Forecasting intermittent data with complex patterns

Ivan Svetunkov, John Boylan, Patricia Ramos and Jose Manuel Oliveira

ISF 2018

19th June 2018

Marketing Analytics and Forecasting | Lancaster University Management School

Introduction

In the previous episodes (ISFs)...

(Pictures are from 'Rick and Morty' by Roiland and Harmon, 2013–2017)

Introduction

Svetunkov and Boylan (2017) proposed an intermittent multiplicative state-space model.

We showed that this model underlies Croston (1972) and Teunter et al. (2011) methods.

We extended that model, presenting at ISF2017 the idea of using ETS(M,N,N), ETS(M,M,N) or ETS(M,Md,N) for demand sizes.

Introduction

We showed how to select between different ETS models in this context.

The approach worked well on a WF wholesale data from Johnston et al. (1999).

The conclusion was: you can use one model for both intermittent and non-intermittent data.

Motivation

Now we want more!

Motivation

The reason is this:

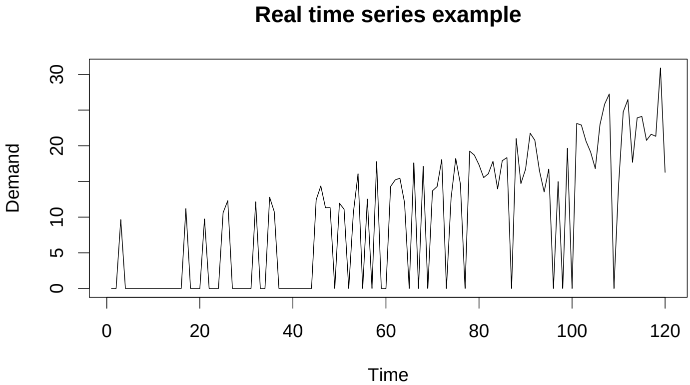

Motivation

How do you deal with this type of data?

There is a trend, but the demand is intermittent.

We can predict the increase in demand sizes with iETS(M,M,N)$_p$...

Motivation

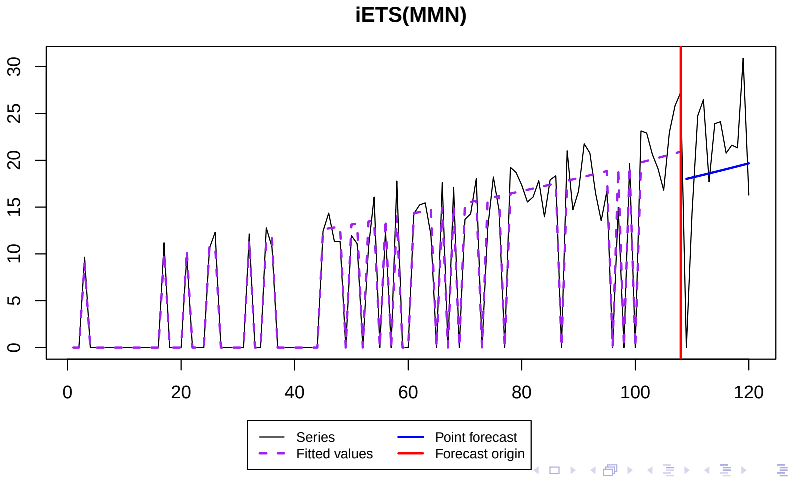
iETS(MMN)
0
5
10
15
20
25
30
0
20
40
60
80
100
120
Series
Fitted values
Point forecast
Forecast origin

Motivation

...Which underforecasts. Because we deal with the following:

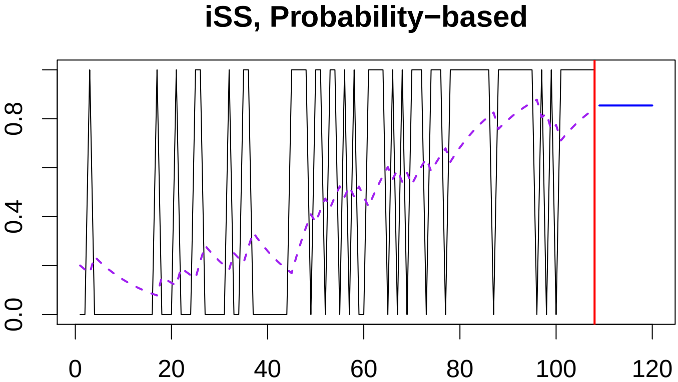

Motivation

We need to capture complex patterns in the occurrence part of demand...

The model

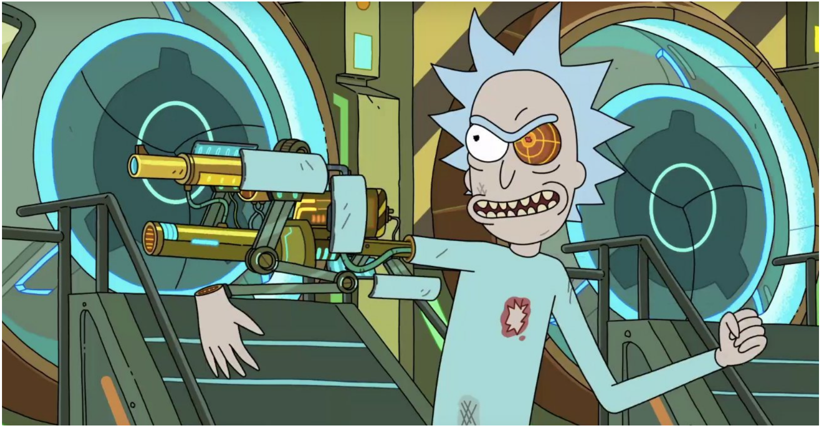

iETS model

$$y_t = o_t z_t, \tag{1}$$

where $o_t \sim \text{Bernoulli}(p_t)$,

z_t is a **statistical model** of our choice

and p_t is another statistical model.

CMAF

Intermittent state-space model

Example. iETS(M,N,N)$_p$ (with probability-based occurrence):

$$
\begin{aligned}
y_t &= o_t z_t \\
z_t &= l_{t-1}(1+\epsilon_t) \\
l_t &= l_{t-1}(1+\alpha\epsilon_t) \\
o_t &\sim \text{Bernoulli}(p_t) \\
p_t &= l_{p,t-1}(1+\epsilon_{p,t}) \\
l_{p,t} &= l_{p,t-1}(1+\alpha_p\epsilon_{p,t})
\end{aligned} \quad (2)
$$

CMAF

Intermittent state-space model

Example. iETS(M,N,N)$_p$ (with probability-based occurrence):

$$\begin{aligned} & y_t = o_t z_t \\ & \left.\begin{aligned} & z_t = l_{t-1}(1+\epsilon_t) \\ & l_t = l_{t-1}(1+\alpha\epsilon_t) \end{aligned}\right\} \text{Demand sizes} \\ & o_t \sim \text{Bernoulli}(p_t) \\ & \left.\begin{aligned} & p_t = l_{p,t-1}(1+\epsilon_{p,t}) \\ & l_{p,t} = l_{p,t-1}(1+\alpha_p\epsilon_{p,t}) \end{aligned}\right\} \text{Demand occurrence} \end{aligned} \tag{3}$$

$1+\epsilon_t \sim \text{logN}(0,\sigma^2)$, which means that z_t is always positive.

$1+\epsilon_{p,t} \sim \text{logN}(0,\sigma_p^2)$.

So far, so good?

CMAF

The problem

But there is a tiny problem...

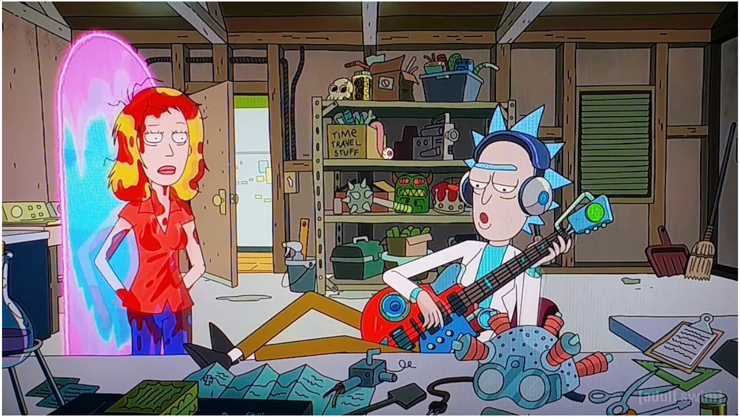

The problem

The problems appear, when we start introducing additional components and variables in the model of p_t:

$$\begin{aligned} p_t &= l_{p,t-1} b_{p,t-1} (1 + \epsilon_{p,t}) \\ l_{p,t} &= l_{p,t-1} b_{p,t-1} (1 + \alpha_p \epsilon_{p,t}), \\ b_{p,t} &= b_{p,t-1} (1 + \beta_p \epsilon_{p,t}) \end{aligned} \quad (4)$$

p_t should be in [0, 1]

But if trend is positive, p_t might become greater than one.

Cutting off values is inhumane...

Logistic transform

The solution - use a different model for p_t.

If we knew the true p_t, then we could use logit transform:

$$q_t = \log \left(\frac{p_t}{1 - p_t} \right) \tag{5}$$

q_t is defined on $(-\infty, \infty)$.

We can use any model for q_t.

CMAF

Logistic transform

For example we can use ETS(A,A,N):

$$\begin{aligned} q_t &= l_{q,t-1} + b_{q,t-1} + \epsilon_{q,t} \\ l_{q,t} &= l_{q,t-1} + b_{q,t-1} + \alpha_q \epsilon_{q,t}, \\ b_{q,t} &= b_{q,t-1} + \beta_q \epsilon_{q,t} \end{aligned} \tag{6}$$

where $\epsilon_{q,t} \sim \mathcal{N}(0, \sigma_q^2)$.

We can extend this model with exogenous variables or seasonal components.

This would mean that in some cases the probability of occurrence increases / decreases.

Logistic transform

In fact, we don't need to know either p_t or q_t, we only need to know $\epsilon_{q,t}$, initial values of $l_{q,0}$ and $b_{q,0}$, and smoothing parameters values.

The latter four can be estimated... IF we have $\epsilon_{q,t}$

But it is unobservable, so...

Logistic model

One can only dream... right?

Logistic transform. The rise of errors

There is a solution...

Use the inverse transform if the value $\hat{q}_t$ is known:

$$\hat{p}_t = \frac{\exp(\hat{q}_t)}{1 + \exp(\hat{q}_t)} \quad (7)$$

Compare the predicted probability $\hat{p}_t$ with the outcome o_t:

$$u_t = o_t - \hat{p}_t \quad (8)$$

The problem now is to translate this error into $\epsilon_{q,t}$.

CMAF

Logistic transform. The rise of errors

u_t lies in $(-1, 1)$.

We transform u_t, so that it lies in $(0, 1)$:

$u'_t = \frac{1+u_t}{2}$.

and then use logit transform to obtain an estimate of error:

$$e_{q,t} = \log \left(\frac{1 + o_t - \hat{p}_t}{1 - o_t + \hat{p}_t} \right) . \tag{9}$$

If $o_t = \hat{p}_t$, then $e_{q,t} = 0$ (because o_t is binary).

CMAF

iETS$_l$ model

So the final model is:

$$\begin{aligned} y_t &= o_t z_t \\ z_t &\sim \text{ETS(Y,Y,Y)} \\ o_t &\sim \text{Bernoulli}(p_t) \\ p_t &= \frac{\exp(q_t)}{1+\exp(q_t)} \\ q_t &\sim \text{ETS(X,X,X)} \\ e_{q,t} &= \log\left(\frac{1+o_t-\hat{p}_t}{1-o_t+\hat{p}_t}\right) \end{aligned}, \quad (10)$$

where ETS(Y,Y,Y) is a multiplicative ETS, and ETS(X,X,X) is an additive one.

CMAF

iETS$_l$ model

How does iETS$_l$ work?

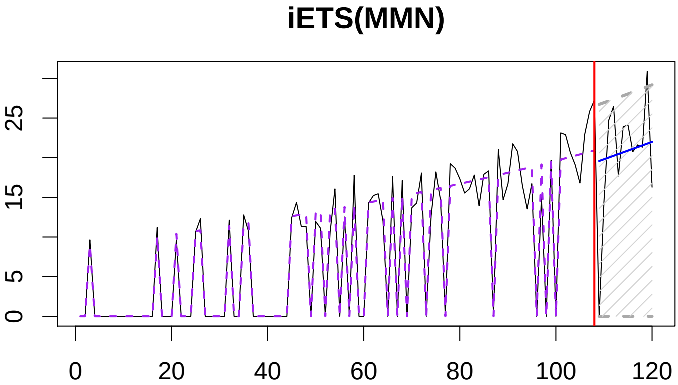

iETS$_l$ model

The occurrence part of iETS$_l$

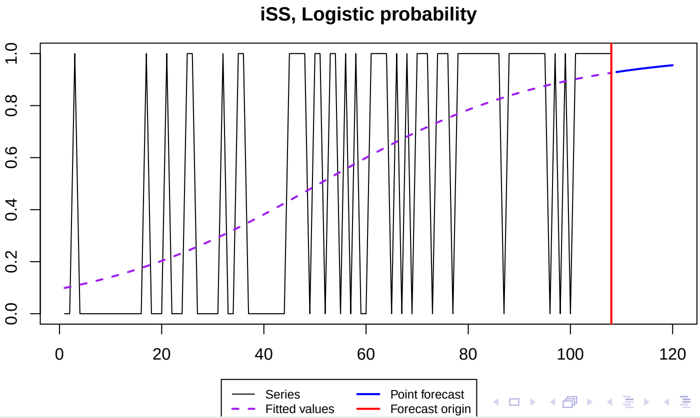

iETS$_l$ model

So now we have:

- an extendable model...
- ...that captures complex patterns for demand sizes...
- ...AND demand occurrence part.

CMAF

Experiment

Yes! Now it is time for experiments...

Experiment

Intermittent data of a Portuguese retailer.

5275 SKUs with at least 4 non-zero demands each.

Weekly data, 173 observations.

$h = \{1, 2, 3, 4\}$.

Rolling with 52 origins.

sMSE, sME, sAPIS (Petropoulos and Kourentzes, 2015).

CMAF

Experiment

Croston, TSB and iMAPA from `tsintermittent` package

es function from `smooth` package with:

- ETS(A,N,N),
- iETS(M,Y,N)$_f$ - fixed probability, selection of trend,
- iETS(M,Y,N)$_i$ - Croston style,
- iETS(M,Y,N)$_p$ - TSB style,
- iETS(M,N,N)$_l$(A,N,N) - logistic with ETS(A,N,N) for occurrence,
- iETS(M,Y,N)$_l$(A,N,N) - similar + selection of trend,
- iETS(M,Y,N)$_l$(A,X,N) - similar + selection for occurrence.

Results

Model	sME	sMSE	sAPIS
iETS(MYN)$_l$(AXN)	0.460	**44.865**	**1.447**
iETS(MNN)$_l$(ANN)	0.428	54.706	1.520
iETS(MYN)$_l$(ANN)	0.471	54.744	1.512
iETS(MNN)$_p$	0.596	55.391	1.493
ETS(AAN)	**0.176**	55.970	1.627
iETS(MNN)$_i$	0.565	56.073	1.514
iMAPA	0.986	58.693	1.644
iETS(MNN)$_f$	1.328	61.397	1.667
TSB	1.092	61.669	1.684
Croston	1.102	61.937	1.692

CMAF

Results

Based on the iETS$_l$:

	Demand occurrence		
Demand sizes	No trend	Trend	Overall
No trend	31.9%	35.7%	67.6%
Trend	19.0%	13.4%	32.4%
Overall	50.9%	49.1%	100%

CMAF

Conclusions

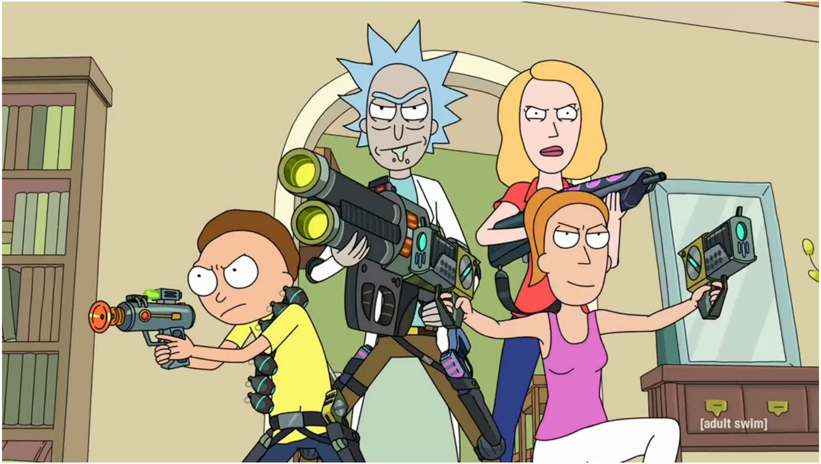

Conclusions

- We now have a more general modelling framework;
- We have a new type of model for intermittent data;
- We can capture complex patterns in intermittent data;
- The approach seems to work well in practice.

What's next?

- More thorough analysis of results driven by the data;
- Go multivariate – vector intermittent models:
  1. Extend the iETS$_l$ to iVES$_l$;
  2. Group time series based on the characteristics;
  3. Capture seasonality across similar time series;
  4. Introduce exogenous variables;
  5. Forecast groups of time series.
- In the end we should have a universal time series forecasting approach...

CMAF

And that's how it's done!

Thank you for your attention!

Ivan Svetunkov

i.svetunkov@lancaster.ac.uk

https://forecasting.svetunkov.ru

twitter: @iSvetunkov

Marketing Analytics and Forecasting

Lancaster University Management School

Croston, J. D., sep 1972. Forecasting and Stock Control for Intermittent Demands. Operational Research Quarterly (1970-1977) 23 (3), 289.
URL
http://www.jstor.org/stable/3007885?origin=crossref

Johnston, F. R., Boylan, J. E., Meadows, M., Shale, E., dec 1999. Some Properties of a Simple Moving Average when Applied to Forecasting a Time Series. The Journal of the Operational Research Society 50 (12), 1267–1271.
URL http://wrap.warwick.ac.uk/13847/http://www.jstor.org/stable/3010636?origin=crossref

Petropoulos, F., Kourentzes, N., jun 2015. Forecast combinations for intermittent demand. Journal of the Operational Research Society 66 (6), 914–924.
URL http://dx.doi.org/10.1057/jors.2014.62{%}7B{%}25{%}7D5Cn10.1057/jors.2014.62http:

CMAF

//link.springer.com/10.1057/jors.2014.62

Roiland, J., Harmon, D., 2013–2017. Rick and morty. Adult Swim.

Svetunkov, I., Boylan, J. E., 2017. Multiplicative State-Space Models for Intermittent Time Series.

Teunter, R. H., Syntetos, A. A., Babai, M. Z., nov 2011. Intermittent demand: Linking forecasting to inventory obsolescence. European Journal of Operational Research 214 (3), 606–615.
URL http://linkinghub.elsevier.com/retrieve/pii/S0377221711004437