

CASE STUDY 1

MAS2903: Introduction to Bayesian Methods

Dr. Lee Fawcett, Semester 2 2019–2020

SPEED CAMERAS AND REGRESSION TO THE MEAN

1. BACKGROUND

In 2007, official statistics revealed that 222,146 people were reported as injured as a result of road traffic accidents in Great Britain. Of these, 2,222 people were killed and 24,690 were seriously injured placing a huge economic and human cost on society. Road casualty reduction is therefore a key aim of government road safety policy with new road safety measures continually being tried and tested in an attempt to reduce the number and severity of casualties.

In 1996, a government report concluded that speed cameras could be an effective weapon in reducing casualty frequencies. However, the relatively high implementation and running costs were felt to prohibit their widespread deployment. In 1998, the government took the decision to allow local traffic authorities to recover the cost of installing and operating speed cameras from the revenues generated from speeding offences detected by the cameras. In April 2000 a two year pilot programme commenced involving eight “road safety camera partnerships”. Results at the end of the first year prompted the government to take an earlier-than-expected decision to introduce legislation in 2001 to enable the national rollout of safety camera partnerships across Great Britain.

2. CONTROVERSY

One of the first road safety camera partnerships was that linked to the Cleveland Police Force. The following newspaper headline was printed in the *Hartlepool Mail* in April 2003:

“Speed Cameras reduce road deaths by 70%”

Similar messages were reported in the local/national media praising the use of speed cameras, including the claim that

“Speed cameras save lives”

in one national newspaper. However, the rapid growth in speed camera activity, and subsequent increase in members of the public being punished for speeding offences, prompted a vigorous debate over the value of speed cameras. Opponents trying to discredit the operation of speed cameras (e.g. *SafeSpeed*) focussed on a range of issues in an attempt to have the scheme abandoned. In particular, they argued that speed cameras were

“...just another government tax”,

generating extra revenue for the treasury (see Figure 1); they also claimed to have the support of leading statisticians when they argued that the way in which the effectiveness of speed cameras is assessed is flawed.



FIGURE 1: Cartoon from *The Independent*, April 2001



FIGURE 2: Vandalism to Speed Cameras

3. ASSESSMENT OF EFFECTIVENESS

The road safety camera partnerships selected sites for speed cameras based on their previous casualty record. For example, the Cleveland Safety Camera Partnership had identified 72 “blackspots” – locations which had observed at least 5 “serious” casualties over a two year observation period (2000–2002). At each of these sites a speed camera was installed, and then the number of casualties was observed over a two year “treatment” period (2002–2004). The results seemed conclusive: in the “before” period, there were 361 casualties; in the “after” period, there were 108, giving the 70% reduction quoted in the *Hartlepool Mail* (see Section 2).

Organisations like *SafeSpeed* claimed such before-after studies were flawed, as they did not account for the phenomenon of **Regression To the Mean**, or **RTM**. In fact, they claimed that such studies exaggerated the effectiveness of speed cameras. Indeed, they had the support of leading statisticians and RTM made the headlines of leading national newspapers.

4. REGRESSION TO THE MEAN: A SIMPLE EXPERIMENT

The Northumbria Safety Camera Partnership was formed in 2001, and I was asked to provide statistical support in their assessment of the effectiveness of speed cameras across the northeast. I came across the debate for and against speed cameras during my research; in particular, I thought about the role of the RTM phenomenon and how we could account for this. One challenge was to explain to my colleagues in the Partnership what RTM actually was!

Suppose a group of students are asked to roll a die 20 times and count the number of sixes they observe. The two students who roll the highest number of sixes are then asked to roll their dice a further 20 times - but this time on top of a piece of paper; these two students once again record the number of sixes. I actually did this with a group of A level students at a "Maths & Stats engagement talk" at South Tyneside College recently, and obtained the following results:



Student	1	2	3	4	5	6	7	8	9	10
#6's	3	3	1	1	0	7*	4	5*	2	2
#6's						2		1		

Has the piece of paper caused a reduction in the number of sixes? Our "high-scorers" were students 6 and 8, and in total the piece of paper has reduced the number of sixes from 12 to 3 – a 75% reduction!

Of course the above experiment is a load of rubbish; we have selected two students with an unusually high number of sixes – in any repeat of the procedure, their number of sixes is bound to be lower! Now suppose each student is a potential speed camera site, and the number of sixes are the number of casualties in a "before" period". The piece of paper is the speed camera, which is applied to the students (sites) with an unusually high number of sixes (casualties). ***If we don't believe the piece of paper has reduced the number of sixes, why should we believe that speed cameras reduce the number of casualties?***

5. ACCOUNTING FOR RTM

So, the argument is that, because speed cameras are installed at sites with an abnormally high casualty record, these sites would probably see a reduction in casualties in the after period anyway, even without the cameras. Perhaps the cameras do have a role to play, but we shouldn't really attribute the *entire* reduction to the cameras – some of this might have happened anyway. Figure 3 below illustrates potential RTM at a speed camera site in Australia.




FIGURE 3: Casualties at the Tullamarine Freeway, Aus.

The Northumbria Safety Camera Partnership had recorded 154 casualties at 17 blackspots in South Tyneside in their "before" period". After the implementation of a speed camera at each, the number of casualties in the "after" period was 41. Rather than just compare before and after figures, I thought I would try to account for any RTM effect when analysing the data for the Northumbria partnership – surely RTM has made a contribution to this 73% reduction?

5.1 A statistical model for casualty frequency

Suppose we assume that the random variables Y_i , $i = 1, \dots, 17$, represent casualty frequencies at each blackspot site i in South Tyneside, and that $Y_i \sim \text{Po}(\theta_i)$. Also, let y_i be the (abnormally high) observed number of casualties from these Poisson distributions in the before period. How can we estimate the number of casualties that we would *normally* expect to see at each of our speed camera sites – i.e. when the casualty rate is *not* abnormally high?

5.2 Formulating a prior distribution for θ_i

A gamma prior distribution for θ_i might be a reasonable assumption, i.e.

$$\theta_i \sim \text{Ga}(g, h_i); \quad (1)$$

we cannot have a negative rate of casualty frequency, and the gamma distribution is defined over the positive real line (also, the gamma distribution is the **conjugate prior** for the Poisson – see Chapter 3 of lecture notes). Notice that the gamma shape parameter g is constant for all sites but the rate h_i is site-specific; this is common in applications such as this. But how do we choose the parameters g and h_i ?

Northumbria Police had recorded the number of casualties at many sites, not just the blackspot sites. We can treat the 67 non-camera sites as a "reference" set at which more "normal" casualty frequencies are observed. Analysing these data shows that the number of casualties depends quite strongly on the average observed speed of vehicles, with casualty frequency increasing as average speed increases. In fact, we can obtain the following simple linear regression equation, known as a **Predictive Accident Model** in the road safety literature:

$$Y_j = -4.288 + 0.157X_j + \varepsilon_j, \quad (2)$$

where X_j and Y_j are the average observed speed and number of casualties (respectively) at reference site j , $j = 1, \dots, 67$. The idea is to link our 17 blackspot sites to the reference sites through this predictive accident model. Equation (2) has been formulated from data at sites observing more "normal" casualty frequencies: we then apply this model to the average speed of vehicles observed at our blackspot sites to estimate a more "normal" casualty frequency for these sites, and we will use this estimate as our prior mean for θ_i in Equation (1), labelling this μ_i for each blackspot site i . Specifically, we use

$$h_i = g/\mu_i$$

in (1), which gives $E(\theta_i) = \mu_i$ and $\text{Var}(\theta_i) = \mu_i^2/g$. In fact, in our application, $g = 2.5$ is the "best" value to use (as provided by a road safety expert).

5.3 Obtaining the posterior distribution for θ_i

Using the above gamma prior for θ_i , $i = 1, \dots, 17$, results in

$$\theta_i | y_i \sim Ga(g + y_i, g/\mu_i + 1), \quad (3)$$

the mean of which is

$$E(\theta_i | y_i) = \alpha_i \mu_i + (1 - \alpha_i) y_i, \quad (4)$$

where


$$\alpha_i = g / (g + \mu_i).$$

Notice that the posterior mean is a weighted sum of the number of casualties we would usually expect to see at site i (μ_i) and the unusually high number of casualties we have actually observed at that site in the before period (y_i). This is often referred to as the **Bayes Linear Rule** (see Example 2.6 in lecture notes).

5.4 Assessing the effectiveness of the speed cameras

The idea is that instead of using a straightforward comparison of the number of casualties before and after the implementation of the speed cameras, we compare the after figure with the weighted sum given by Equation (4). We have not “thrown away” the unusually high before figure at each site i – notice y_i features in Equation (4) – but we “tone this down” a bit by combining it with what we might normally expect to see at this site. In other words, we have combined our **prior beliefs** about the casualty rate at each site i with the unusually high number of casualties we have actually **observed** at these sites, to form our **posterior beliefs** which are a synthesis of the two – i.e. a Bayesian formulation!

The difference between the number of casualties in the before period and the posterior mean given by Equation (4) can be seen as the amount by which we would expect casualties to reduce by anyway, even if no speed cameras had been introduced. The difference between the posterior mean given by Equation (4) and the number of casualties in the after period is then the reduction due to the cameras:



5.5 Results for speed camera sites in Northumbria

The Table below shows results for some speed camera sites in the Northumbria study.

Site	y_i	μ_i	α_i	$E(\theta_i y_i)$	$y_{i, after}$	Observed	After RTM
1	20	4.35	0.36	14.29	0	-20	-14.29
10	8	2.31	0.52	5.04	3	-5	-2.04
Tot.	154			90	41	-113	-49

For example, at site 1 we have observed 20 casualties in the before period. Actually, based on the average observed speed of vehicles at this site, we'd normally expect to see between 4 and 5 casualties. Combining the two via Bayes' Theorem gives just over 14 casualties. We conjecture that the number of casualties would have fallen from 20 to about 14 anyway, had no speed camera been used; thus, after accounting for this regression to the mean effect, the real effect of the speed camera has been to reduce casualties from about 14 to 0, rather than from 20 to 0.

Summing across all sites, we see that the number of casualties would have fallen from 154 to 90 anyway, even without the speed cameras; thus, we might say that the remaining reduction to 41 in the after period – 49, instead of the full 113 – are due to speed cameras.

Figure 4 shows the changes in our beliefs about the rate of casualties at site 1, after we observe 20 casualties in the before period. Notice how the posterior is a compromise between what we believe should normally happen at this site (based on the average observed speed of vehicles here), and the actual value observed.




FIGURE 4: Prior and posterior densities for θ_1 (dashed/solid curves); observed casualty frequency (vertical line).

6. SUMMARY

In this case study we have considered the Bayesian approach to analysis as a means of combining prior information about the expected casualty frequency at designated blackspot accident sites, with the actual observed casualty frequency at these sites.

- The Bayesian framework provides a means of accounting for the phenomenon of RTM
- Without doing so, and just comparing before and after figures, is bound to exaggerate the effectiveness of the speed cameras
- RTM has now become a recognised phenomenon in studies such as this, and the Bayesian approach outlined here has become an official part of the Government's assessment procedures

Fawcett, L. and Thorpe, N. (2013). Mobile Safety Cameras: Estimating Casualty Reductions and the Demand for Secondary Healthcare. *J. App. Stat.*