CS-BIGS 2(1): 1-8 http://www.bentley.edu/csbigs/vol2-1/aragon.pdf

Estimation of Air Traffic Loss at Toulouse Blagnac Airport after September 11, 2001 Attacks

Yves Aragon
GREMAQ – Université Toulouse I, France

Sandrine Gnassou
SOGETI HighTech, Blagnac, France

Given a time series that reacts to an intervention, this article illustrates two points: (1) how simulations, either based on a normality assumption or through bootstrapping, can help us measure the impact of the intervention not necessarily on the mean but on functionals of the time series; and (2) how backcasting the series can help finding the time interval necessary for the series to recover a regular dynamic after the intervention. The methodology has been developed using an ARIMA model, but could equivalently be based on alternative models such as basic structural models, or on exponential smoothing.

1. Introduction

The September 11 attacks on the United States and the ensuing security measures generated perturbations on air transportation all over the world. The goal of this article is to study the incidence of these attacks on the traffic at Toulouse-Blagnac airport in the South-West of France. The data that have been made available to us (series of arrivals and departures) encompass the period of January 1993 to October 2007. We consider the traffic (departures and arrivals) of scheduled and irregular flights merged. Our goal is to measure the impact of the attacks of September 11, 2001, on this air traffic. An intervention model would consist in modelling the transition from the mean level before the event to the mean level after, under the assumption that the dynamic of the error remains unchanged before and after the event. Here, we adopt a point of view which, in some sense, is opposite to the one just mentioned. We regard the situation afterwards as data and examine the difference between this reality and what the traffic would have been in the absence of this event. We have to assume that a good forecast of what should have happened can be derived from the series before the event.

We begin with a descriptive examination and a comparison of the series before and after the event. We consider also the link between arrivals and departures. Then, we search for a suitable model to predict the series from its observations before September 2001 and, backcasting the series from the end of the observation period, determine the time, after September 2001, needed for the series to reach again a regular dynamic. Lastly, using the model obtained for the period before September 2001, we carry out a large number of simulations of the series beyond this date. The distribution of the loss, that is of the difference between simulations and the realization, gives a measurement of the September 11 influence on the activity of the airport.

2. Air Traffic at Toulouse-Blagnac

The airport of Toulouse-Blagnac is the fourth largest provincial airport in France with 5,956,552 passengers in 2006 (compared to 5,612,559 passengers in 2004 and 5,799,108 in 2005, a rise of 2.71% from 2005 to 2006), behind Nice, Lyon and Marseille. It is however the most important provincial freight platform with 58,720 tons handled in 2006, an increase of 4.1% over the level in 2005. A similar rise has been noticed on all fronts, including both domestic traffic (+2.5%) and international traffic (+2.8%). The airport is served by 46 airlines.

Figure 1. Blagnac Airport.

Scheduled flights include:

– fifteen national destinations: Ajaccio, Bastia, Brest, Clermont-Ferrand, Lille, Lyon, Marseille, Metz/Nancy, Mulhouse/Bâle, Nantes, Nice, Paris/Charles de Gaulle, Paris/Orly, Rennes and Strasbourg;
– twenty European destinations: Amsterdam, Belfast, Birmingham, Bremen, Bristol, Brussels, Dusseldorf, Frankfort, Geneva, Hamburg, Leeds/Bradford, Lisbon, London/Gatwick, Madrid, Malta, Manchester, Milan/Malpensa, Munich, Rome and Zaragossa ;
– seven intercontinental destinations : Algiers, Casablanca, Marrakech, Montréal/Pierre Trudeau (ex-Dorval), Oran, La Réunion, and Tunis.

In terms of non-scheduled flights, destinations such as Senegal, Andalusia, Austria and Ireland are served from Toulouse Blagnac International Airport.

2.1. Link between arrivals and departures

The yearly traffic, sum of arrivals and departures, is given in Table 1.

As Figure 2 shows, arrivals and departures are tightly linked each month. The vertical line indicates September 01. After this date we can observe a shift in the growth of the traffic and, after some rather irregular months, a modification of its dynamic. We will examine in the next section the length of time needed for the series to find a stable behavior again. We now focus on the link between arrivals and departures. In order to understand this link, we perform an Ordinary Least Squares (OLS) regression of monthly arrivals over departures and identify the residuals.

Table 1. Yearly passenger traffic at Blagnac Airport (thousands of passengers)

Year	Arrivals	Departures
1993	1,563.0	1,555.3
1994	1,639.8	1,631.5
1995	1,834.3	1,833.1
1996	2,042.6	2,045.9
1997	2,148.4	2,151.2
1998	2,305.0	2,307.1
1999	2,484.8	2,489.4
2000	2,625.4	2,622.2
2001	2,592.3	2,594.6
2002	2,641.1	2,647.4
2003	2,631.2	2,626.4
2004	2,780.2	2,783.1
2005	2,870.5	2,878.5
2006	2,940.4	2,954.0

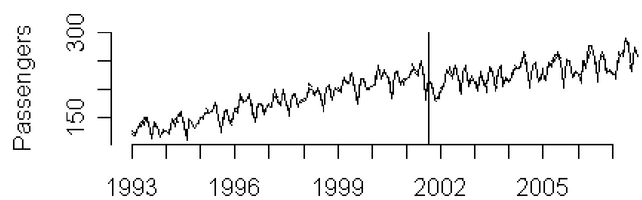

Figure 2a. Arrivals and Departures

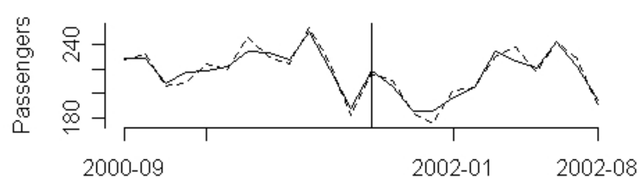

Figure 2b. Arrivals (solid) and Departures (dashed) around September 11, 2001.

Examining the autocorrelation function (ACF) of the OLS residuals and the Partial Autocorrelation Function (PACF) which almost vanishes after Lag 12 (see Figure 3), we choose a Seasonal Autoregressive Moving Average (SARMA) model to fit these residuals. We call "noise" the residuals of this second model. If the SARMA model is correctly chosen, this noise should be white noise.

Starting from this observation, in a second step, we estimate an ARMAX model (the X refers to the presence of an independent variable in the model) of arrivals; precisely we fit a linear model for arrivals against departures with a SARMA error. After some

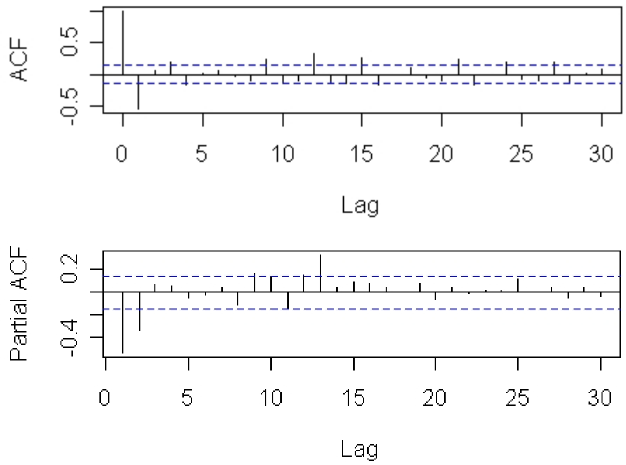

Figure 3. Autocorrelations of OLS residuals: ACF and PACF.

exploratory steps we obtain a SARMAX(1,4)(2,2)$_{12}$ with some coefficients constrained to 0. In this notation, (1,4) refer to the AR and MA parameters for the error series, and (2,2)$_{12}$ to the AR and MA parameters for the seasonal component of the error series. Table 2 displays the results of estimating this model, and the resulting equation is given below. Classical tools such MINIC (Minimum Information Criterion) which suggest tentative orders for the autoregressive and the moving average terms are not very helpful when seasonality is present[1]. The estimation and the t-statistics are given in Table 2, and the equation is displayed next.

$$Arrivals_t = 10.363 + 0.969 Departures_t + \frac{(1-0.330B^2-0.186B^4)(1-0.704B^{24})}{(1+0.776B)(1-0.371B^{12}-0.561B^{24})} Z_t$$

Here B is the Backward or Lag operator defined by $B Z_t = Z_{t-1}$, $B^2 Z_t = Z_{t-2}$, etc. The estimated variance of Z_t is 13.7. The quality of the fit is measured through the whiteness of the noise Z_t. The Ljung-Box statistic, whose p-value is plotted on Figure 4, is based on the chi-square distance between a vector of *k* empirical autocorrelations and their theoretical counterpart which is 0 if the process is white noise. We conclude from Figure 4 that the Z_t error of this model can be considered as white noise. We see that the intercept is significantly different from 0 while the t-statistic to test that the coefficient of the variable on departures is unity, equals (0.969 – 1) / 0.011= -2.8182. This allows us to conclude that there are significantly less arrivals than departures each month.

Table 2. ARMAX model of arrivals against departures

Coefficient	Value	Standard error	t-statistic
ar1	-0.776	0.056	-13.86
ma2	-0.330	0.085	-3.87
ma4	-0.186	0.075	-2.49
sar1	0.371	0.078	4.75
sar2	0.561	0.089	6.29
sma2	-0.704	0.081	-8.68
Intercept	10.363	4.163	2.49

A rough check of non existence of unit root in the autoregressive non seasonal part is given by the value of the autoregressive equation at 1 : 1+0.776= 1.776 which is far from 0 and for seasonal unit roots by the value of the seasonal part at 1 : 1 - 0.371- 0.561 = 0.068 which is not too close to 0. We can consider that the series on departures and arrivals co-integrate.

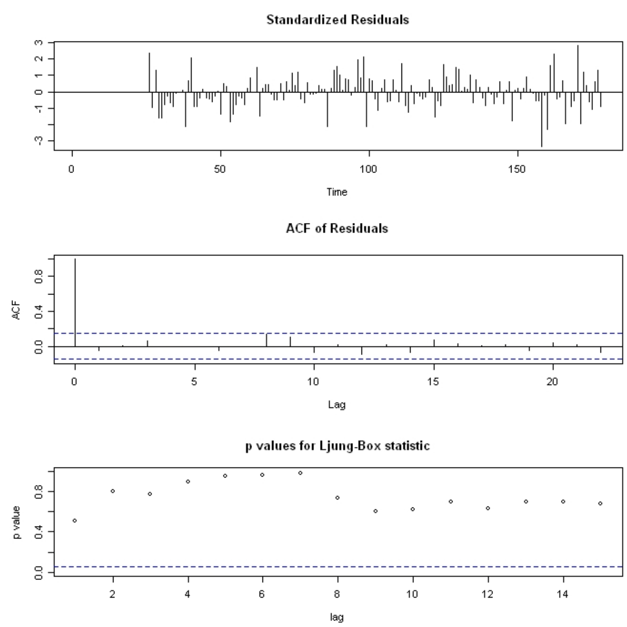

Figure 4. SARMA errors Z_t in the regression of arrivals over departures.

2.2 Descriptive study of the traffic series

From now on, we consider the series of passenger traffic at Toulouse-Blagnac airport, that is the sum of arrivals and departures each month, measured in thousands of passengers. We observe that it presents an increasing tendency noticeably attenuated after September 2001. The series also has a marked seasonality (Figure 5). Differences between monthly activities can be best understood on a *monthplot* (Figure 6). On this plot, each

[1] The minic method was proposed by Hannan and Rissanen in 1982. A practical presentation of it can be found in Box and al. (1994).

time plot corresponds to a month over all the available years. The horizontal bar is the mean of the subseries. August is a low and regular activity month. We see also that the influence of 9/11 is very important on the months such as December, January, February and May while it is light on the months of July and August

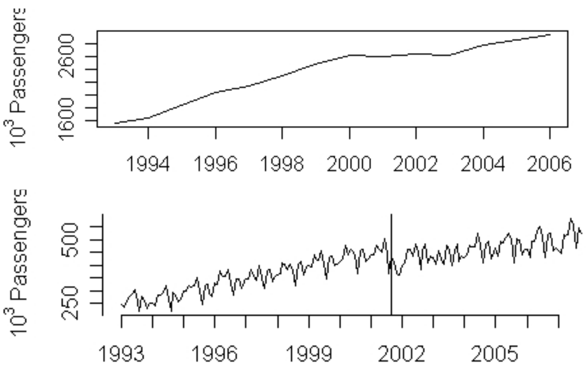

Figure 5. Yearly (top) and monthly (bottom) passenger traffic

The normality of the series is examined through the Jarque-Bera normality test. For the full series, the series before 09/2001 and the series after, the p-values are respectively: 0.04282, 0.1249, and 0.7346. To get a better understanding of the differences between before and after 9/11, we perform a Seasonal-Trend Decomposition[2] of the series (Table 3), restricting the series before to Sept 1996 – Aug 2001 and the series after to Oct 2001 – Sept 2006, so that both series have the same length. The interquartile ranges (IQR) of the components are given here in Table 3.

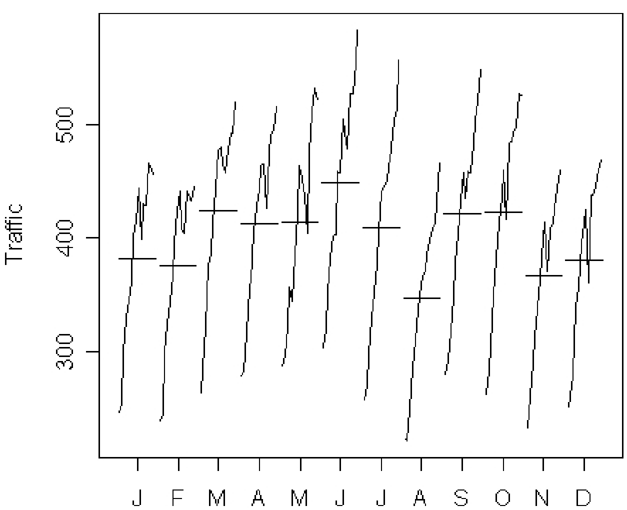

Figure 6. Monthplot of traffic at Blagnac airport

Table 3. Seasonal-Trend Decomposition of the arrivals; interquartile ranges

	Seasonal	Trend	Remainder	Data
Before	23.733	65.54	7.6294	77.157
After	46.336	40.32	12.758	58.121

We observe that the seasonal IQR before is half the seasonal IQR after and the trend IQR before is 1.5 times that after. Thus we suspect that it is rather difficult to find a unique model able to capture the dynamic of the full series with only minor modifications from one period to another. The difference between the two periods is also apparent on the empirical autocorrelation functions (Figure 7. ACF of the traffic series, before and after 9/11.). The series before is clearly non stationary while the series after seems more stationary.

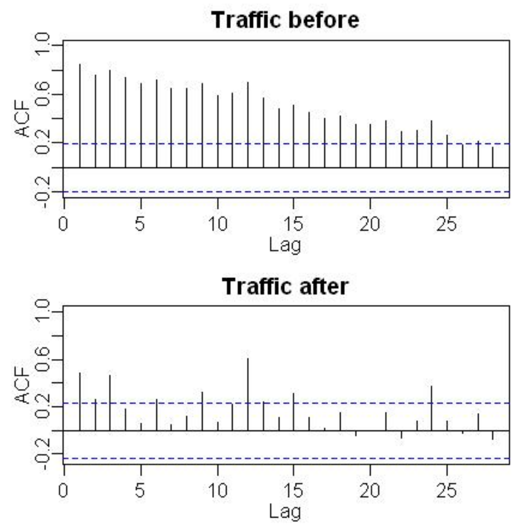

Figure 7. ACF of the traffic series, before and after 9/11**.**

3. Modeling the traffic time series

3.1 Traffic before September 2001

In order to evaluate the stability of our model specification we first model the series over the period 01/1994 - 08/2000, then, predict the 12 next months and last, compare the prediction with the observed series. After some trials we arrive at a SARIMA(1,1,2)(0,1,1)$_{[12]}$ specification (displayed in the equation below) where the first order MA term is constrained to zero. We perform a Maximum Likelihood Estimation of the model; the Portmanteau test does not reject the whiteness of the noise but a Jarque-Bera normality test on the residuals gives a p-value of 0.007918. So, although the series does not display any obvious heteroscedasticity we try to model the logged series.

[2] See Cleveland et al. (1990) and the stl() function of R.

We obtain the following SARIMA model and estimated error variance:

$$(1-B)(1-B^{12})(1+0.769B)\log(y_t)$$
$$=(1-0.525B^2)(1-0.438B^{12})Z_t$$
$$Var(Z_t)=0.00126$$

The p-value for the normality test of residuals is now 0.1236. To understand the predictive ability of this model for the interval September 2000 to August 2001, we simulate 1000 forecasts of these months and determine the percentile of the observed traffic with respect to the distribution of the corresponding forecasts. The simulation can rest on the normality assumption of the residuals on the log model or on bootstrapping the residuals. Results are summarized in Table 4; the tenth and ninetieth percentiles of the simulated predicted series, as well as the observed series are displayed in Figure 8. Table 4 contains percentiles of the observed value: (1) when the series is not transformed and the estimation is achieved by Maximum Likelihood under the erroneous assumption of normality and errors are drawn from the estimated white noise distribution, (2) under the same assumption as (1) but with errors sampled with replacement in the observed residuals, (3) same steps but the series is log transformed, (4) under the same transformation as (3) with errors drawn from the estimated white noise distribution.

We see that log based estimation gives usually less extreme percentiles for the observed series. We now re-estimate the model using the series from Jan 1994 to Aug 2001. The results of both estimations are summarized in the following table (Table 5). In order to evaluate the stability of the model we examine the position of the estimates from the first model (estimation 1), with respect to the 80% confidence intervals obtained from the second model (estimation 2). We see that estimations are clearly within the confidence intervals, thus we may consider that the model is rather stable.

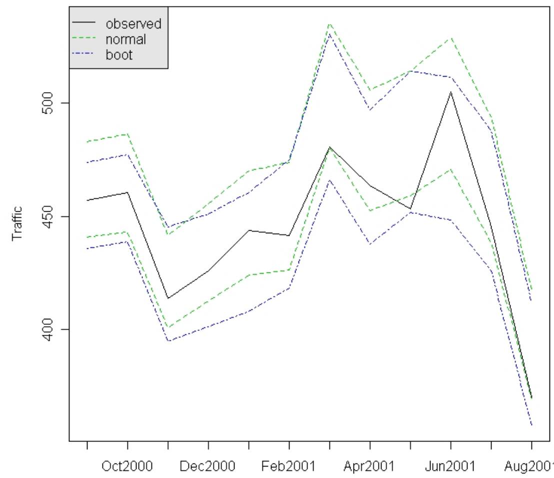

Figure 8. 10% of 90% Quantiles for Prediction and Realization of Traffic, September 2000 to August 2001.

3.2 Traffic after September 2001

We examine now the series from October 2001 to October 2007. The time plot (Figure 2) suggests that it takes some time to go back for the traffic to return to a rather regular evolution. In order to find the date until which the series is not regular, we would like to fit a model on the regular part, then backcast the series until October 2001 and examine the discrepancy between the realized and predicted series. As is known and easily checked, if a series follows some ARMA model, the series obtained by time reversion follows the same model. Thus a simple way of achieving the desired backcast is to reverse the time, adjust a model on the beginning of this series and predict the end, that is the series which finishes on October 2001.

Table 4. Percentiles of Sept. 2000-Aug. 2001 observed traffic in simulated distributions from different prediction methods

Month	1	2	3	4	5	6	7	8	9	10	11	12
Normality based (1)	30.5	40.0	10.5	27.7	54.1	15.0	2.6	15.3	0.3	84.4	4.7	0.9
Bootstrap based (2)	35.4	39.8	19.2	35.6	55.4	21.3	8.6	20.9	1.5	77.1	11.3	3.7
Log (3)	45.3	47.6	33.4	48.5	67.4	37.2	17.3	41.9	6.9	83.8	24.4	21.3
Log boot (4)	45.4	50.6	36.7	49.5	68.0	41.6	24.4	45.0	10.7	85.6	32.1	24.4

Table 5. Estimated models for traffic excluding (Est.1) and including (Est. 2) the period September 2000-August 2001

	Lag	Est. 1	s.e. of est. 1	Est. 2	s.e. of est. 2	Lower bound	Upper bound
ar1	1	-0.769	0.121	-0.741	0.107	-0.85178	-0.63087
ma2	2	-0.525	0.168	-0.486	0.148	-0.65624	-0.31666
sma1 (seasonal)	12	-0.438	0.132	-0.515	0.127	-0.68724	-0.34341
sigma2 (error variance)		0.00126		0.00114			

On the first 67 months of the reversed series, (from October 2007 to April 2002) we identify a SARIMA(1,1,3)(0,1,1)$_{12}$ model. Its maximum likelihood estimation is given by:

$$(1-B)(1-B^{12})(1+0.862B)y_t =$$
$$(1+0.434B-0.774B^2-0.660B^3)(1-0.506B^{12})Z_t$$

The estimated standard errors of the parameters are 0.107, 0.190, 0.193, 0.144, 0.168 and the estimated variance of the error is 199. The Jarque-Bera normality test of the residuals gives a p-value of 0.1029. Thus we predict the series for the next six months, that is for March 2002 until October 2001, and we compute the percentiles of the realization with respect to the distributions of the forecast for each lead time. Low or high percentiles suggest a bad predictive ability of the model.

From Figure 9, where time is reversed and ending on October 2001, and Table 6, we can consider that the series after 9/11 goes back to a regular behavior from January 2002 onwards.

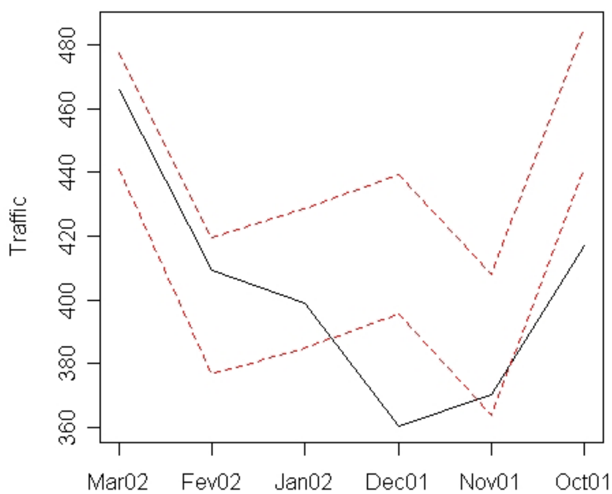

Figure 9. Backcast of traffic, and 80% prediction intervals; estimation based on October 2007 to April 2002.

Table 6. Backcast of traffic from March 2002 to October 2001

Date	Observed traffic	Prediction	S.E. of prediction	Percen-tile
2002 Mar.	465.65	459.25	14.317	0.67
2002 Feb.	409.14	398.23	16.684	0.74
2002 Jan.	398.98	406.79	16.938	0.32
2001 Dec.	360.46	417.32	16.985	0.0004
2001 Nov.	370.17	386.04	17.148	0.18
2001 Oct.	417.17	463.14	17.166	0.004

4. Measure of the influence of September 11 on the activity of the airport

4.1 Introduction

We want to measure the loss in traffic over a given period, caused by September 11. A classical tool for measuring the impact of an event on a series is intervention analysis. An example of intervention is the introduction of compulsory wearing of seat belts introduced on 31 Jan 1983 in Great Britain. This intervention causes a switch in the level of the series. The methodology known as *Intervention Analysis* is a tool to measure the effect of an event on the mean of a series. Intervention analysis is concerned not only with the estimation of the shift on the mean of a series but also with the description of the evolution of the mean of the series after an intervention. It assumes that, up to a non constant mean, the series follows an ARMA model, which is not modified by the intervention. In this framework, intervention analysis looks for a model for the mean which may be a simple step function or a linear combination of rational fractions of step and impulse functions, called *transfer function*. An intervention analysis can be seen as an ARMAX model where the mean follows a non linear model. Intervention analysis was popularized after the work of Box and Tiao (1975). Work by Abraham (1980) and Ledolter-Chan (1996) give some of the many applications of this technique. The results of an intervention analysis are included in the parameters of the estimated transfer function whose significance measures the influence of the intervention.

However our aim when we measure the impact of 9/11 on airport traffic is not to model the adaptation of the traffic after 9/11 but to estimate the loss of traffic for each month over one or two years or the loss of traffic for some particular period. We may want also to evaluate the distribution of the loss for some lead time. To achieve this goal we proceed as follows. First, we make the assumption that the estimated model before the event would have been valid after it if the event had not happened. Next, we consider the time series after the event as given, and we then look for a good model of the series before the event and forecast the series at some lead time. The difference between the forecast and the series is the loss for this lead time. From the adjusted model we may simulate a large number of trajectories and measure the loss for each. We thus get the distribution of the loss for each lead time. Losses can be cumulated for several times.

Our treatment assumes that the estimated model holds for the period of measurement. This model may be corrected with the help of economic forecasts established before the event. The simulation rests on the ability to simulate the white noise underlying the model. If the normality assumption does not seem to hold, it may be preferable to bootstrap the error rather than simulate it from a normal distribution. A similar treatment might be considered, backcasting the series before the event from estimation after it and considering the "loss" between the backcast and the realization before the event. Now we proceed to the different steps of this approach.

4.2 Prediction beyond September 2001

To predict the traffic until December 2003, we use the model estimated over the interval, January 1993 – August 2001, referred to as Est. 2 in Table 5. Of course such prediction is valid only if the dynamic of the series is not modified by an intervention or more predictable events, for instance, events based on macroeconomic forecasts.

The point prediction of the traffic series y_t, t = Sept 2001 to Dec 2003, which we abridge into y_{2001-2003}, is the vector of the estimated conditional means of these random variables given : (1) the series before and up to Aug 2001, (2) the estimated model on the log series and, (3) the normality assumption. However we want to predict not only the mean but also the conditional distribution of y_{2001-2003}. Given our framework, two paths may be considered: (1) rely entirely on the normality assumption to draw errors, simulate the log transformed series and take the exponential of it or (2) bootstrap the residuals and then simulate the paths. We shall consider both approaches. For the simulations we write the recurrence relation giving $\log(y_t)$ as a function of its past and the error with initial values taken, for the AR part, from the observed series and, for the MA component, the residuals from the estimation step. The error will be drawn from the estimated normal distribution in one case and bootstrapped from the residuals in the other case. For each month we thus get a number of simulated traffic measures that we can compare to the observed traffic. We therefore obtain a distribution of the monthly or yearly loss. Table 7 shows that bootstrap based and normality based simulations give rather similar results, which strengthens the normality assumption.

Yearly traffic forecasts for 2002, based on different hypotheses have been established by the *Institut du Transport Aérien* (ITA) and communicated to us by the Chamber of Commerce of Toulouse. They are: 4790, 5540, 5550, and 6060 thousand passengers. We draw them on the histogram of the simulated traffic (Figure 10).

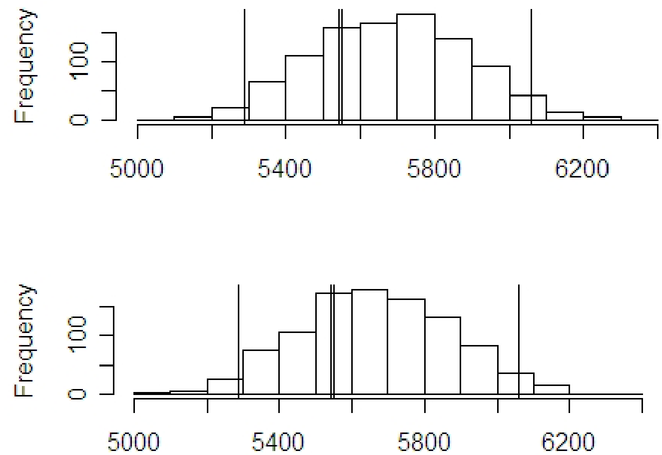

Figure 10. Distribution of simulated 2002 traffic; vertical bars from the left: observed traffic and 3 macro economic forecasts; a more pessimistic forecast is at 4790. Simulated errors in top graph, bootstrapped errors in bottom graph.

Thus, considering the simulation under a normality asumption, we find that the year 2002 loss lies between 250 000 et 541 000 passengers at a 50% confidence level. The 2002 loss is about 7% of the expected traffic.

Computational remarks. Computations were run with the ARIMA function in R (2008). Shumway and Stoffer (2007) warn us about the peculiar behavior of the R arima() function when the model contains an integrated part. The function simulate() was not used for the simulations because it does not seem to integrate initial conditions on the noise.

5 Conclusion

Intervention analysis tries to model the change on the mean of a series caused by an intervention. However consequences of an intervention might be of interest and may be difficult to derive from the result of such an analysis. In this article we suggest an alternative framework for studying the impact of an event on a series. First, we model the series before the event, next we simulate many trajectories from the estimated model for a time period starting from the date of the intervention, and we then compare the realized series with its intervention-free simulations. Contrary to intervention analysis, our approach does not assume any link between the model before and the model after the event. This approach may be used when some stable model can be assumed for the series. Simulating trajectories is a simple

Table 7. Simulated loss and traffic (thousand passengers)

		Loss - Simulations under normality assumption			Loss - Bootstrap simulations		
Year	Observed traffic	1st Qu.	Mean	3rd Qu.	1st Qu.	Mean	3rd Qu.
2002	5288.5	250	394	541	238	380	524
2003	5257.6	447	704	958	439	699	959
		Traffic - Simulations under normality assumption			Traffic - Bootstrap simulations		
Year	Observed traffic	1st Qu.	Mean	3rd Qu.	1st Qu.	Mean	3rd Qu.
2002	5288.5	5540	5680	5830	5530	5670	5810
2003	5257.6	5700	5960	6220	5700	5960	6220

tool to draw inference on linear and non linear functionals of the series or to circumvent non normal situations. The forecasting method should be not too sensitive to the most recent values of the series. In that respect, ARIMA models seem better suited than structural time series models such as Basic Structural Models (BSM) or exponential smoothing models.

Acknowledgements. The data necessary to the study were kindly communicated to us by the Chamber of Commerce and Industry of Toulouse, which also authorized their publication in support of the study.

Correspondence: aragon@cict.fr

REFERENCES

Abraham, B. 1980 , Intervention Analysis and Multiple Time Series. *Biometrika*, 67(1): 73-78.

Box, G. E. P., Jenkins, G. M., and Reinsel, G. C. 1994 . *Time Series Analysis: Forecasting and Control* (3rd ed.), Prentice-Hall.

Box, G.E.P. and G. C. Tiao 1975 , Intervention Analysis with Applications to Economic and Environmental Problems. *Journal of the American Statistical Association*, 70(349): 70-79.

Cleveland, R. B., W. S. Cleveland, J.E. McRae, and I. Terpenning 1990 , STL: A Seasonal-Trend Decomposition Procedure Based on Loess. *Journal of Official Statistics*, 6: 3–73.

Ledolter, J. and K.S. Chan 1996 . Evaluating the Impact of the 65 mph Maximum Speed Limit on Iowa Rural Interstates. *The American Statistician*, 50(1): 79-85.

R Development Core Team 2008 . R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0, URL http://www.R-project.org.

Shumway R. H., D. S. Stoffer 2006 . *Time Series Analysis and Its Applications With R Examples* (2nd ed.), Springer.

Shumway R. H., D. S. Stoffer 2007 http://lib.stat.cmu.edu/general/tsa2/Rissues.htm