

# Bangor Business School

## Working Paper



PRIFYSGOL  
**BANGOR**  
UNIVERSITY

**BBSWP/15/03**

**Forecasting, Foresight and Strategic Planning for Black Swans<sup>1</sup>**

**By**

**K. Nikolopoulos<sup>a,2</sup>, & F. Petropoulos<sup>b</sup>**

<sup>a</sup> forLAB, Bangor Business School, Prifysgol Bangor, Bangor, LL57 2DG, Wales, U.K.

<sup>b</sup> Cardiff University Business School, Prifysgol Cardiff, Cardiff, CF10 3EU, Wales, U.K.

**March, 2015**

**Bangor Business School  
Bangor University  
Hen Goleg  
College Road  
Bangor  
Gwynedd LL57 2DG  
United Kingdom  
Tel: +44 (0) 1248 382277  
E-mail: [business@bangor.ac.uk](mailto:business@bangor.ac.uk)**

---

<sup>1</sup> Earlier versions of this work has been presented at: a) Nikolopoulos, K. & Syntetos, A.A. (2013), "Forecasting Black, Grey & White Swans...", 2013 DSI Annual Meeting, November 16-19, 2013, Baltimore, Maryland, USA, and b) Nikolopoulos, K., Syntetos, A.A., Batiz-Lazo, B., (2013), "Forecasting Black (& White) Swans...", OR55, 3-5 September 2013. The University of Exeter, Exeter, England, United Kingdom. Keynote Presentation for the 'Forecasting stream'. We thank Professor Syntetos from Cardiff University Business School for his views on earlier versions of this paper.

<sup>2</sup> Corresponding author: [kostas@bangor.ac.uk](mailto:kostas@bangor.ac.uk)

# Forecasting, Foresight and Strategic Planning for Black Swans

## Abstract

In this research essay we propose a methodological innovation through: a) advocating for the broader use of OR forecasting tools and in specific *intermittent demand estimators* for forecasting black and grey swans, as a simpler, faster and quite robust alternative to econometric probabilistic methods like EVT; b) demonstrating the use in such a context of a rather popular forecasting paradigm: the Naive method (*forecasting* short horizons) and the SBA method (*foresight* long horizons) through the ADIDA non-overlapping temporal aggregation method-improving framework, and c) arguing for a new way for *deciding the strategic planning horizon* for phenomena prone to the appearance of black and grey swans.

**Keywords:** Forecasting, Black Swans, Intermittent Demand, Temporal Aggregation, Forecasting Horizon

# Forecasting, Foresight and Strategic Planning for Black Swans

## White, Gray and Black Swans...

In this research essay we follow Taleb's (2007) classification of events into:

- White Swans: the mainstream... *"indeed the normal is often irrelevant..."* and definitely predictable with time series as well as causal methods
- Gray Swans: rare but expected... *"somewhat predictable, particularly to those who are prepared for them and have the tools to understand them"*
- Black Swans: the unexpected... *"fewer remote events, but are more and more extreme in their impact..."*

We particularly focus on the latter two classes and we:

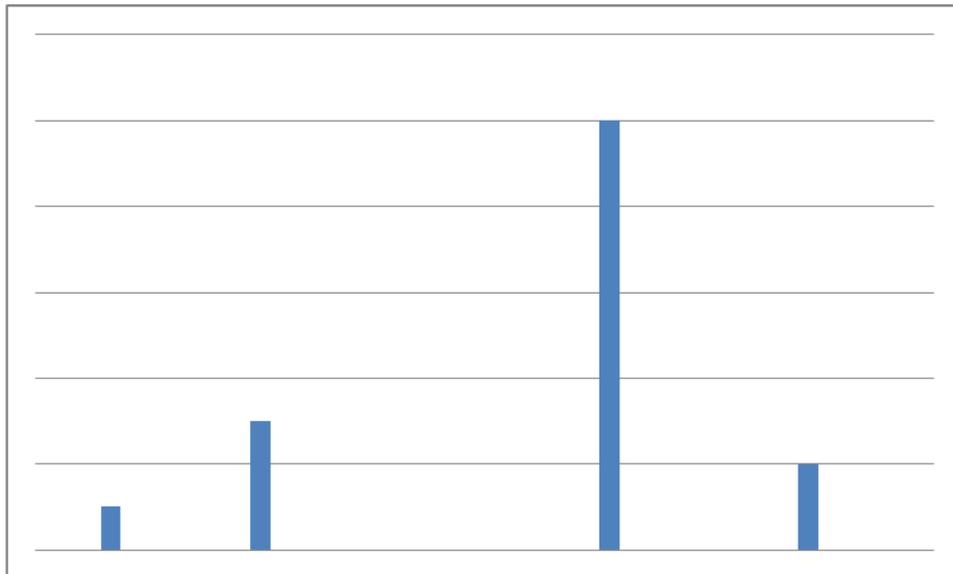
- a) advocate for the use of Operational Research (OR) forecasting tools i.e. intermittent demand estimators for forecasting black and grey swans, as a simpler, faster and quite robust alternative to econometric probabilistic methods like Extreme Value Theory (EVT) (Leadbetter, 1991),
- b) demonstrate the use in such a context of a rather popular forecasting paradigm: the Naive method (for short horizons or *forecasting* tasks) and the SBA method (Syntetos and Boylan, 2001, for long horizons or *foresight* tasks) through the ADIDA non-overlapping temporal aggregation method-improving framework (Nikolopoulos et al., 2011), and
- c) propose a new way for deciding the strategic planning horizon for phenomena prone to the appearance of black and grey swans.

The remain of the paper is structured according to the aforementioned objectives.

## Forecasting Black and Gray Swans with intermittent demand estimators

Intermittent demand patterns are evident very often in real life and have been studied for thousands of years: 2500 years ago in Ancient Greece Archimedes described the intermittent nature of earthquake occurrences, and we are pretty sure he was not the first one to use such modeling in the past. During the last fifteen years there has been a significant increase in the interest in intermittent demand and ways to classify, model and forecast in this context (Syntetos et al., 2005); to attest to that, the seminal work from Croston (1972) has received in this recent period more than 100 citations!

We argue that this interest could be further augmented, via a time series decomposition lens in a similar fashion like Leadbetter (1991) isolating 'Peaks over Threshold' data points; this would lead to a new intermittent demand series with identical characteristics with the series that Croston aspired to mathematically model and forecast back in 1972. In essence one could use intermittent demand modeling and forecasting techniques so as to tackle much more difficult and important phenomena and problems coming from finance, business, economics and even social sciences.



**Figure 1.** Intermittent Demand

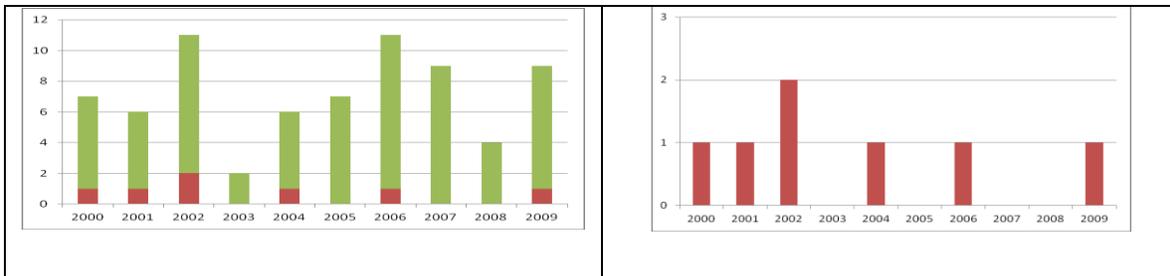
Intermittent Demand by and large it looks like figure 1, that is periods with demand interchanging with periods of no demand at all, and on top of this even the demand volume (when realized) comes with significant variation. It is obvious that the forecasting problem under consideration, is at least twice as hard as a traditional time series forecasting problem, as there are two things to forecast:

- When the next demand period is going to be realized?
- Whenever demand is realised, what will be the volume of this demand?

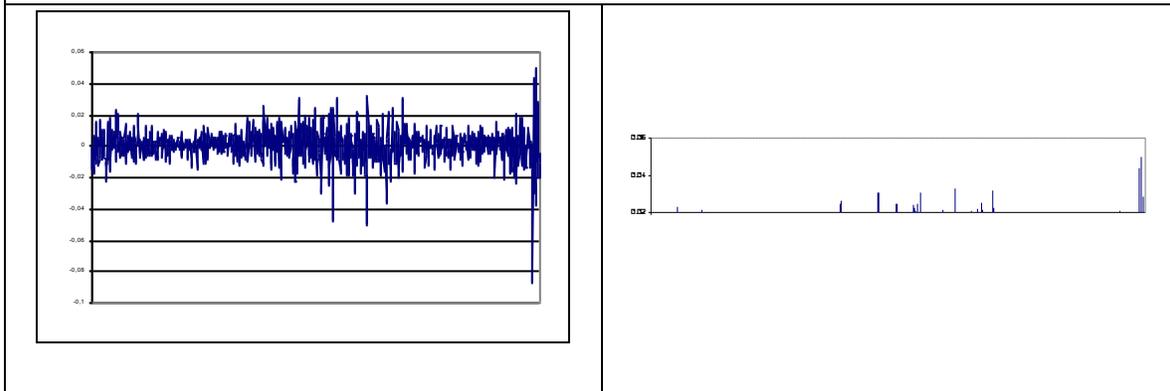
### ***The Decomposition lens***

In this section we argue through a series of illustrative examples, how one could decompose very 'important' series into two subseries: one containing the baseline (*white swans*) and one containing the extremes (*gray* and *black swans*). For the former traditional forecasting approaches could be used. For the latter, intermittent demand approaches may be the way to go. And bare in mind, all the risk, all the uncertainty lay by these very extremes! So it is very important to be able to effectively model and forecast them. In figure 2 we provide the motivation for such treatment of time series via historic data for earthquakes, stock markets and epidemics.

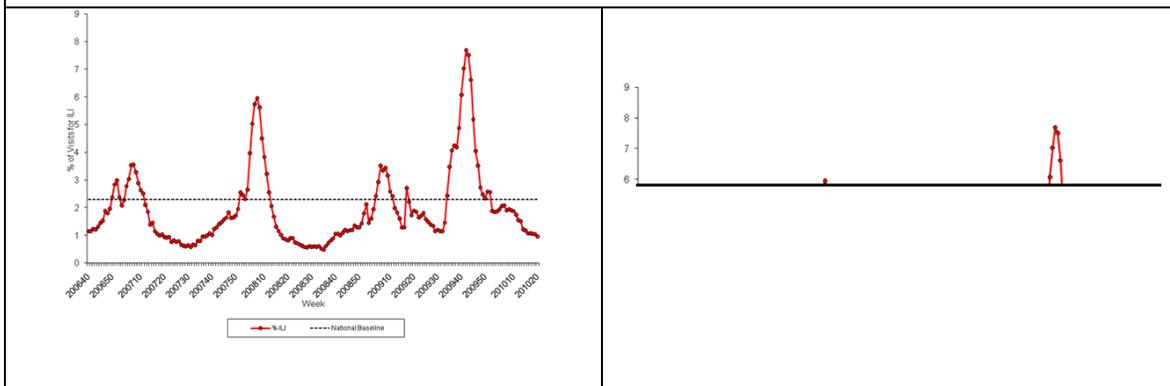
It is evident that all three graphs on the right column of figure 2 have similar characteristics with what is traditional perceived and classified as intermittent demand (figure 1) and thus intermitted demand forecasting approaches could be perfectly employable: from the original inception of Croston (1972) up to more advanced and recent developments (Willemain et al., 2004, Nikolopoulos et al., 2011; Petropoulos and Kourentzes, 2014). For an extensive review for methods specifically designed for and empirically proved to perform well on intermitted data the reader may follow either Petropoulos et al. (2014) for insights on simulated data or Syntetos et al. (2015) for empirical evaluation on real data.



**2(a). Earthquakes** : the graph on the right is the lower 'dark-red' part of the graph on the left i.e. the earthquakes with magnitude more than 7 points in the Richter scale in large densely-populated areas in USA).



**2(b). Stock Markets** : S&P 500, LogDiffs of Weekly Closing Values. [1/1/1990-1/2/2009] The graph on the right is the upper part of the graph on the left – anything over the 0.02 points limit line.



**2(c). Epidemics - Influenza** : the graph on the right is the upper part of the graph on the left – anything over the seasonal influenza level of 6 points limit line (source: <http://www.cdc.gov/h1n1flu/updates/us/>). The variable measured is the percentage of Visits for Influenza-like Illness as reported from the United States Centers for Disease Control and Prevention.

**Figure 2.** Examples of an Intermittent Demand series as Points-Over-Threshold part of a

regular time series

It is even more evident how important it is to be able to forecast in such contexts: In figure 2(a) it is those big earthquakes that are potentially responsible for major catastrophes including tsunamis and nuclear disasters like the recent Fukushima incident reminds us. In figure 2(b) it is those periods in the Stock markets that are extremely volatile and thus arguably considered as periods of high uncertainty. While in 2(c) it is those periods where normal periods of virus activity converts suddenly into epidemics and bring entire national health systems into their knees...

All these phenomena, all these special events as illustrated in the right column of figure 2 - and those are just a few examples - can arguably been considered as *black swans*; or just *gray swans* if we assume that black ones are even more rare and of more extreme impact!

### ***What can really be forecasted?***

Given that none of the Intermittent Demand forecasting approaches will give you the exact timing of the forthcoming extreme event, but rather it will provide a *cumulative* forecast over a long period of time, the question still is what can you do with this forecast?

If this *cumulative* or aggregate or 'demand per period estimate' will help you in taking an informed decision, then yes there is merit in applying a decomposition approach and using intermittent demand forecasting techniques for the extrapolation of the new time series: constructed from the extremes of a regular time series.

For each of the aforementioned exemplar application areas, let's think if there is merit in and what are the implications in such contexts there from achieving a more accurate *cumulative* forecast:

- Healthcare (Epidemics): clear implications as you could proactively source resources, get temporal staff in time etc.
- Stock Markets (Volatility/Uncertainty): such a forecast would be an estimate of the aggregate risk over a certain and typically long period of time and thus necessary hedging actions could be taken in advance.

- Earthquakes (Catastrophes): although even if you know that a major earthquake is going to hit a region it is hard to decide to evacuate cities, still you can influence and legislate the way buildings are built and increase the awareness and readiness of the public as well as make sure there is capital on hold so as to cope with the aftermath of the disaster

***If not Croston et al... then what?***

What is the alternative...? Or more accurately what is the current doctrine when forecasting in such a context: advanced probabilistic models. These methods typically require lot's of data and typically try reconstruct the distributions of the underlying phenomena -these come with a merit but definitely with lot's of constraints as well: big *datasets* needed, increased *computational* time, high mathematical *complexity* (of methods), and not clear to practitioners how these methods do actually work and thus less *acceptance*. For the sake of completeness we do make a short presentation of these tools.

Extreme Value Theory (EVT) by and large is dealing with the extreme deviations from the median of probability distributions, in order to estimate the probability of events that are more extreme than anything observed in the past. Two approaches mainly exist:

- deriving max/min series as an initial first step, that is generating an "Annual Maxima/Minima Series" (AMS) usually leading to the use Generalized Extreme Value Distribution as the most appropriate distribution to be fitted. Very often the number of relevant random events within a year usually is quite limited if any, very often analyses of observed AMS data lead to other distributions.
- isolating the values of a normal series that exceed a threshold, usually referred to as "Point Over Threshold" (POT) method and yet again can lead to a few only or no values at all being extracted in any given year. The analysis involves fitting one distribution for the number of events in a basic time period (typically Poisson) and a second distribution for the size of the resulting POT values (typically a Generalized Pareto Distribution).

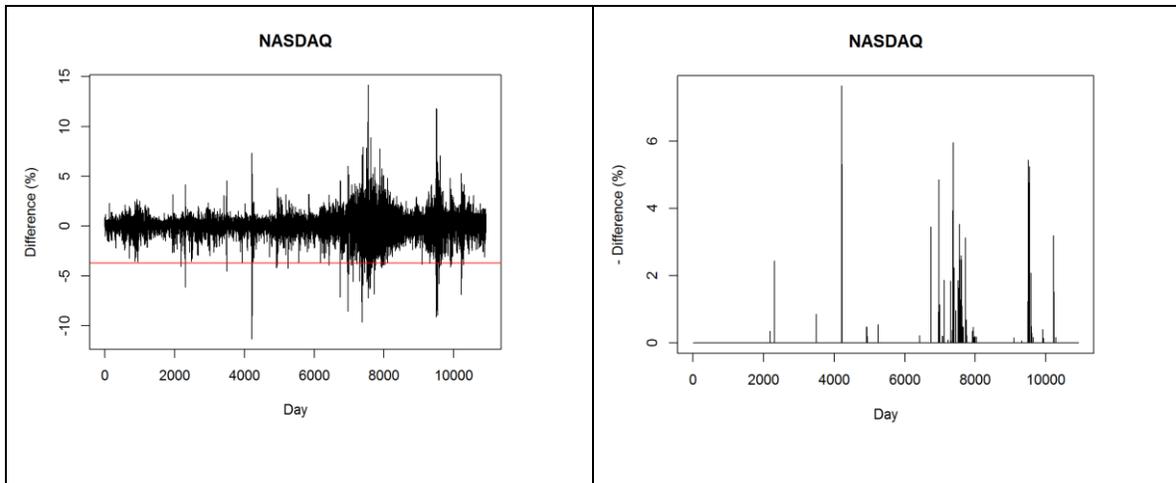
This latter approach clearly resembles in principle the logic of the Croston decomposition, however Croston's approach is not trying to reconstruct the distributions of the two independent series (inter-demand arrivals and demand volumes) rather than just forecast the two independent series with simple and very robust exponential smoothing approaches and then combine those two forecasts into a ratio forecast for future demand per period.

**Implications for practitioners: forecasting and foresight Black swans via intermittent demand methods through the non-overlapping temporal aggregation method-improving framework**

We have already advocated in the previous section for the use of intermittent demand forecasting methods. Then we can further advocate that for a short term forecasting task Naive is an appropriate approach while for the longer-term forecasting and foresight horizons SBA might be a more appropriate one, following selection protocols discussed in the literature (Petropoulos et al., 2014; Syntetos et al., 2015). These methods have been proven empirically to work better through the ADIDA non-overlapping temporal aggregation method-improving framework (Nikolopoulos et al., 2011).

***The data***

We have collected for the NASDAQ index, Daily Returns  $(=100\%*(CP_t-CP_{t-1})/CP_{t-1})$ , since the origination of the index in 5th of Feb 1971 till the 13th of June 2014, totaling more than 10500 Daily Returns from the respective trading days of the last 44 years.

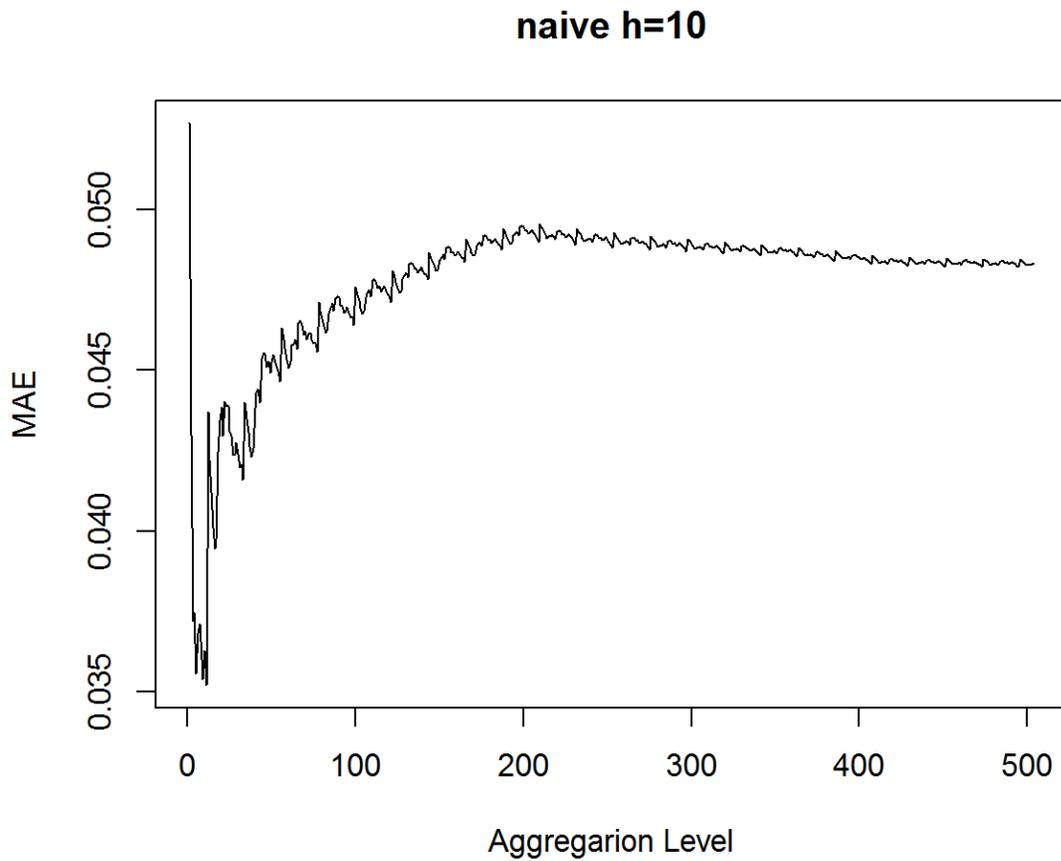


**Figure 3.** NASDAQ Daily returns from 5th Feb 1971 to 13th June of 2014: Negative Daily returns of more than 4% are illustrated on the right part of the graph.

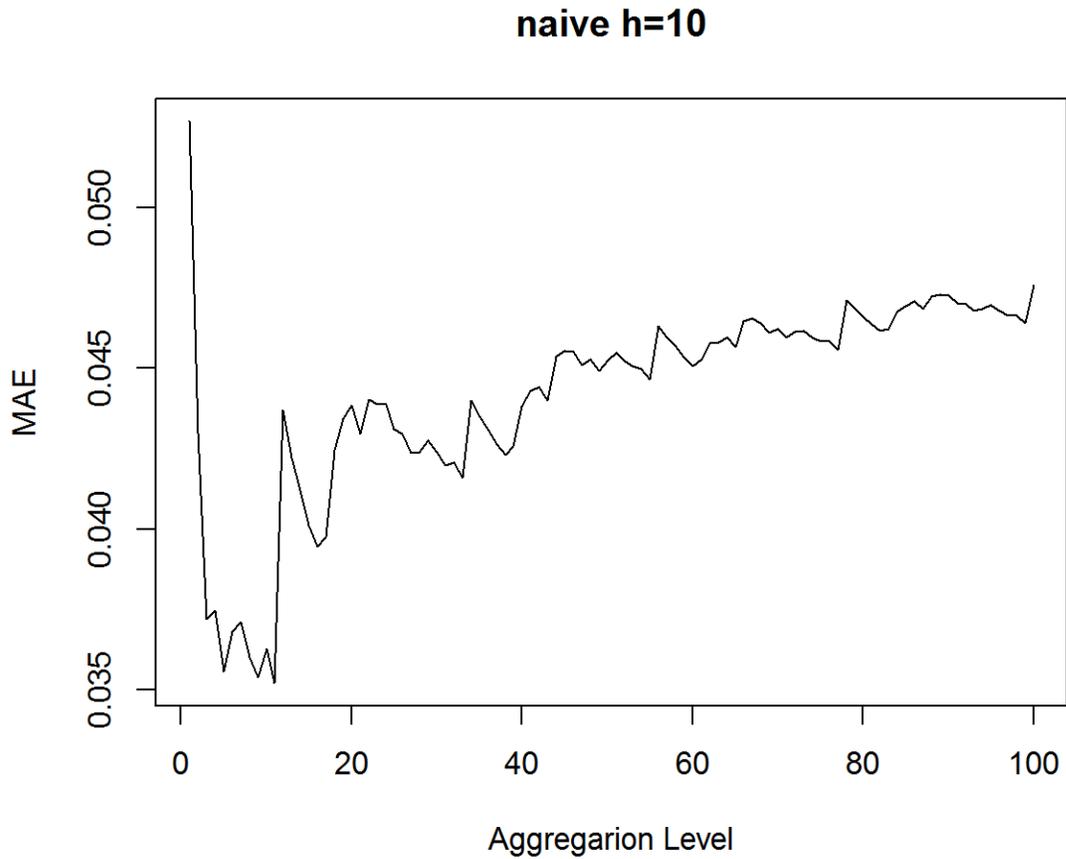
We get points under threshold - this includes the daily returns that are negative (where money are lost!) and over 5 standard deviations from the mean, that sets a threshold of about -4%: the resulting series (mirrored) is illustrated on the right part and is an 'intermittent demand' in nature series, a series illustrating all the days in the market where significant amounts of money were lost and as such reflect very risky periods of time. A cumulative forecast in such a context would proxy aggregate risk over a long period of time and this is exactly what we will try to forecast...

### ***Forecasting***

We illustrate forecasting for 2 calendar weeks ahead (10 trading days). We apply Naive over the ADIDA temporal aggregation framework and measure the average MAE over a 10-day window. This is illustrated in figures 4 and 5 (zoomed).



**Figure 4.** NASDAQ Negative Daily returns (over a threshold of -4%) forecasted with ADIDA(Naive). The graph illustrates the evolution MAE over a 10-day period along various temporal aggregation windows.

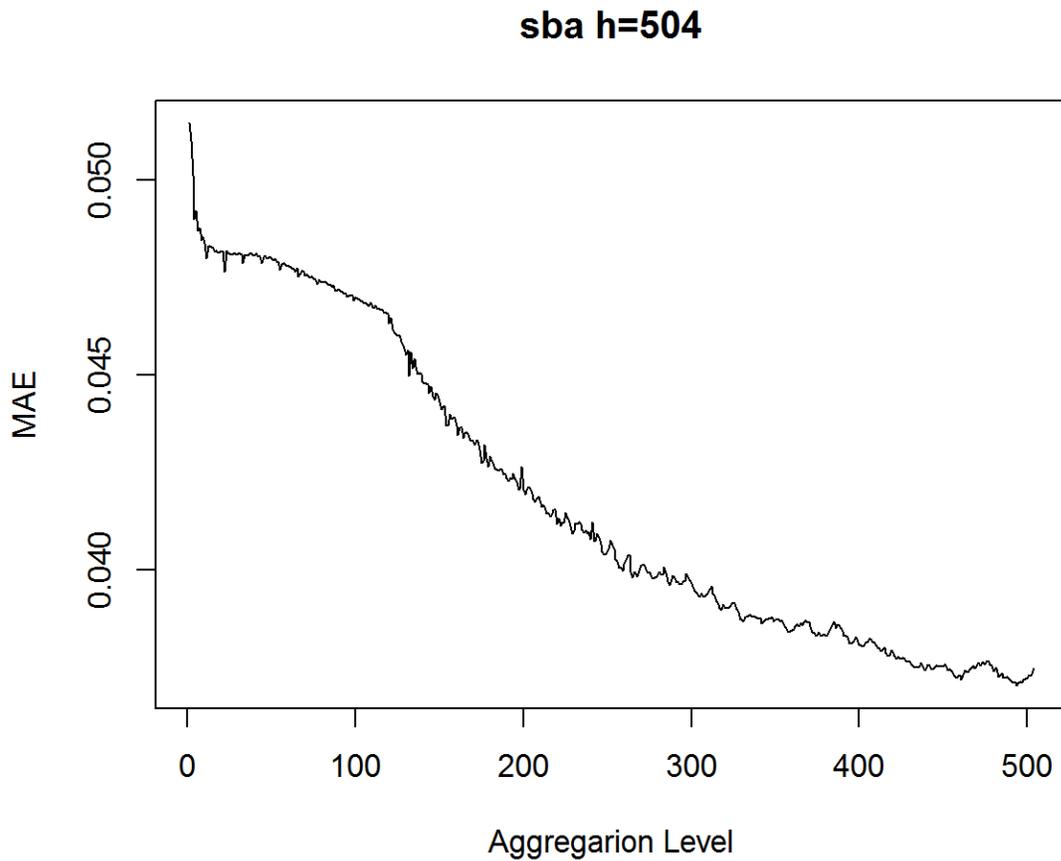


**Figure 5.** NASDAQ Negative Daily returns (over a threshold of -4%) forecasted with ADIDA(Naive). The graph illustrates the evolution MAE over a 10-day period along various temporal aggregation windows [Zoomed for a max of a window of 100 days]

It is clear that the forecasting performance of the Naive method is improved as the temporal aggregation window increases, and an empirical minimum error seems to be achieved for a window of about 10 days that matches the targeted forecasting horizon.

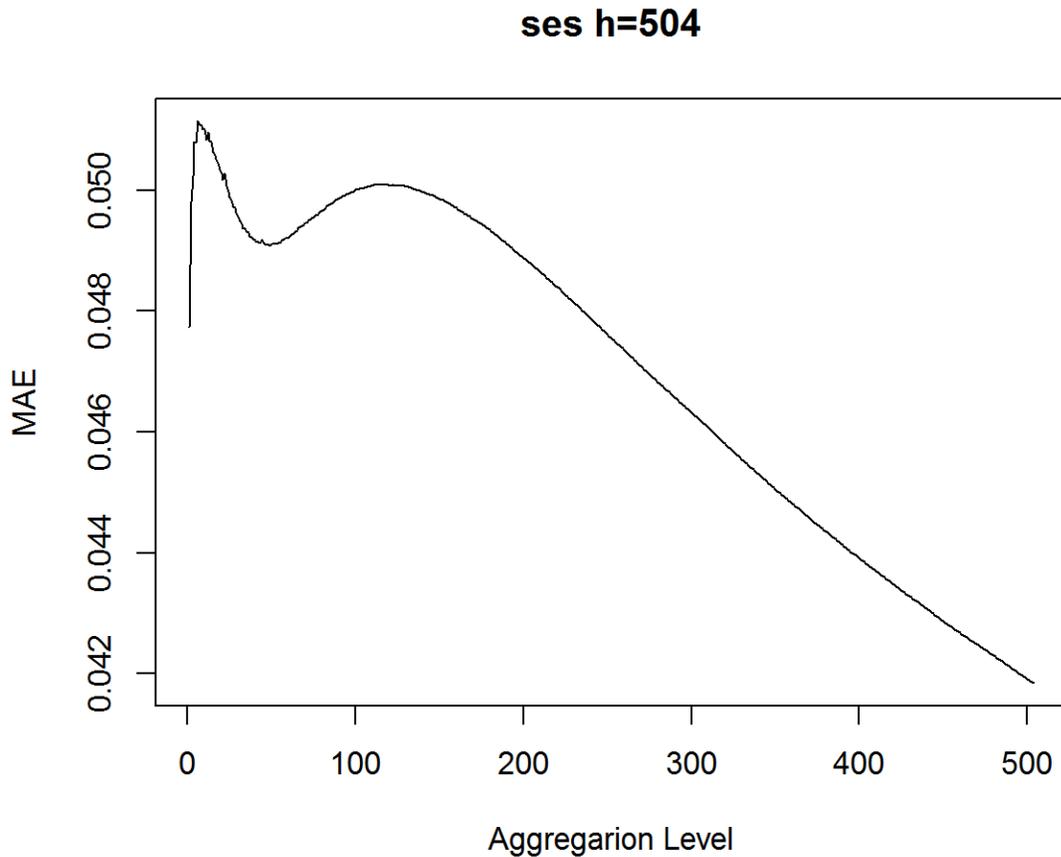
### ***Foresight***

We illustrate foresight/long-term-forecasting for 2 calendar years ahead (504 trading days). We apply SBA and SES over the ADIDA temporal aggregation framework respectively and measure the average MAE over a 504-day window. This is illustrated in figures 6 and 7.



**Figure 6.** NASDAQ Negative Daily returns (over a threshold of -4%) forecasted with ADIDA(SBA).

The graph illustrates the evolution MAE over a 504-day period along various temporal aggregation windows.



**Figure 7.** NASDAQ Negative Daily returns (over a threshold of -4%) forecasted with ADIDA(SES). The graph illustrates the evolution MAE over a 504-day period along various temporal aggregation windows.

It is clear that the forecasting performance of both the SBA and SES keeps on improving as the temporal aggregation window increases, and an empirical minimum error does not seem to be achieved. Our simulation could not be extended for aggregation of more than 504 trading days and as such there is still chances that there was an optimum level of aggregation to be identified about this aggregation level. Also SES has a local minimum for a window of about 10 days that matches the targeted forecasting horizon much earlier on about an aggregation window of 50 days but this is not the global one. In both cases however ADIDA helps the performance of both SES and SBA.

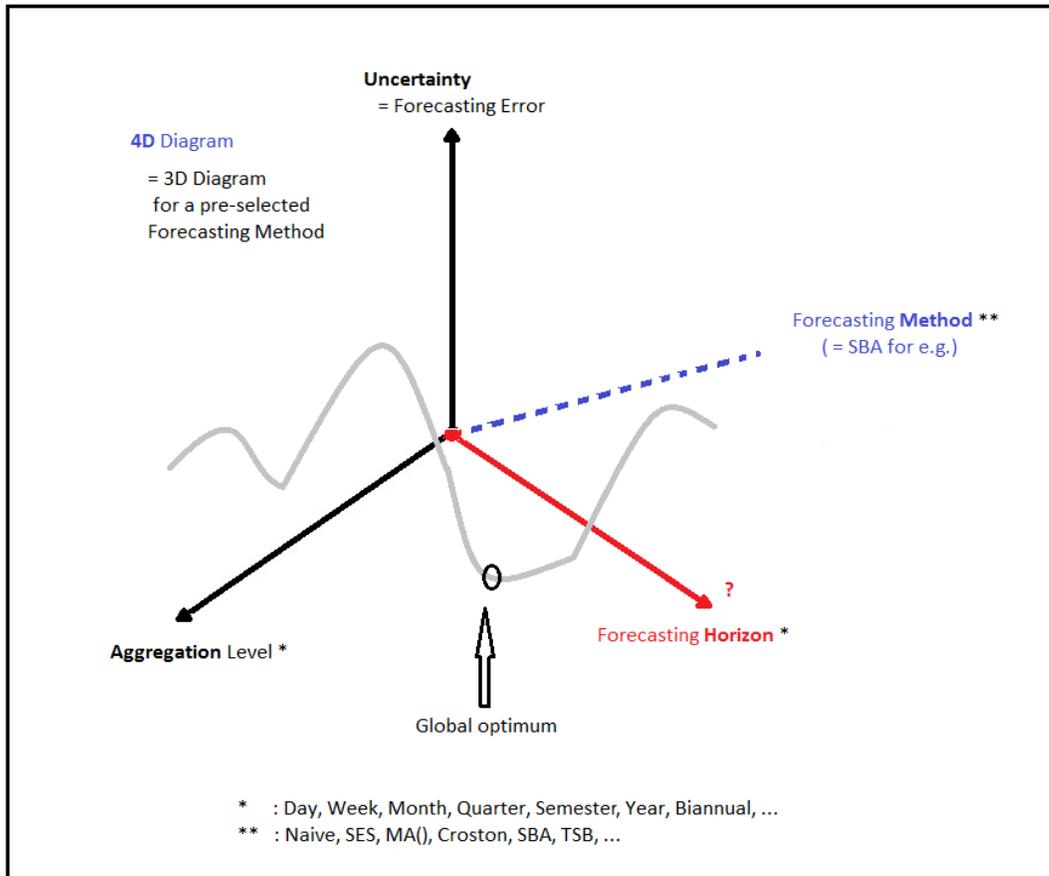
## Strategic planning in the presence of Black Swans

Ok... we do have the *forecasting horizon* set in any problem. This may be imposed by our customers, our suppliers, or any other aspect of the supply chain; sometime it is imposed by our IT systems. But is this the horizon where our forecastability is maximized? and thus that minimizes the uncertainty that prevails our forecasts and respective decisions?

If we follow this argument, we may be able to rethink about the forecasting/foresight horizon we could be using in our application; In other words: "*Re-thinking of the Strategic Planning horizon*". What if we match it to the forecasting/foresight horizon that minimizes the forecasting/foresight error and thus the respective uncertainty... This becomes more important if we are in a context of high uncertainty, like in the presence of Black and Gray swans, where then we really 'battle against uncertainty' and any more improvement in forecasting accuracy could be proved critical.

One way to do this is by running an exhaustive simulation across: a) all forecasting methods (tailored for intermittent demand data), b) all forecasting horizons, c) all non-overlapping temporal aggregation levels... and try measure the forecasting error we achieve (that maybe MAE or even better a 'super-metric' ranking many popular accuracy metrics e.g. MAE, MASE, MAPE, MSE, ME, etc). This is obviously a 4D problem and a attempt to illustrate that in figure 8 is given where the fourth dimension (forecasting method) is 'frozen' and set to SBA.

If there is such a long-term forecasting/foresight horizon (usually over our operational horizon) in which the forecasting accuracy is improved and thus uncertainty decreasing - whether or not that horizon is close-enough to the one we are using at the moment for strategic planning - we may wish to consider switching our decision and strategic planning horizon to that one... With the benefits being in the reduction of the prevailing uncertainty on the horizon where our decision are taken.



**Figure 8.** Re-thinking of the Strategic Planning horizon: what if we match it to the forecasting/foresight horizon that minimizes the forecasting error and thus the respective uncertainty... This is a 3D snapshot (with Forecasting Method = SBA) of a 4D graph (for all possible methods as a 4th dimension - illustrated by the dotted-blue line).

## The future

This is a methodological proposition we do in this research paper. We do not attempt to show superiority of our approach over the existing 'managerial decision-making' status quo: EVT for forecasting and 'Imposed Forecasting horizons' for Strategic Planning. We do however emphasize the *applicability, simplicity* and *robustness* of our proposition that even in the expense of some optimality in our overall solution, might prove very appealing among practitioners....

## References

- Croston, J.D. (1972). "Forecasting and stock control for intermittent demands". *Operational Research Quarterly* 23: 289-303.
- Leadbetter, M.R. (1991). "On a basis for 'Peaks over Threshold' modeling", *Statistics & Probability Letters* 12 (4): 357-362.
- Nikolopoulos, K., Syntetos, A.A., Boylan, J., Petropoulos, F., and Assimakopoulos, V. (2011). "An aggregate-disaggregate intermittent demand approach (ADIDA) to forecasting: an empirical proposition and analysis". *Journal of the Operational Research Society* 62: 544-554.
- Petropoulos, F., and Kourentzes, N. (2014). "Forecast combinations for intermittent demand". *Journal of the Operational Research Society*. DOI 10.1057/jors.2014.62.
- Petropoulos, F., Makridakis, S., Assimakopoulos, V. and Nikolopoulos, K. (2014). "'Horses for Courses' in demand forecasting", *European Journal of Operational Research* 237 (1): 152-163.
- Syntetos, A.A. and Boylan, J.E. (2001). "On the bias of intermittent demand estimates2. *International Journal of Production Economics* 71 (1-3): 457-466.
- Syntetos, A.A., Boylan, J.E, and Croston, J.D. (2005). "On the categorization of demand patterns". *Journal of the Operational Research Society* 56: 495-503.
- Syntetos A.A., Babai M.Z. and Gardner, E.S. Jr. (2015). "Forecasting Intermittent Inventory Demands: Parametric Methods vs. Bootstrapping". *Journal of Business Research*, forthcoming
- Taleb, N.N. (2007), *The Black Swan: The Impact of the Highly Improbable*. New York: Random House and Penguin.
- Willemain, T.R., Smart, C.N., and Schwarz, H.F. (2004). "A new approach to forecasting intermittent demand for service parts inventories". *International Journal of Forecasting* 20: 375-387.