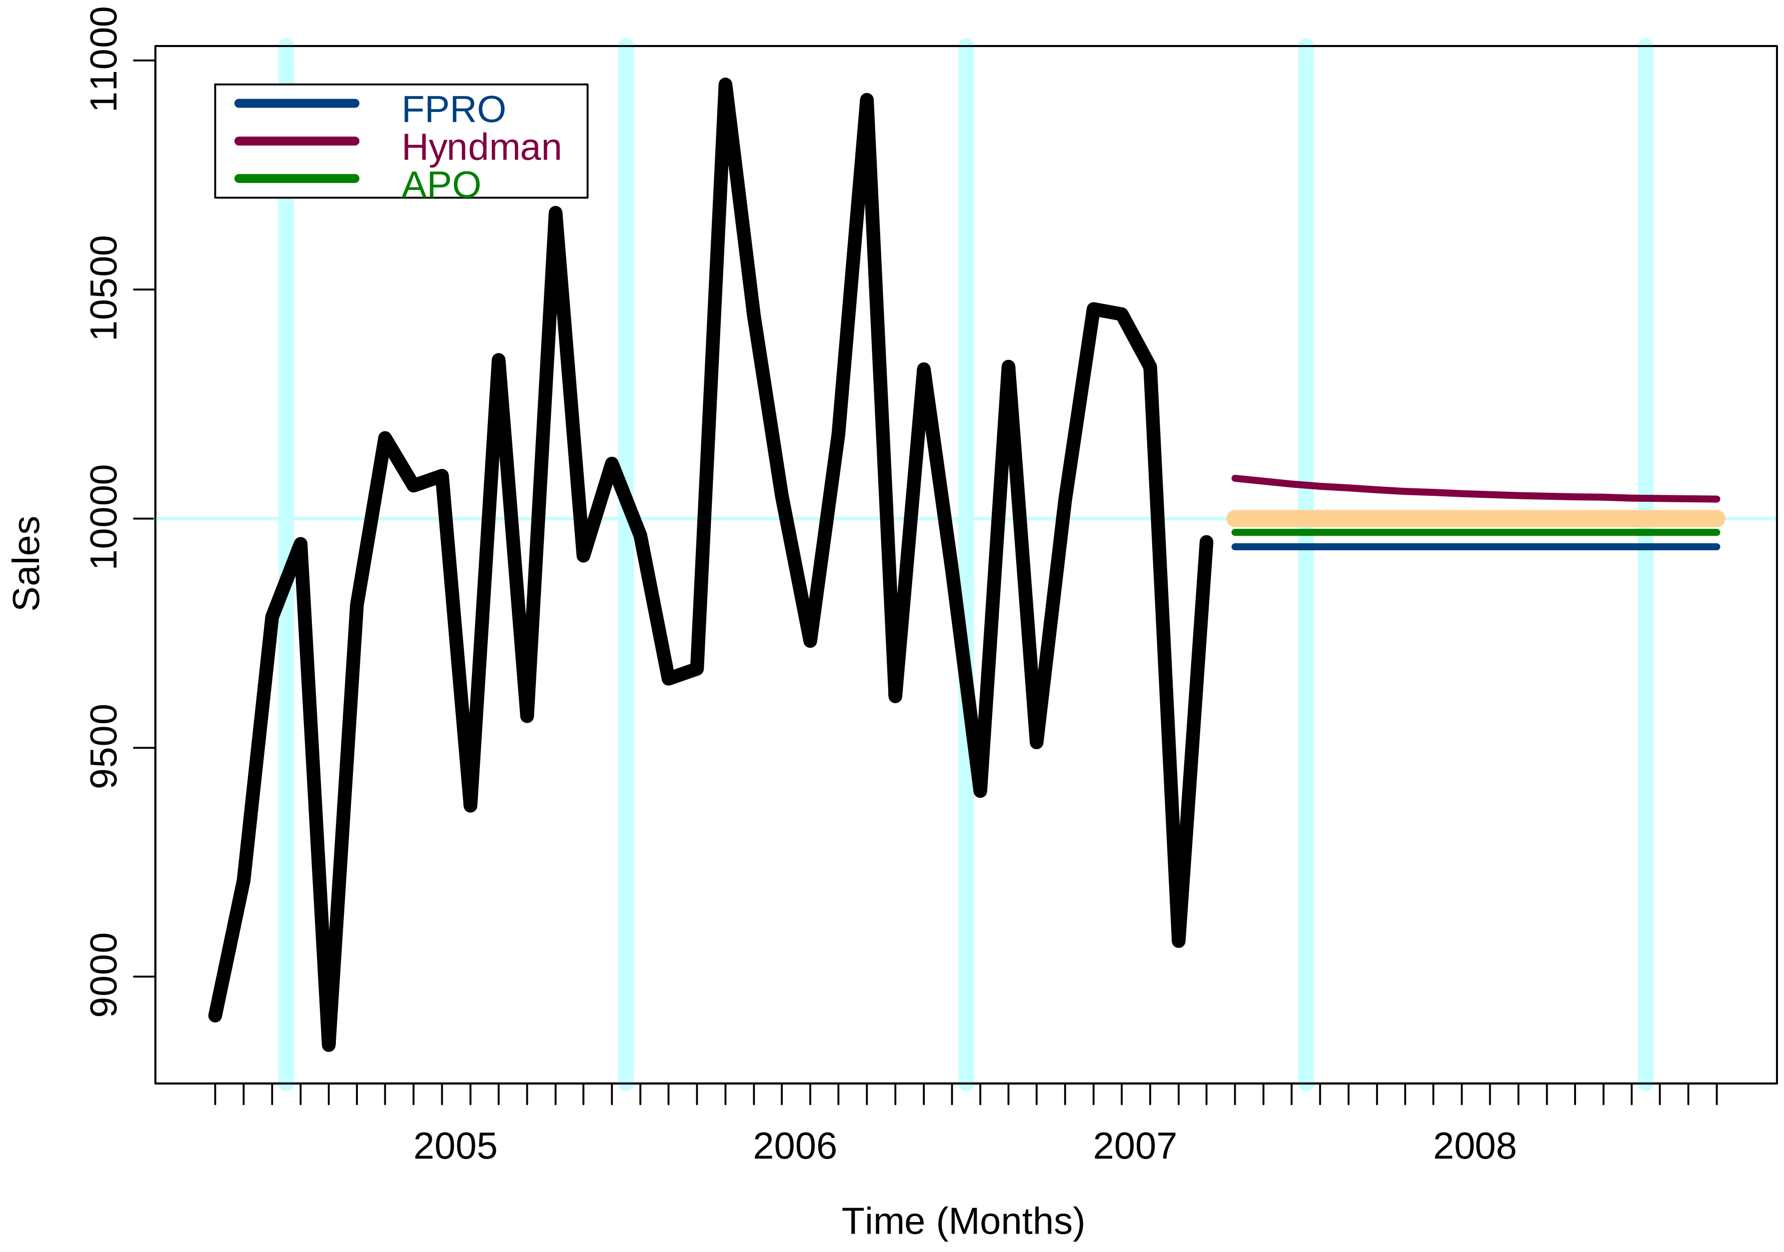
FPRO
Hyndman
APO
11000
10500
10000
9500
9000
Sales
2005
2006
2007
2008
Time (Months)

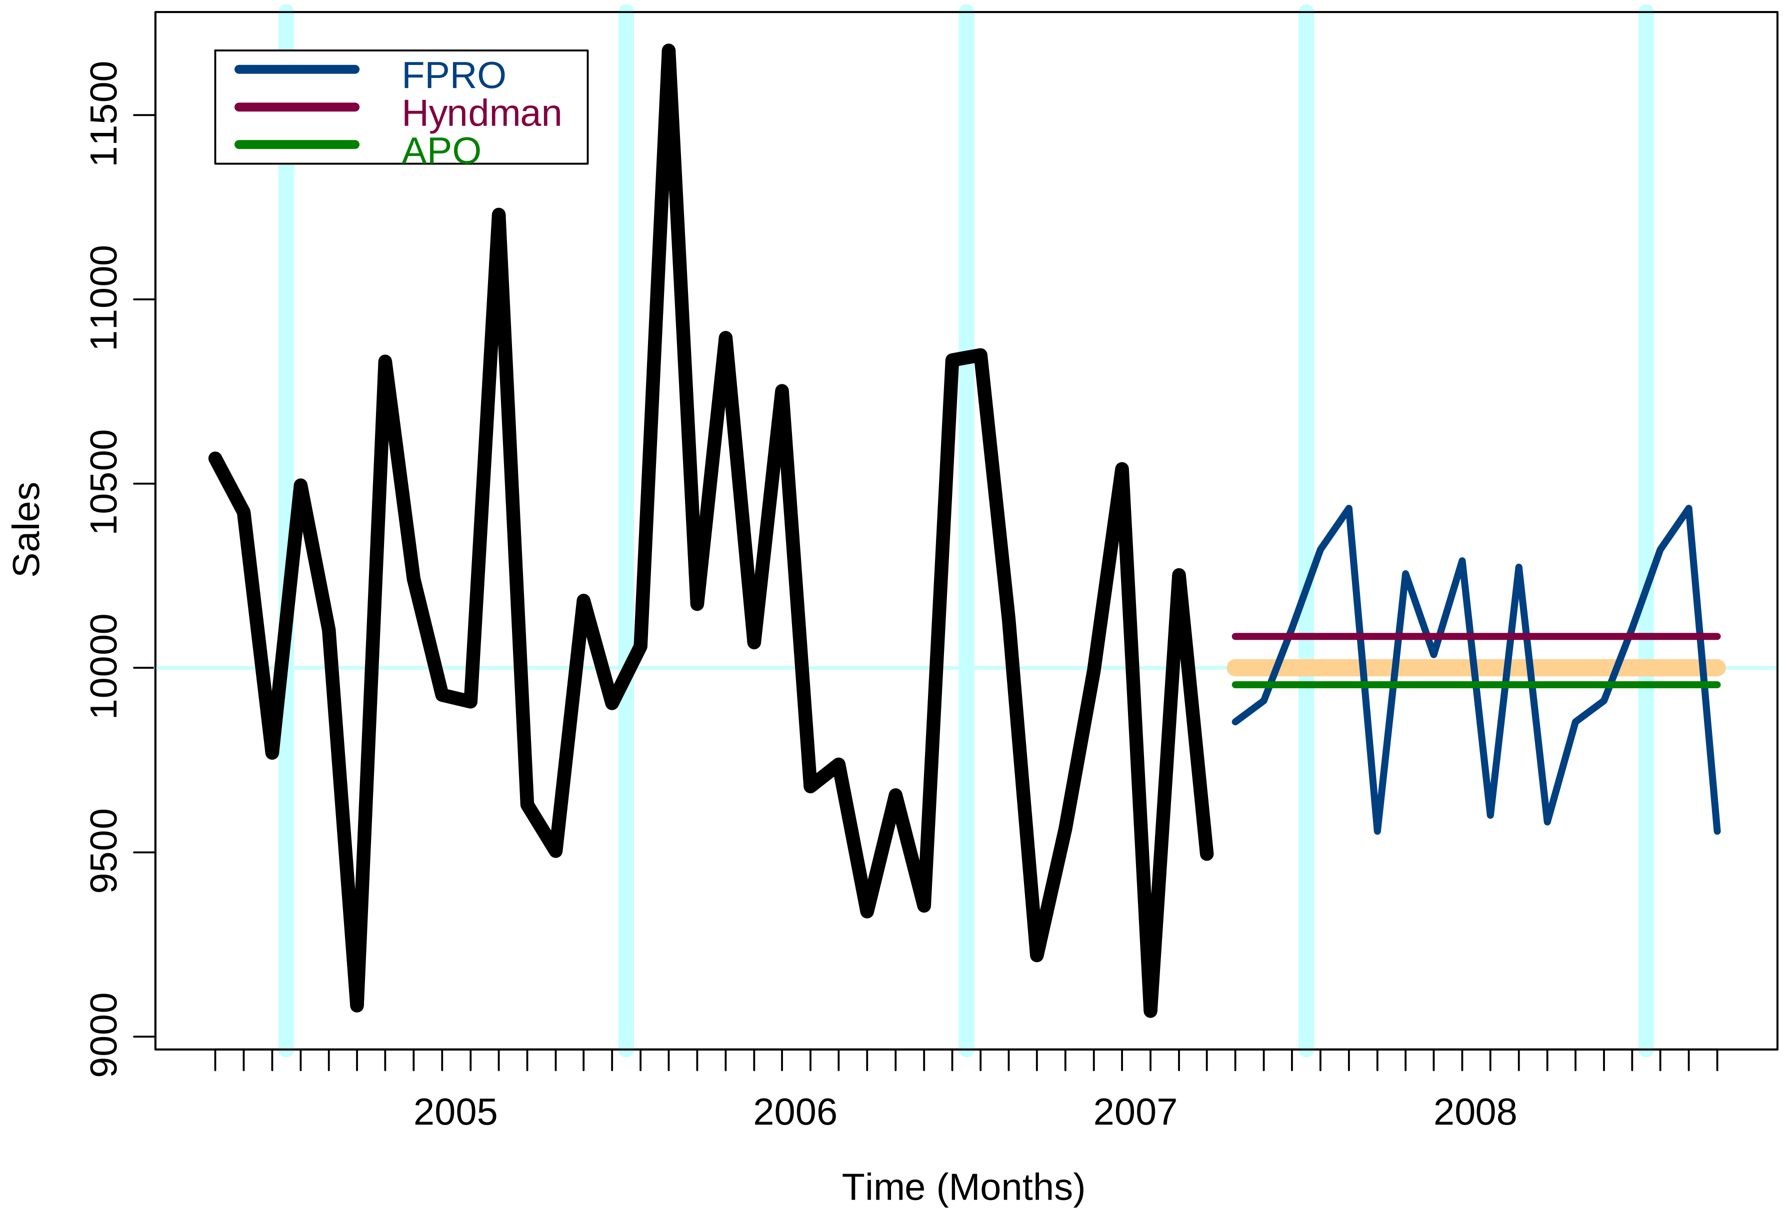

FPRO
Hyndman
APO
Sales
11500
11000
10500
10000
9500
9000
2005
2006
2007
2008
Time (Months)

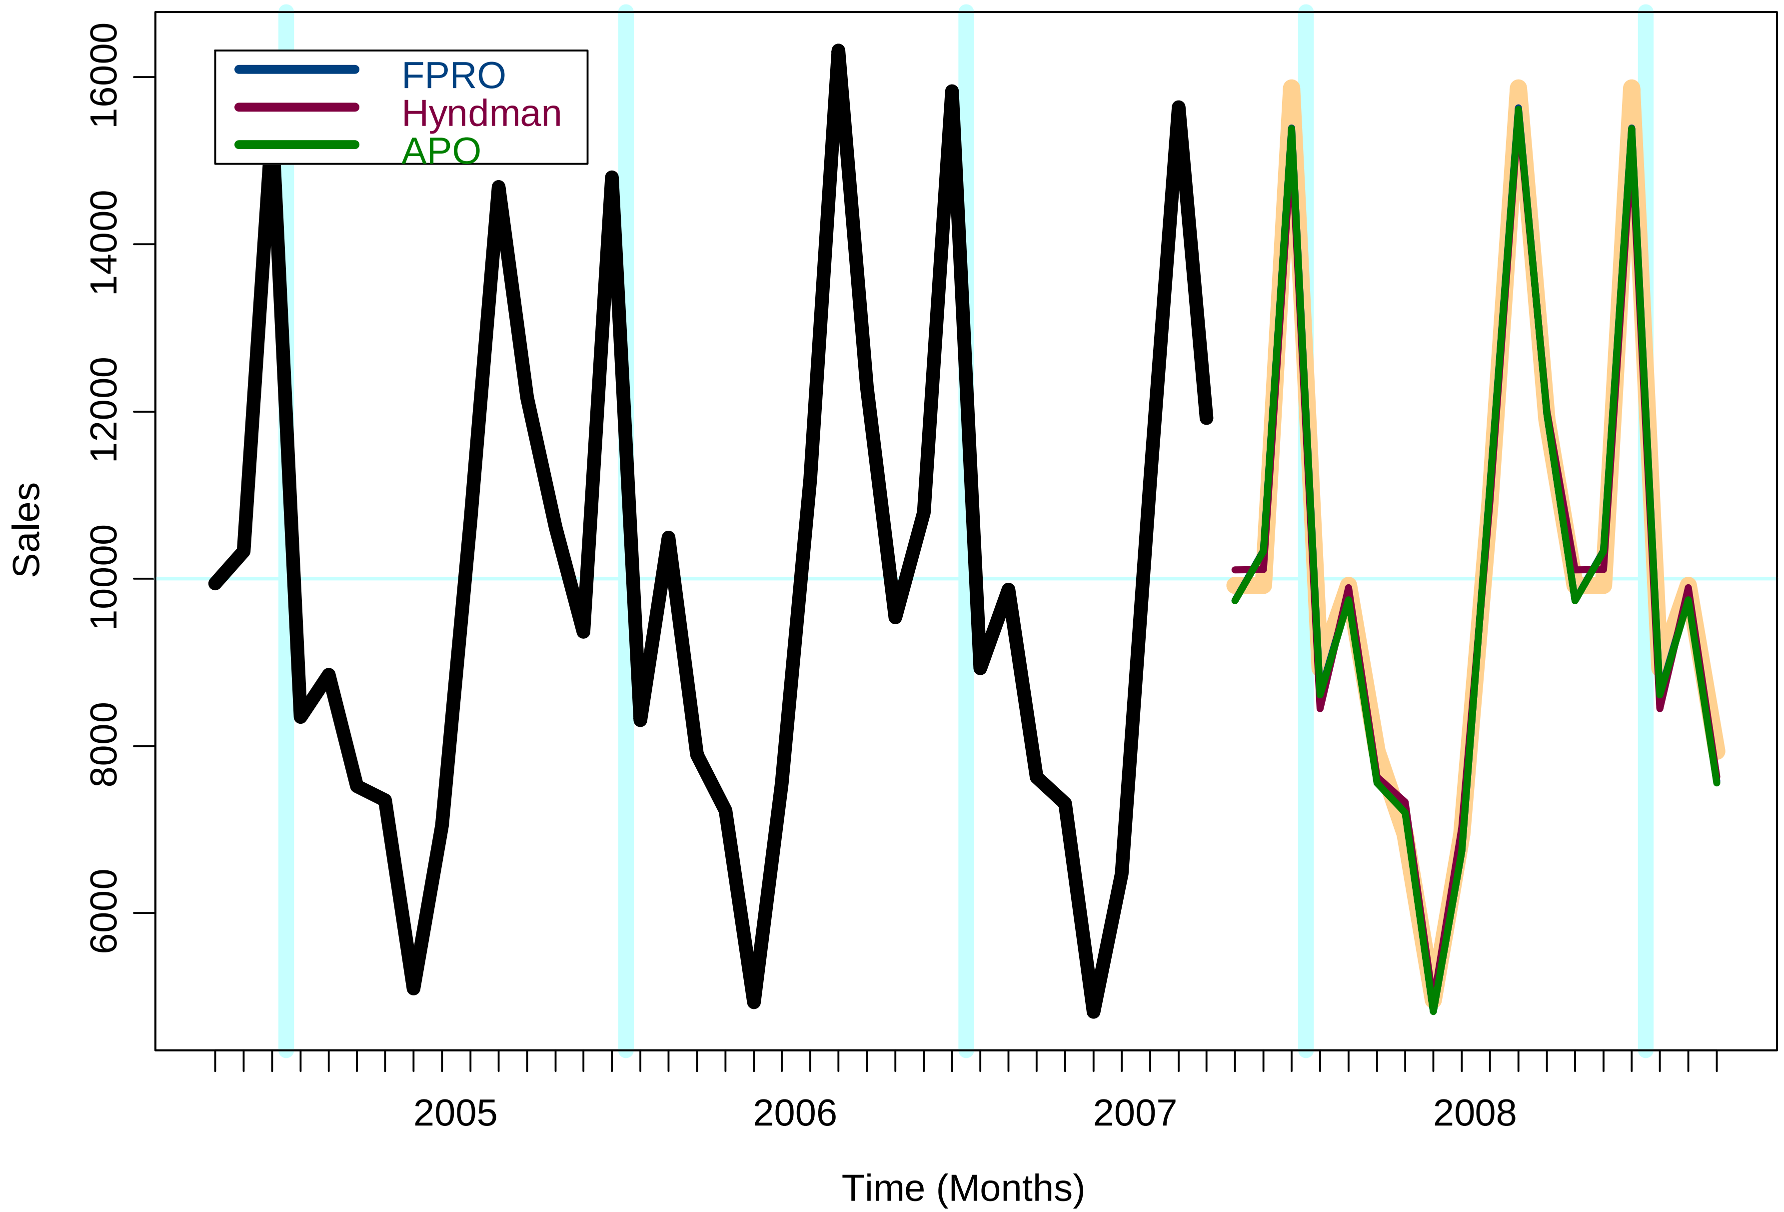
FPRO
Hyndman
APO
Sales
6000
8000
10000
12000
14000
16000
2005
2006
2007
2008
Time (Months)

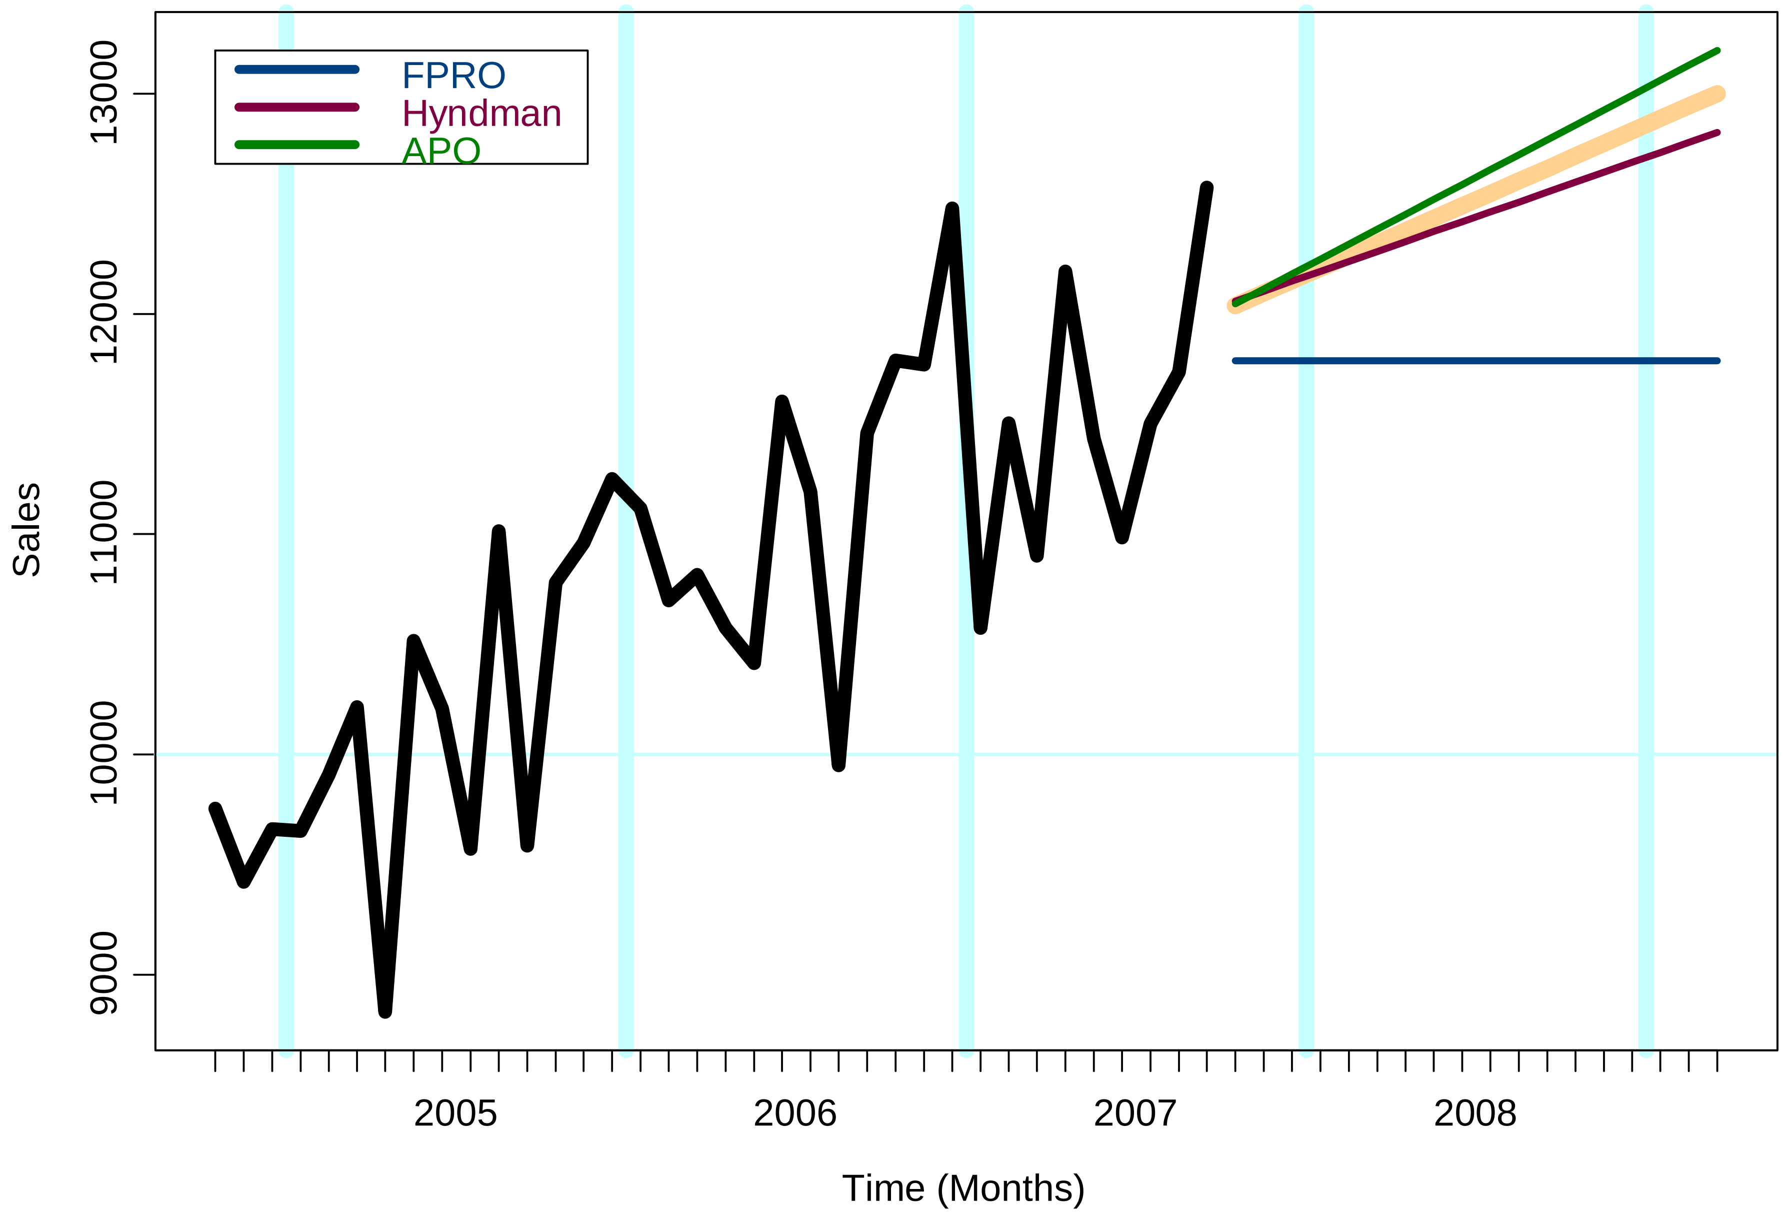
FPRO
Hyndman
APO
Sales
9000
10000
11000
12000
13000
2005
2006
2007
2008
Time (Months)

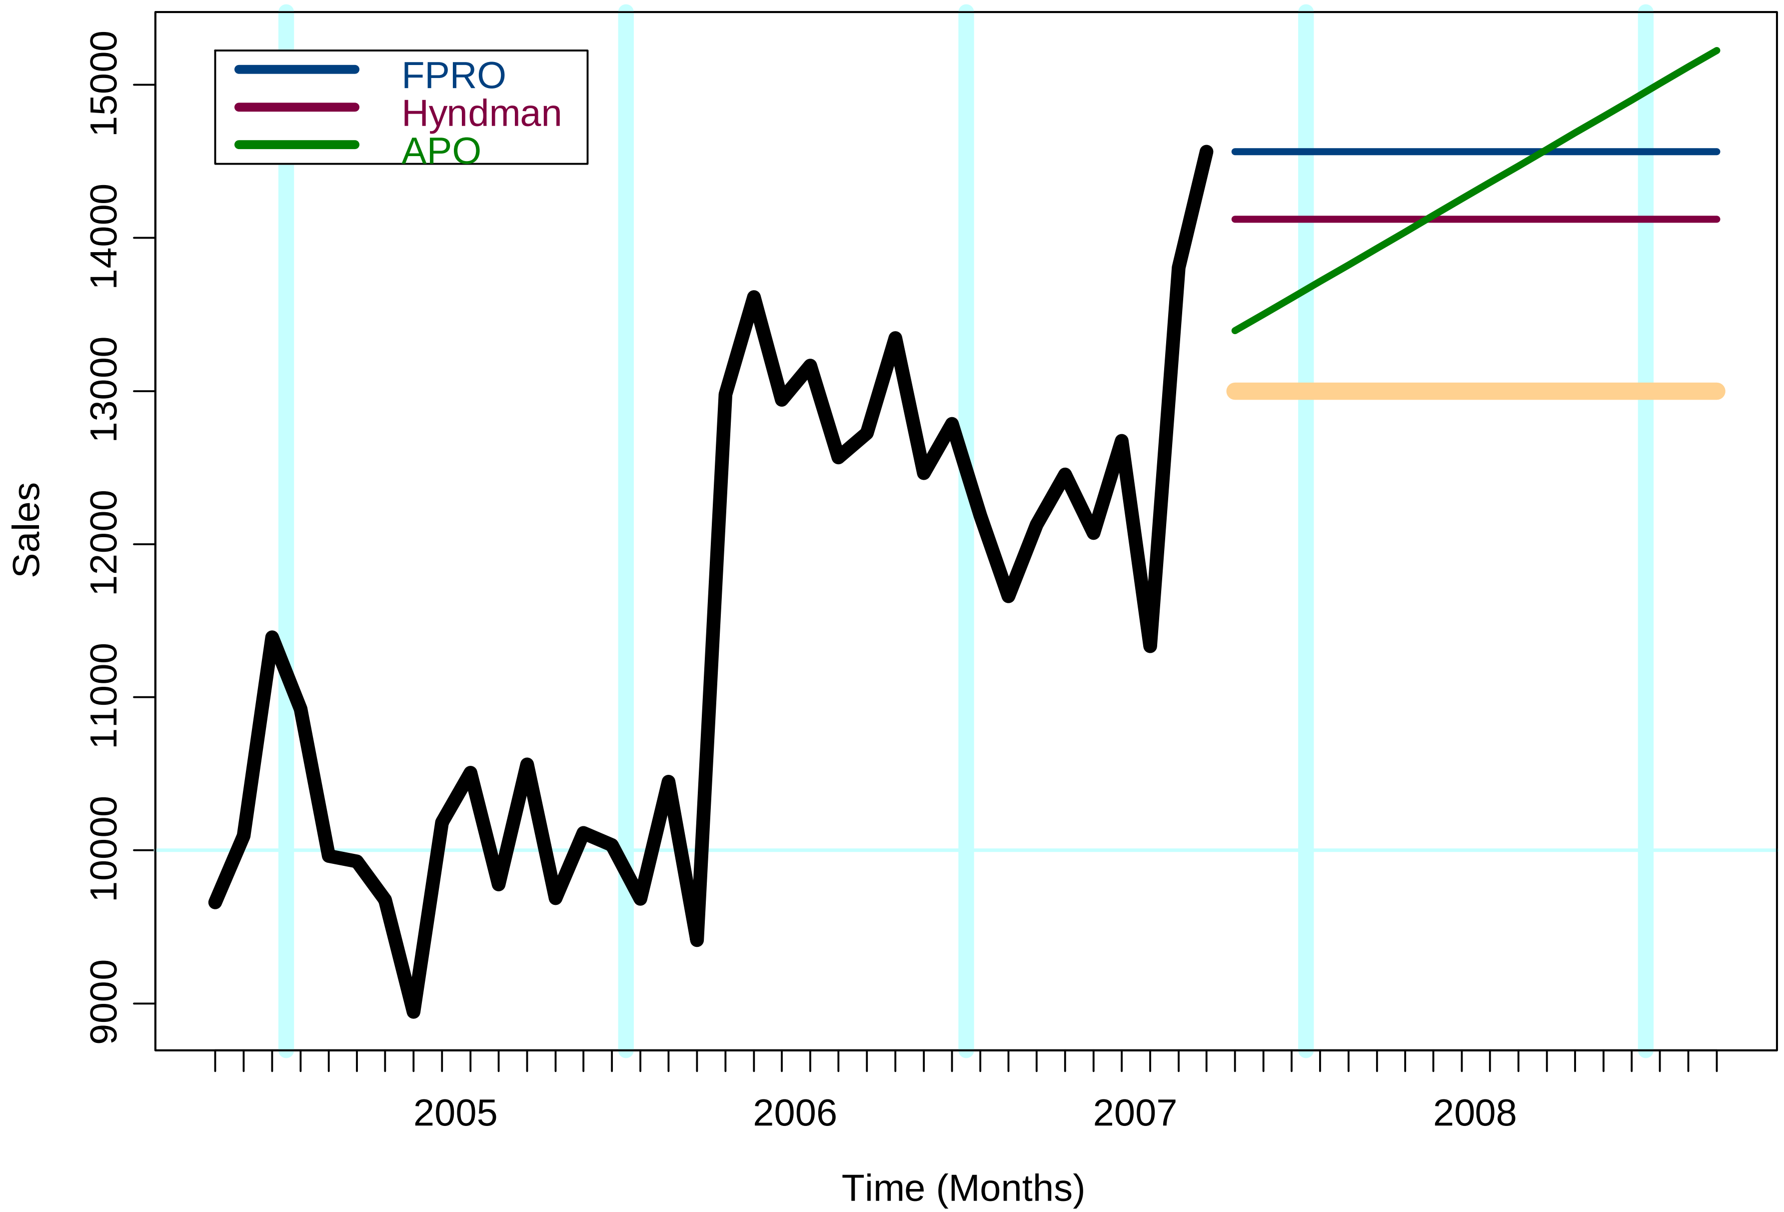

FPRO
Hyndman
APO
Sales
9000
10000
11000
12000
13000
14000
15000
2005
2006
2007
2008
Time (Months)

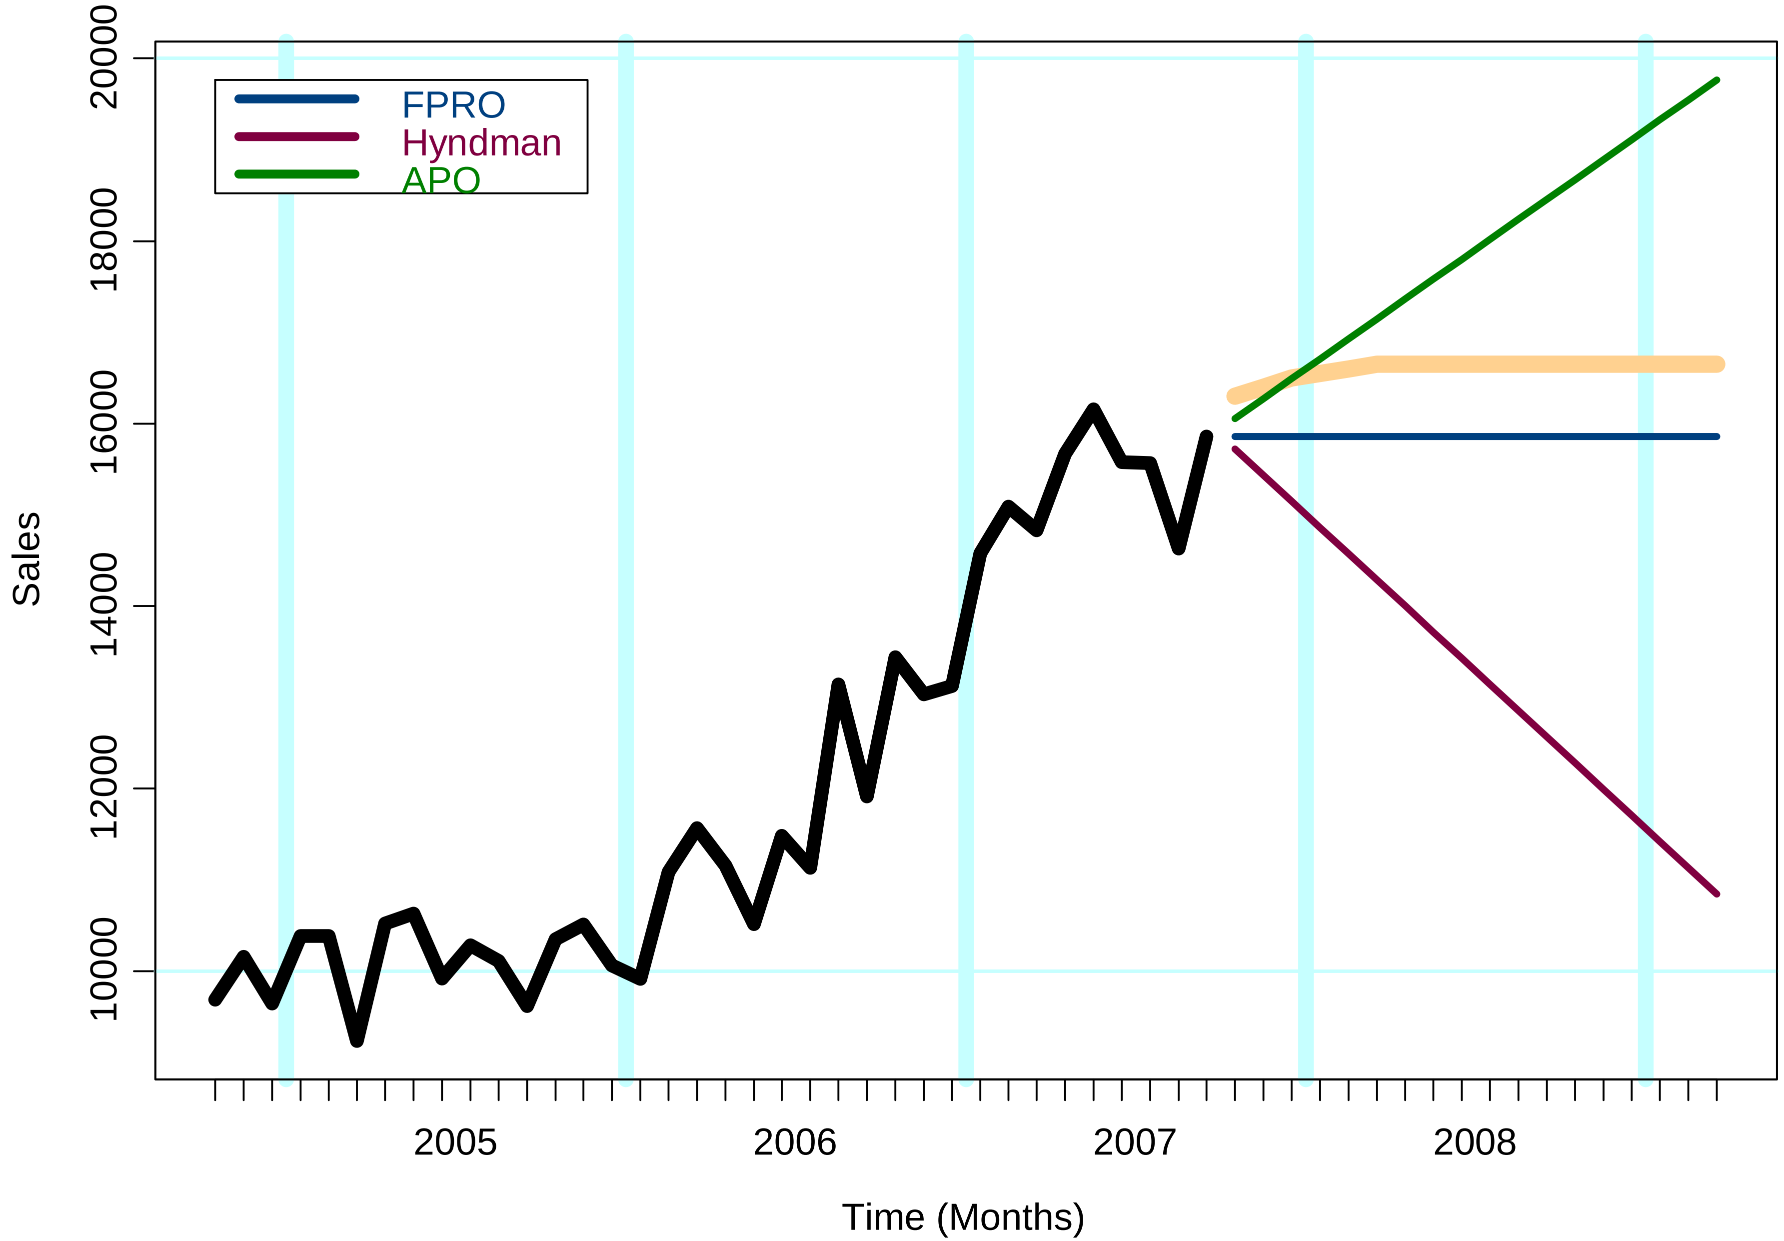
FPRO
Hyndman
APO
Sales
20000
18000
16000
14000
12000
10000
2005
2006
2007
2008
Time (Months)

Real Data

We now show the results of the same forecasting methods on the set of real time series presented earlier.

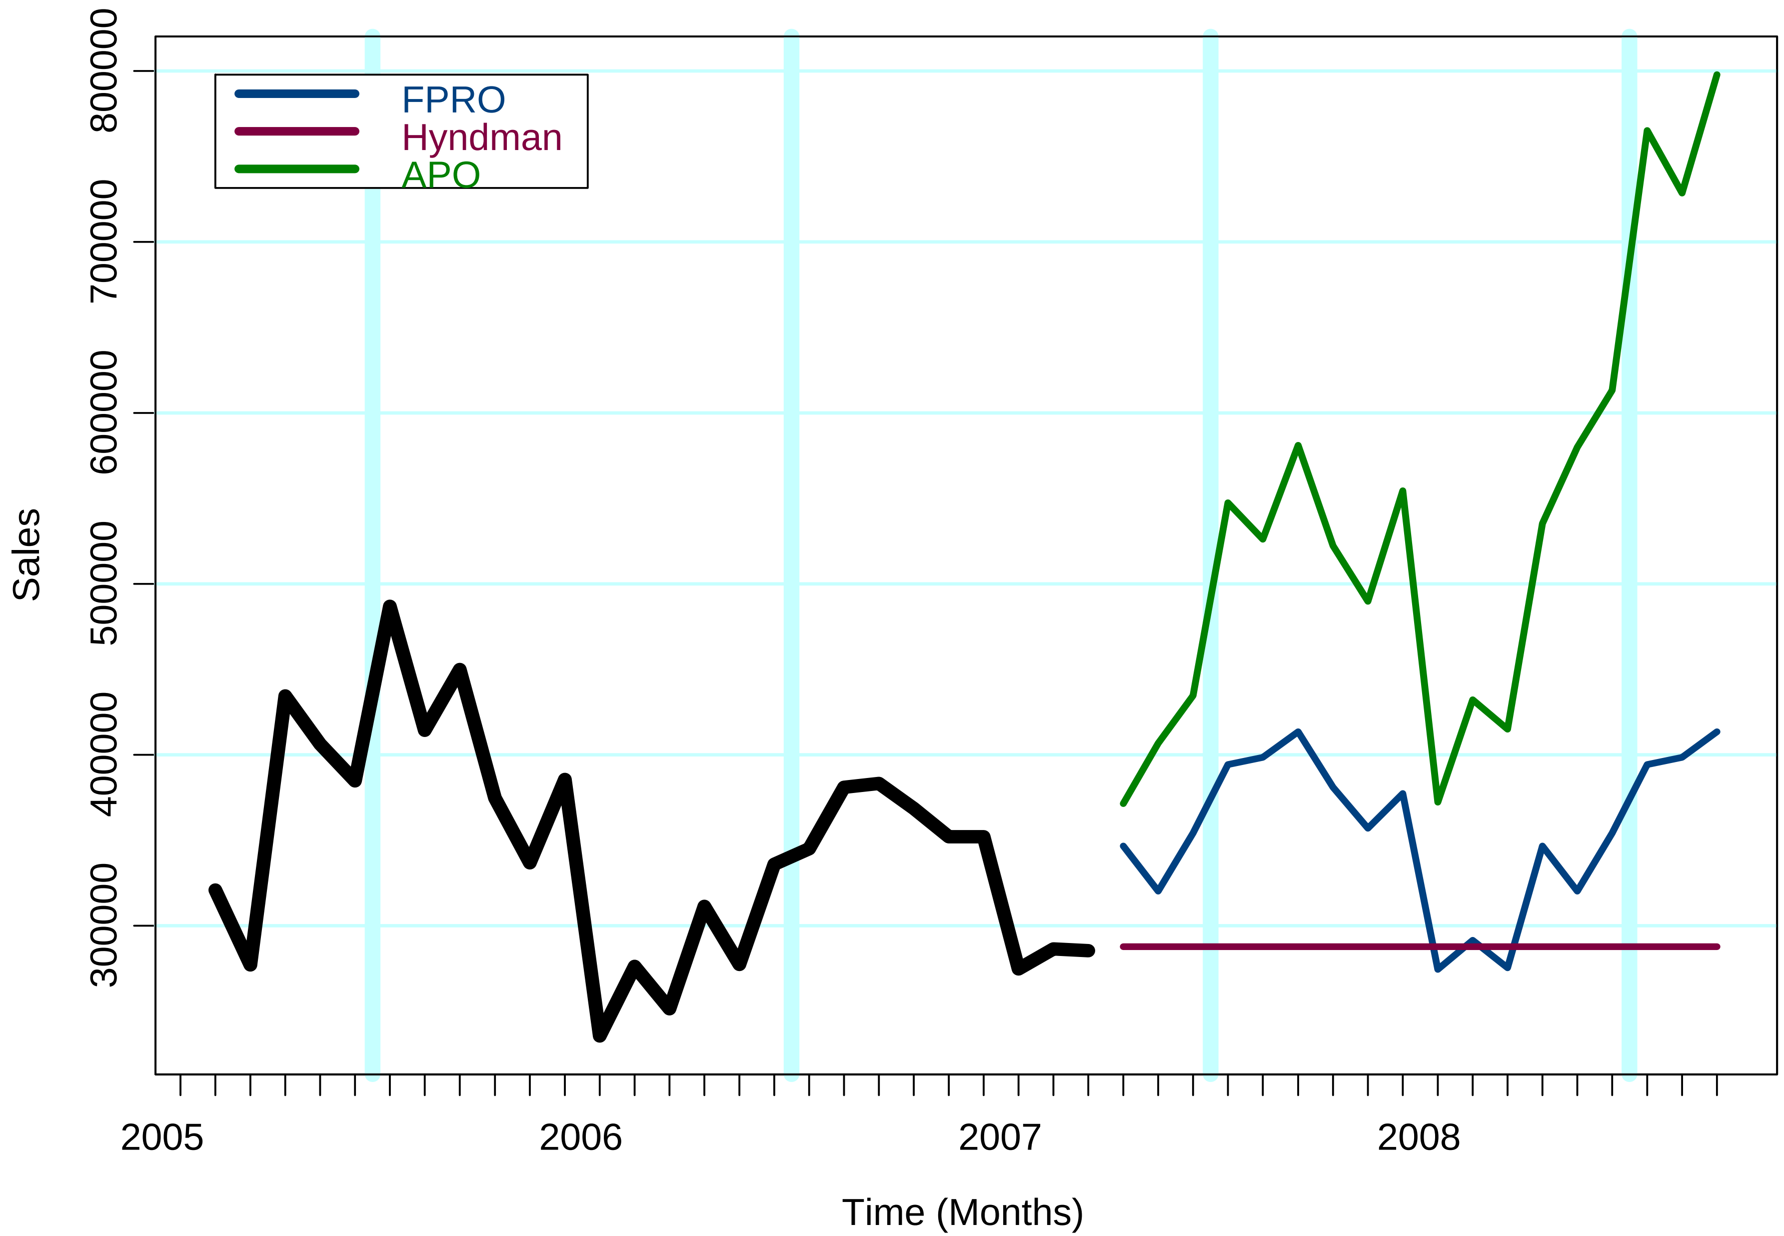

FPRO
Hyndman
APO
Sales
800000
700000
600000
500000
400000
300000
2005
2006
2007
2008
Time (Months)

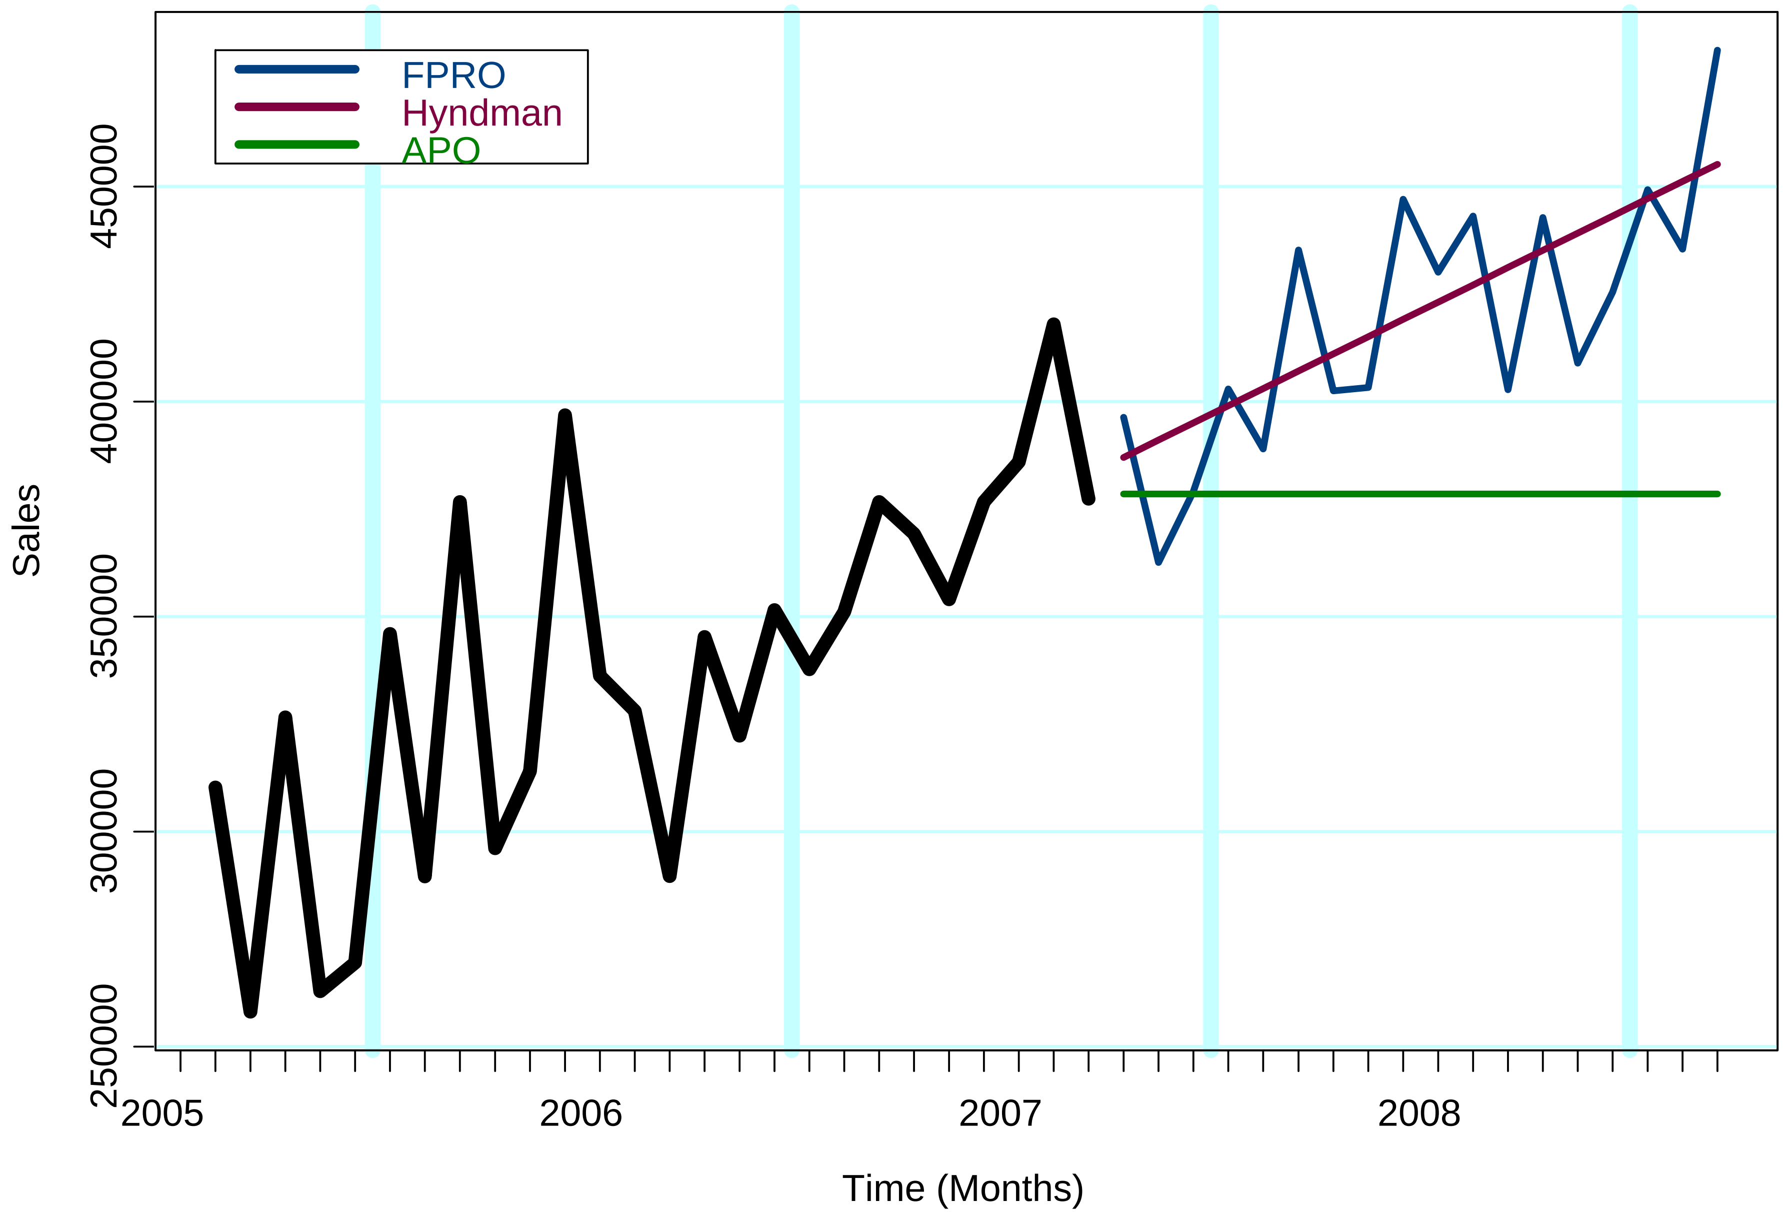
FPRO
Hyndman
APO
Sales
450000
400000
350000
300000
250000
2005
2006
2007
2008
Time (Months)

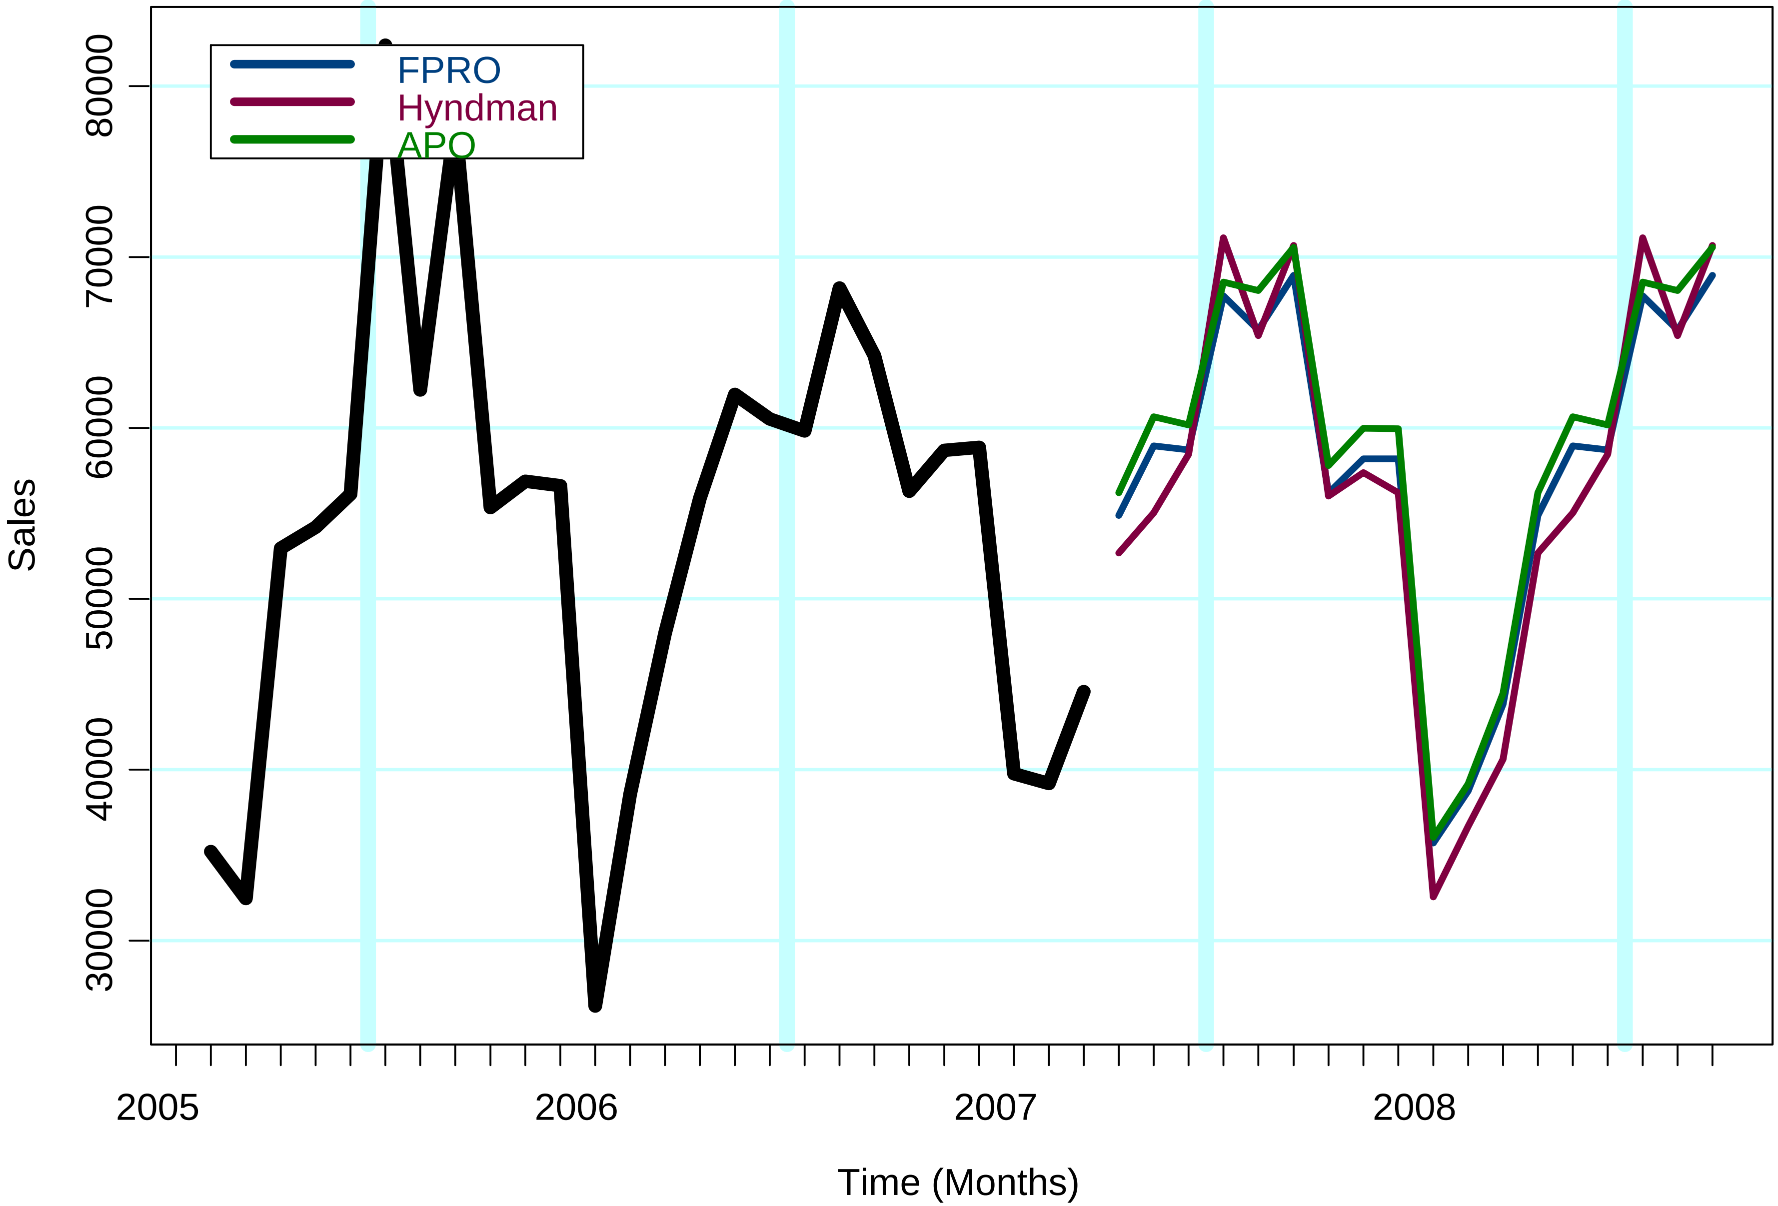

FPRO
Hyndman
APO
Sales
80000
70000
60000
50000
40000
30000
2005
2006
2007
2008
Time (Months)

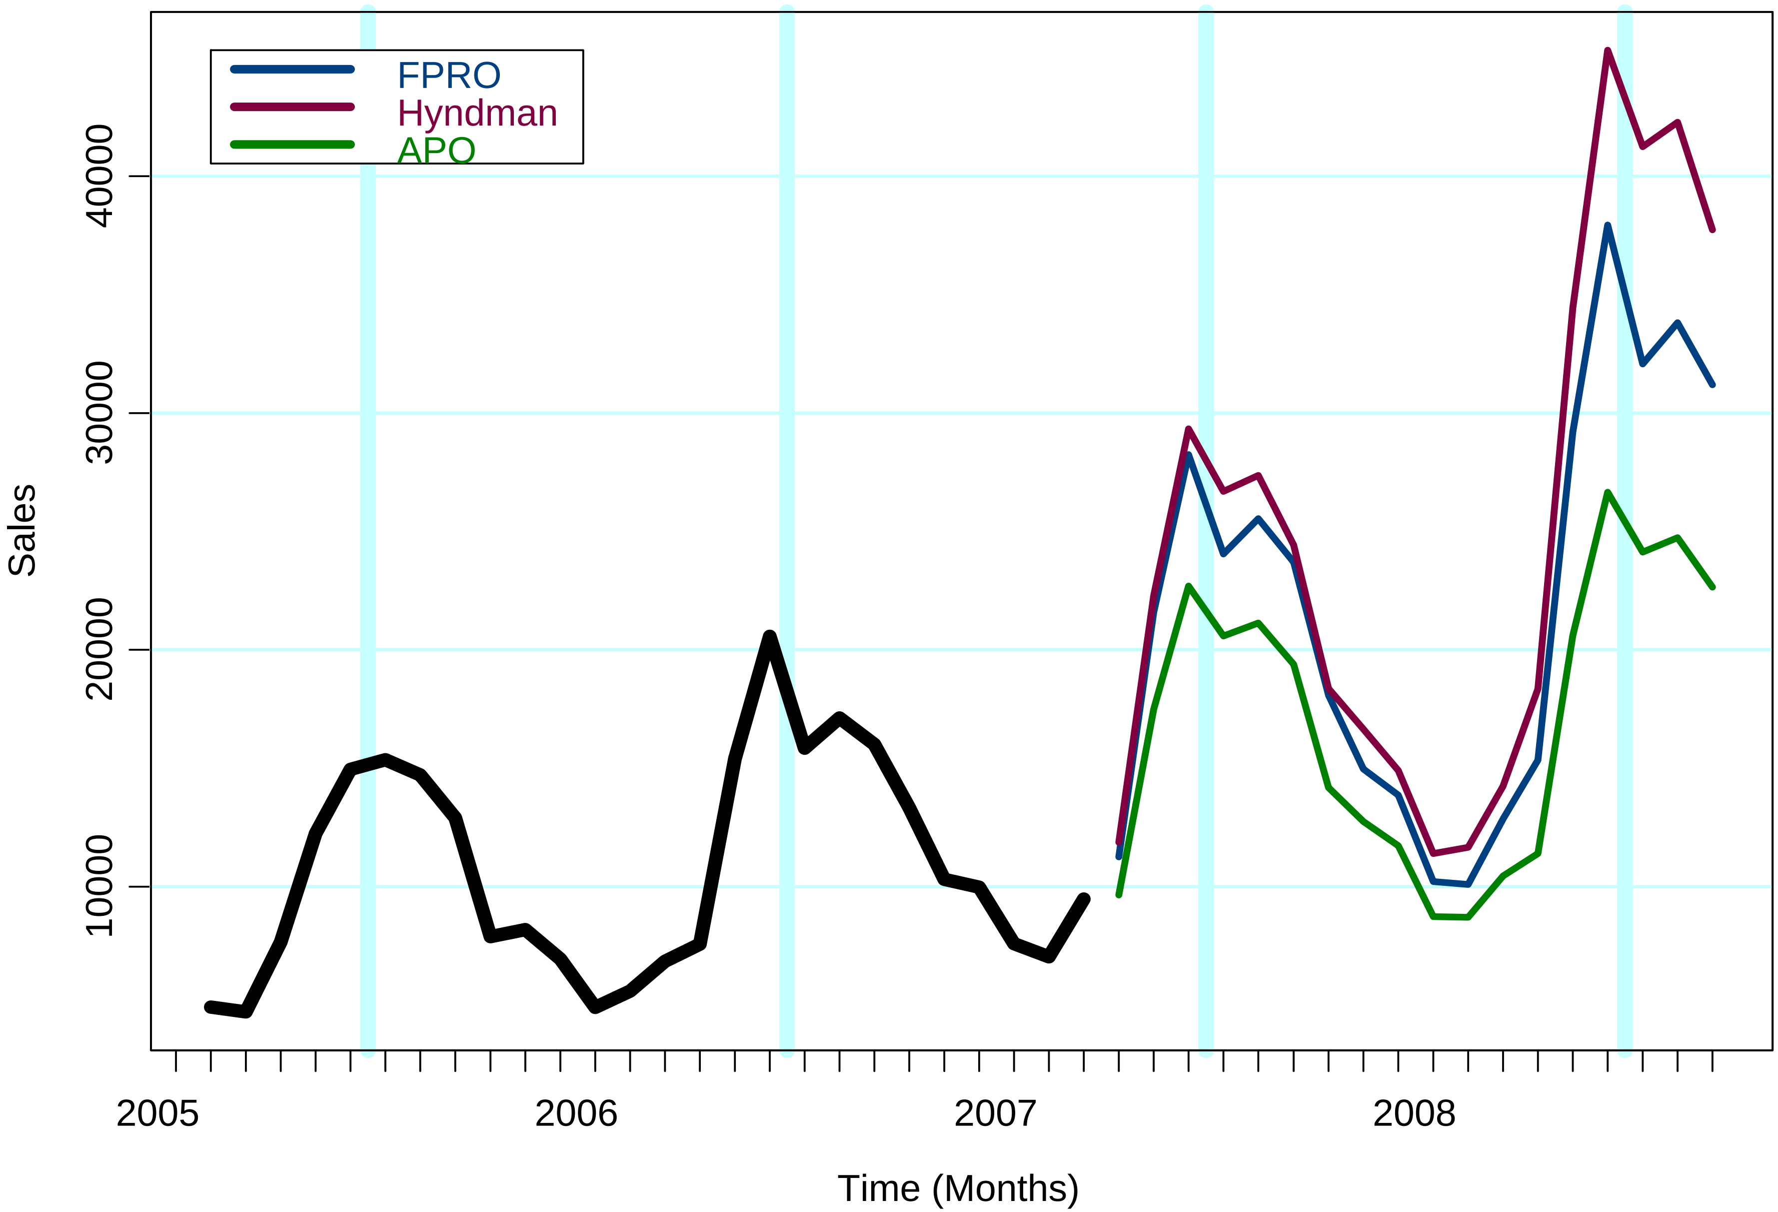
FPRO
Hyndman
APO
40000
30000
20000
10000
Sales
2005
2006
2007
2008
Time (Months)

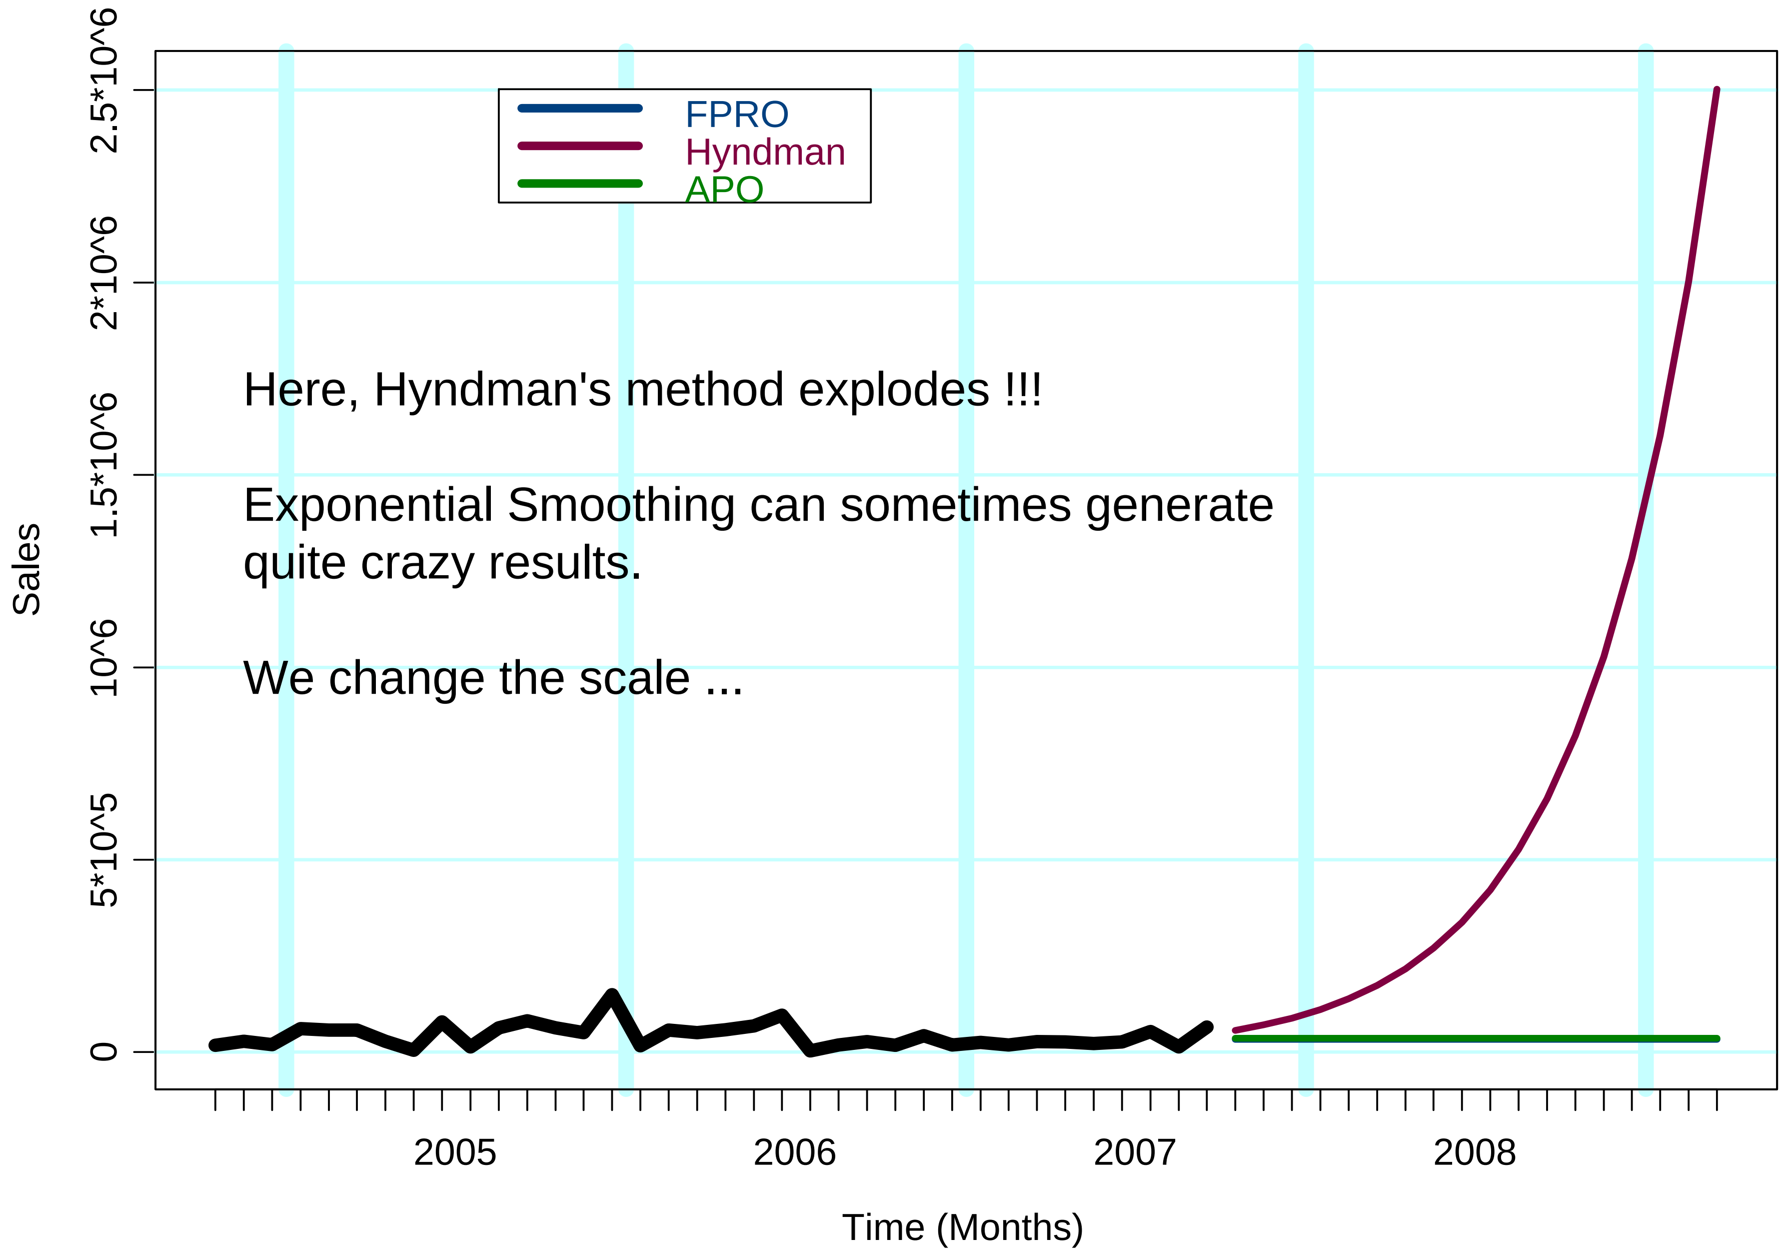

FPRO
Hyndman
APO
Here, Hyndman's method explodes !!!
Exponential Smoothing can sometimes generate quite crazy results.
We change the scale ...
Sales
0
5*10^5
10^6
1.5*10^6
2*10^6
2.5*10^6
2005
2006
2007
2008
Time (Months)

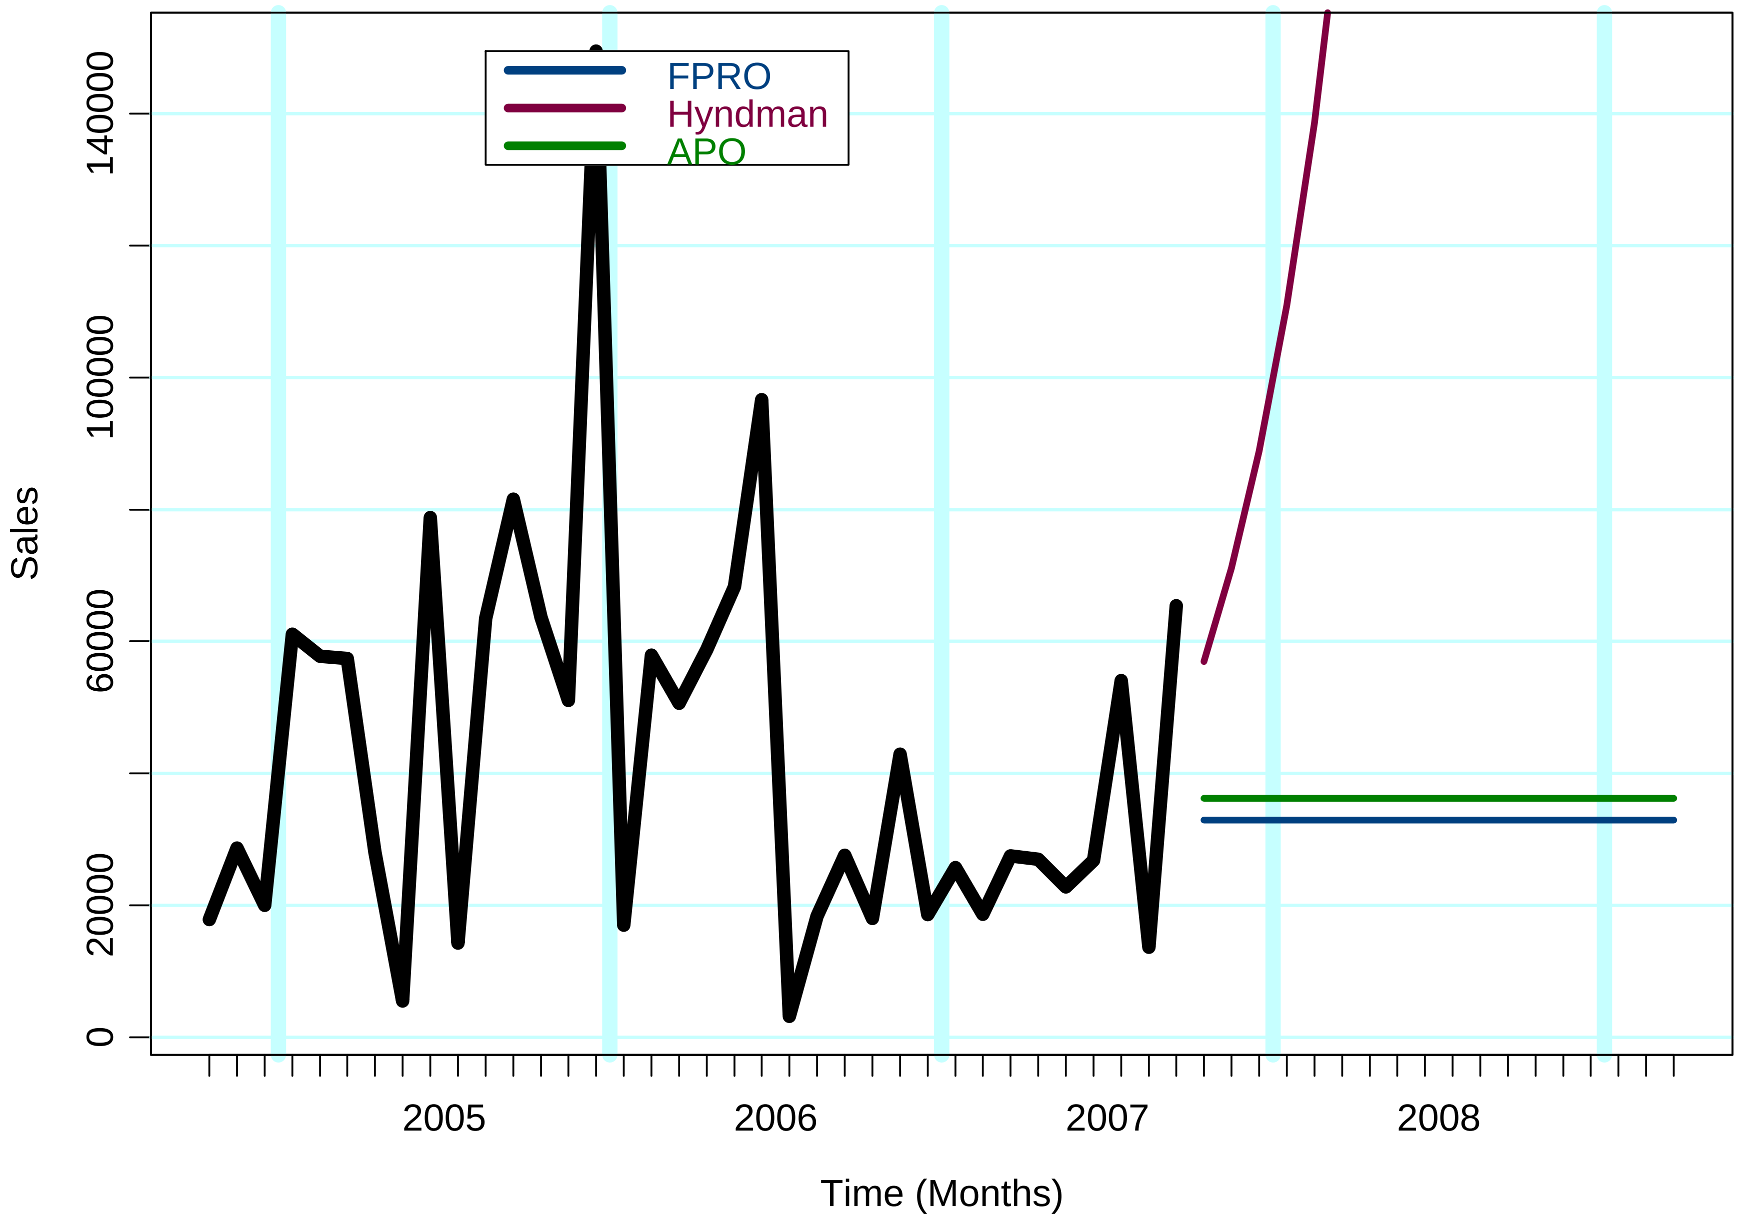

FPRO
Hyndman
APO
Sales
140000
100000
60000
20000
0
2005
2006
2007
2008
Time (Months)

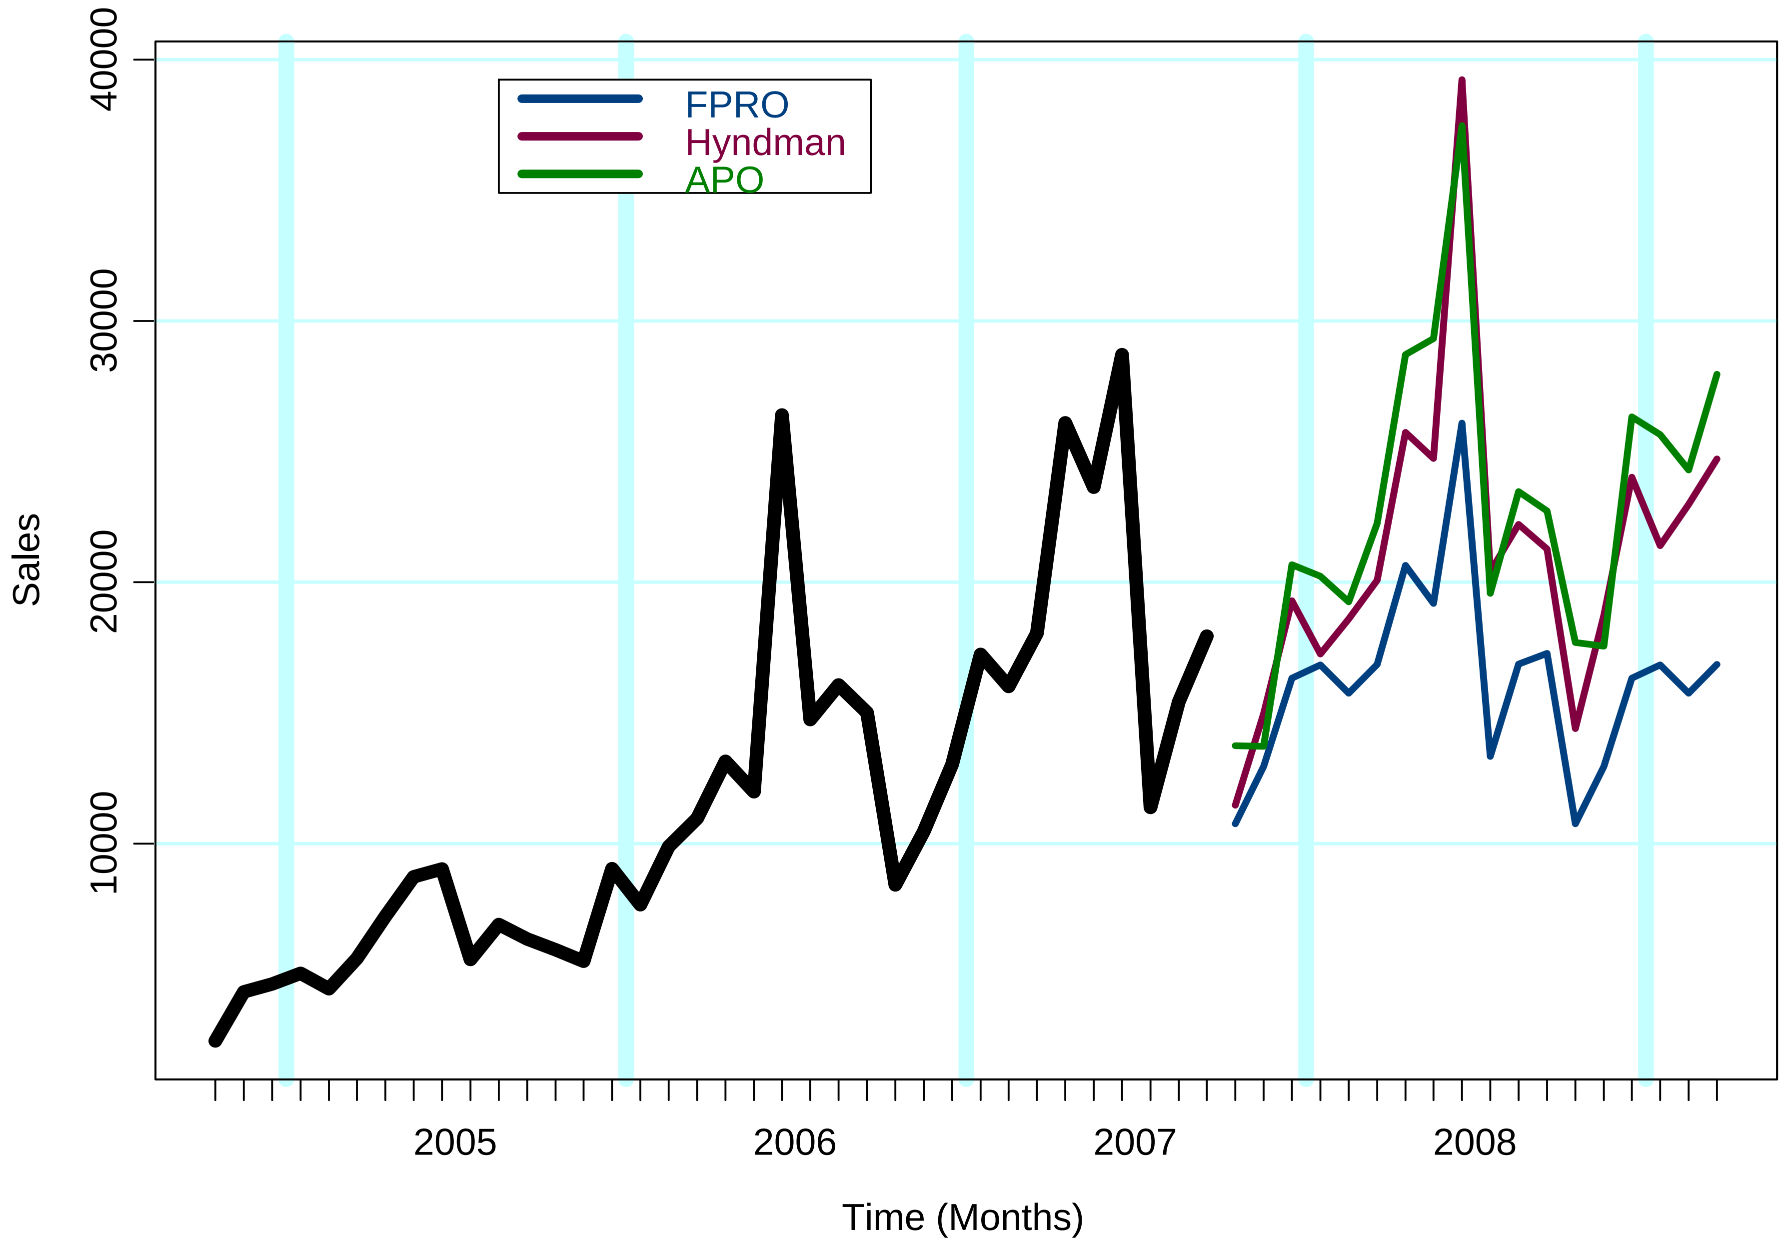
FPRO
Hyndman
APO
40000
30000
20000
10000
Sales
2005
2006
2007
2008
Time (Months)

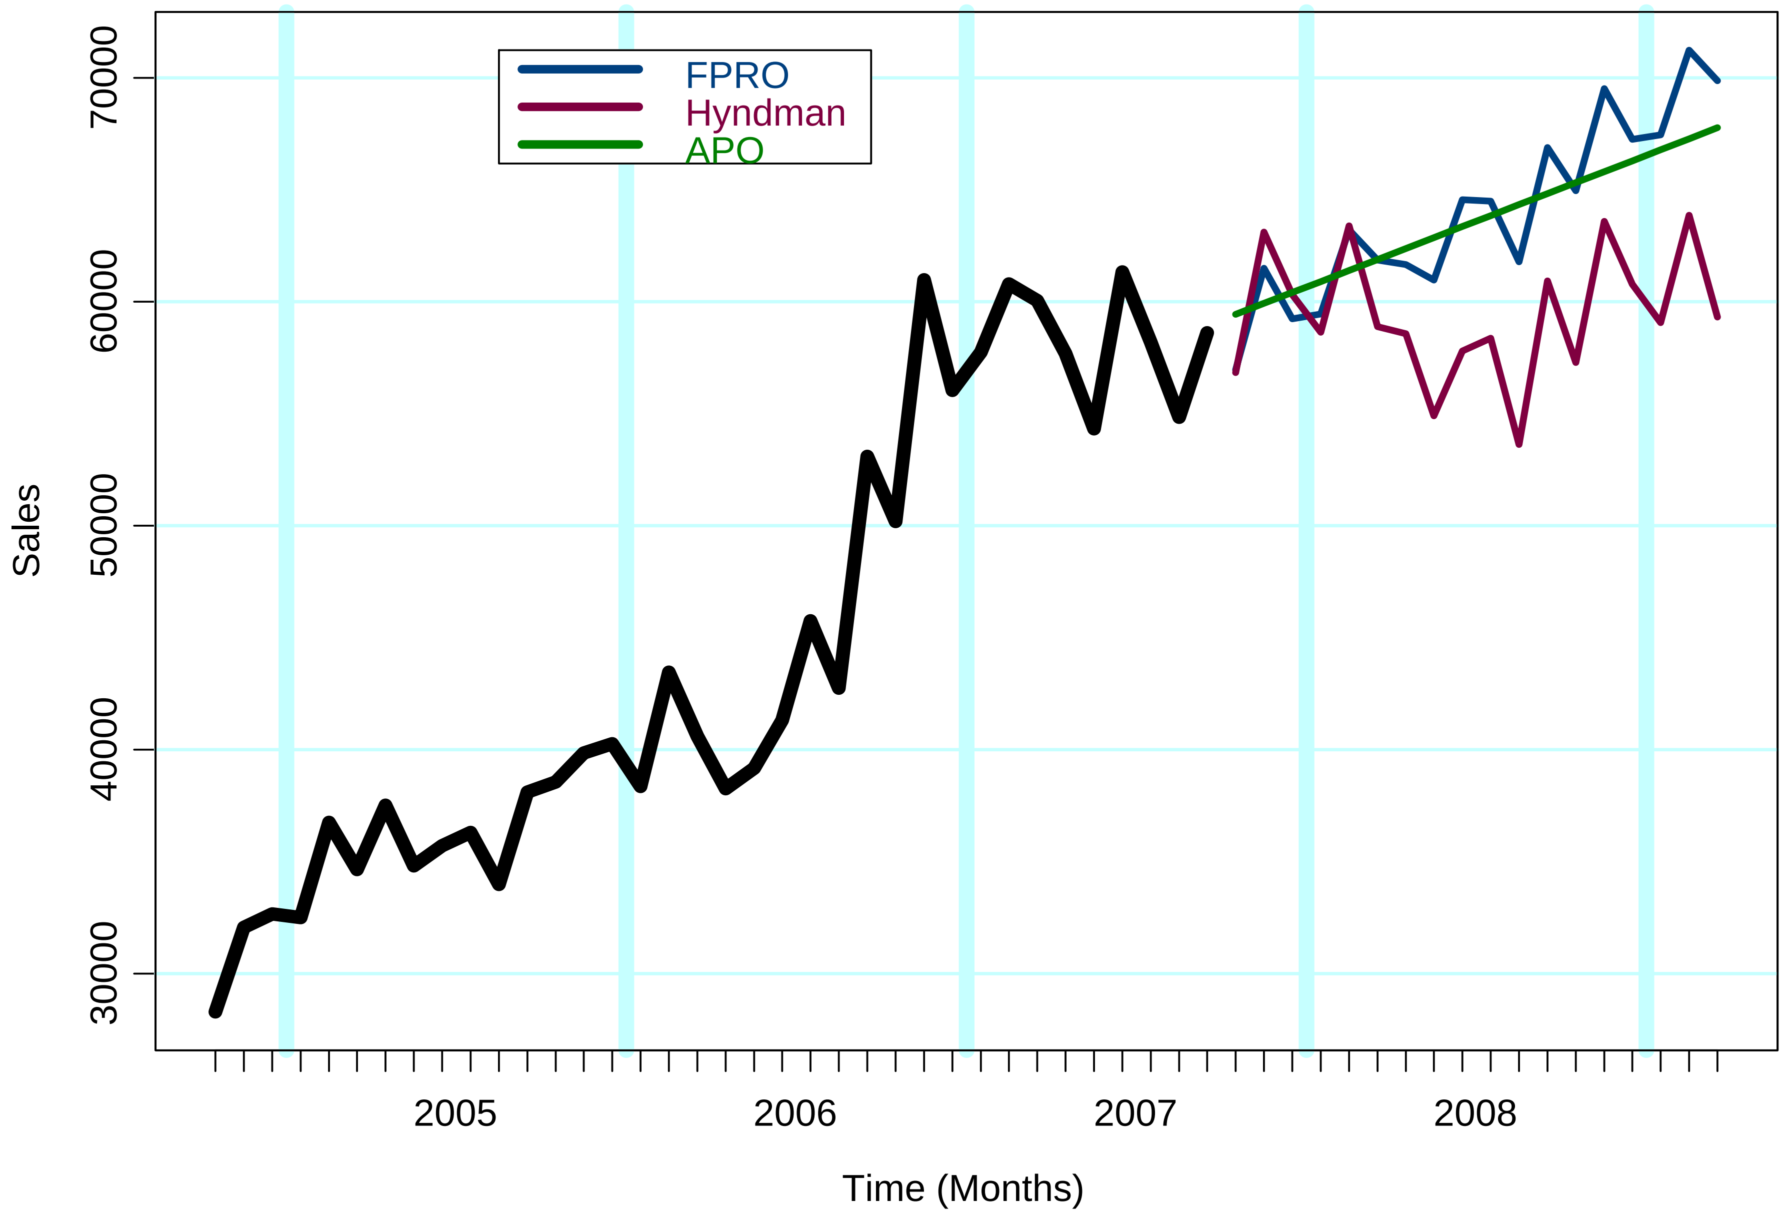
FPRO
Hyndman
APO
Sales
70000
60000
50000
40000
30000
2005
2006
2007
2008
Time (Months)

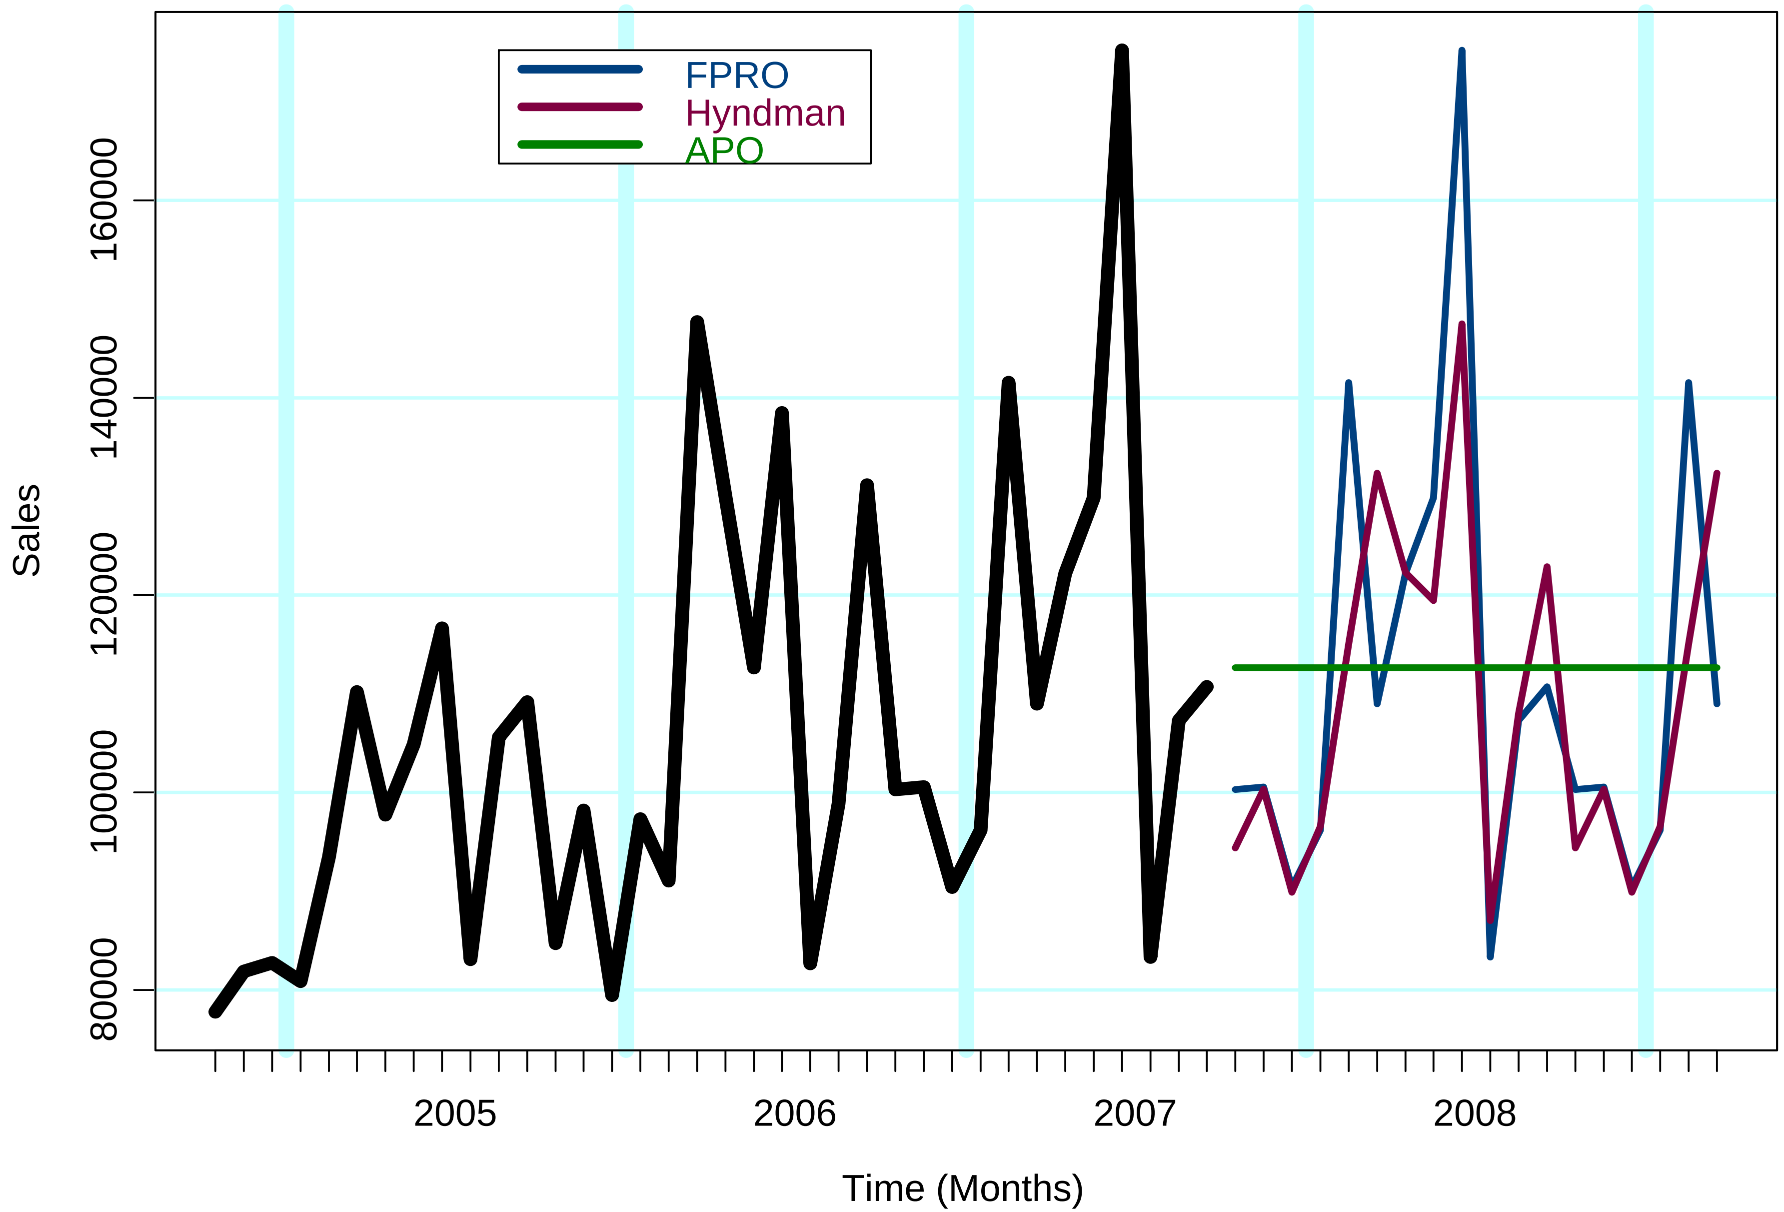
FPRO
Hyndman
APO
80000
100000
120000
140000
160000
Sales
2005
2006
2007
2008
Time (Months)

Conclusions

The different methods behave in a fairly similar way, and with good results for the simulated data.

The more simplistic algorithm of APO also compares favorably with the more advanced approach of ForecastPRO and with Hyndman's method, which has a more solid statistical foundation. For noisy time series, more simple approaches sometimes work even better.

We can thus continue to improve and promote the use of the automatic forecasting method in APO, and follow-up the next developments in the area of state-space models for exponential smoothing.