REPORT BY MR FINNEY INTO MOBILE SPEED CAMERAS IN THAMES VALLEY - 8TH OCTOBER 2013

https://speedcamerareport.co.uk/08_mobile/

INTRODUCTION

Mr. Finney, a long time anti speed cameras campaigner, has published another report into the effectiveness of speed cameras in Thames Valley. I find his research of interest as the data being used was provided by the Thames Valley Safer Roads Partnership (TVSRP) while I was the Operations Manager in charge of day-to-day running of the partnership.

Dave Finney (DF): I have never campaigned against speed cameras. What I have campaigned for is:

1) Honesty regarding the effects of speed cameras on road safety
2) Speed cameras to be operated within scientific trials

Mr Finney and I have had many long conversations about the analysis of camera statistics over the years and I have voiced concerns about the methods used in the past. In order to inform the road safety community about the validity of this new research I offer the following summary of his latest report.

DF: I would like to thank Richard Owen (RO) and his team for checking and verifying the collision data, and for creating and providing the database of collisions at all speed camera sites in Thames Valley. It is this database that allows agreement on the basic data. The issue now is the evaluation of that data and, specifically, how to exclude the effects of RTM (regression to the mean) from the final results.

SITE SELECTION UNDER TVSRP

First of all it is important to note that not all of the mobile camera sites analysed as a part of the research were installed on the basis of high collision rates. Under the various rules of the National Safety Camera Programme[1] there was an exception allowed for up to 15% of enforcement time. These *exceptional* sites actually constituted the majority of mobile sites (64% as at the end of the Thames Valley Safer Roads Partnership on 31st March 2012).

These sites were selected using a 'Community Concern' programme put into place in 2004 and there was no requirement to demonstrate high KSI[2] or PIC[3] collision rates in order to be accepted. Often there were no recorded collisions prior to installation, as long as there was a perceived risk of high speeds and potential collisions.

DF: RO is correct to state that not all sites were selected following high collision rates. Of the 75 mobile speed camera sites in my report, 12 sites had no collisions in the identified SSP (Site Selection Period) prior to speed cameras. Fairly obviously, collision numbers could only go one way (up) and they did. There were 14 collisions in the first 3 years of mobile speed camera operations at these 12 sites.

Information produced at the end of the partnership for mobiles sites in the counties of Berkshire and Buckinghamshire[4] looked at changes in KSI collisions in the three years after installation versus the three years before. For core sites the reduction was 37 KSI collisions but for exceptional sites there was an increase of 1 KSI collision. This shows that any analysis of both site classifications with no differentiation is clearly flawed. Exceptional sites should be analysed separately from the *core* sites in order to investigate the true effect of enforcement.

DF: RTM is an effect of the site-selection process therefore, if "core" sites were often selected following a high number of KSI collisions, then a large reduction due to RTM could be expected. And if "exceptional" sites were often selected following a lower (or normal) number of KSI collisions, then there should be little or no change due to RTM.

The results RO presents are therefore consistent with the expected effects of RTM, rather than speed cameras.

Collision rates decreased at some sites and increased at others after mobile speed cameras were deployed. What matters is the overall effect and there is therefore no need to separate sites into TVSRP's two groups. My report measures the overall RTM effect at all of the sites where they deployed their mobile speed cameras.

1 There were annual 'handbooks' that told partnerships how to operate and this included rules on site selection which changed almost on an annual basis
2 Killed or Seriously Injured
3 Personal Injury Collision
4 Oxfordshire left the partnership in the summer of 2010 and all sites were decommissioned. Without further information about decommissioned dates it is impossible to include these sites in the analysis.

Mr Finney assumes that the camera site polygons used to analyse performance were the same used to carry out site selection which is incorrect. Many of the early (pre 2003) mobile sites were selected due to poor route performance e.g. '3 miles of the A404, Marlow Hill in High Wycombe' or because they were legacy locations previously used by Thames Valley Police prior to the creation of the partnership. When site selection was carried out through the use of cluster analysis these clusters were generally 500m in diameter, roughly half the size of the polygons used for collision performance analysis.

The difference in the collisions used to carry out pre and post-installation analysis is vitally important because of Mr Finney's assertions regarding Regression To Mean (RTM).

DF: Actually, the final results in my report do not assume that analysis and site-selection polygons were the same. Provided that the two polygons overlap such that the SSP can be identified (which clearly was the case), what matters is that the analysis polygons include the stretches of road where the cameras might influence motorists.

Different polygons might affect measurements of RTM though. The RTM effect should be greatest when analysing site-selection polygons and lower when analysing any other polygons. The levels of RTM measured in my report refer to the analysis polygons but, had site-selection polygons been analysed, higher levels of RTM would be expected.

STANDARD ANALYSIS

Within Mr Finney's report he quotes the standard before / after comparison at mobile speed camera sites. These calculations appear to be correct and are broadly in line with previous analysis carried out by the Thames Valley Safer Roads Partnership[5]. These results do not allow for any RTM however and an appropriate statistical method should be used (if possible with such as small sample size) to allow for this effect. TVSRP did allow for general improvements in road safety seen over the period during which cameras were operational to unveil the underlying trend at enforcement locations. This still showed the largest influence on collision frequency rates to be due to the presence of enforcement.

DF: The standard before / after comparison at camera sites produces results that include the effects of site selection (RTM), general influences ("trend") and the cameras. To find the effect of the cameras, the other effects have to be removed. TVSRP never did this, hence my report.

5 http://www.saferroads.org/2010/06/lifting-the-lid-on-safety-camera-stats/

ALTERNATIVE ANALYSIS

Notwithstanding the general problem of analysing different site categories there is a further fundamental problem with the analysis and techniques used in the reminder of the report.

The *four time periods theory* mentioned in section 8.5 of the report is seriously flawed. There is no way of knowing what periods were used for site selection at individual sites without referring to the original analysis which may have formed part of the Operational Case supplied to the DfT prior to sites being accepted into the programme. There may have been delays in commissioning a site which would lengthen the 'ASBiC period'. For a small number of sites there may have been a three year analysis but the most recent data used could have been anywhere between 4 and 24 months old. Simply assuming a 12 month period is inaccurate and misleading. This is highly important because Mr Finney's analysis of RTM effects specifically uses this short time period. This methodology also presupposes his conclusion – because he starts out selecting the period with the highest rate by definition. Of course it is likely to regress.

DF: My four time periods method is "*straightforward and logical*", according to the Department for Transport's head of statistics (9), and it was used by Professor Richard Allsop in his 2013 report for the RAC Foundation.

If RO could refer to his original analysis and email me a list of the SSP start and end dates that he used to select each site, I could check to see if the actual SSPs were all within the identified SSP for the group of 75 mobile speed camera sites in my report. That would be useful follow-up research.

There may have been delays commissioning fixed (Gatso-type) speed cameras (due to difficulties getting power to the sites etc) but my report only includes mobile speed cameras. No power or other installation work should have been required therefore there would be no need to delay the commission date after approval was granted. The time taken to get mobile speed camera sites commissioned (the ASBiC) should therefore have been reasonably consistent, and the evidence suggests that it was.

If data was sometimes 4 to 24 months late, then the database will be complete except for the most recent 4 to 24 months. At the 75 sites, data could only be missing therefore during the period after mobile speed cameras started operating. The increase in casualties at these sites, therefore, may actually be higher than found in my report.

The Site Selection Period defined in the report (SSP) has been calculated to last 2.5 years rather than the three years mentioned in the 'Standard Analysis'. There seems to be no explanation of why this specific period was used in its place. This period does covers the five worst 6-month periods in terms of total KSI collisions and seems to have been created entirely to support a theory proposed once the initial analysis had taken place. Also, using six month blocks instead of 12 month ones exposes the analysis to seasonal variations.

DF: We must not confuse "baseline" with "SSP". The baseline specified in the national criteria was "*the three years prior to the proposal being submitted*" but the SSP was chosen by the partnership. Clearly, if sites were selected using a 3 year SSP, but it took 6 months to prepare a proposal, some sites might not satisfy the criteria by the time the proposal was submitted. By using a 2.5 year SSP, the partnership ensured that every site they selected would still satisfied the 3 year baseline criteria when the proposals were submitted some 6 months later.

Fairly obviously, the SSP does cover "*the five worst 6-month periods in terms of total KSI collisions*". That is what the site-selection process will tend to produce and is good evidence that the SSP was those five 6-month periods.

Using six month periods can lead to seasonal variations but only when using odd numbers of them. The final results in my report use the PreSSP (5 years), the ASBiC (1 year), and after speed cameras (3 years). Because all three periods are multiples of whole years, there should be no seasonal variations in the final results.

He excludes slight injury collisions at sites, despite the fact that they were a large part of the netting off site selection criteria. Using KSIs makes his sample size pitifully small – only about 300 in total over more than 13 years. His stated reason makes little sense and has no basis in fact.

DF: The national guidelines suggested that both KSI collisions and all collisions should influence site selection and the SSP could have been identified from either. The same SSP can clearly be seen in the data for both KSI collisions and all collisions and this is yet further evidence that the SSP used to select these 75 sites was a duration of 2.5 years, ending 1 year before the mobile speed cameras were deployed.

He mentions the KSI reporting adjustment in Thames Valley, but his method of allowing for it is faulty. Firstly, he seems to be calculating the mean of percentage figures where the underlying sample size in each percentage varies in size, which would be arithmetically invalid (although his exact methodology is rather unclear from the text). Secondly, because the figures in each sample come from a different absolute time period (because installation dates vary by up to three years) many of the sample points must contain a mixture of pre and post adjustment KSIs. Thirdly, his principal argument rests on comparing the 'PreSSP' to the 'SSP', and he seems to perform that comparison without correction for the reporting adjustment.

DF: Because the 1999 serious-injury reporting adjustment affected KSI collision reporting across the whole of the Thames Valley area, it is one of the "general influences" that needs to be compensated for. And, because it was a sudden change, the standard approach could not be used. This was the primary reason I developed a new method using relative collision rates. My new method is based on the same concept as the standard approach, but it's more accurate and can compensate for the effects of general influences as they vary, and even when they change suddenly.

There is no indication of how he calculated the comparison to the baseline given that the baseline period is different for each site. Without this information it is not possible to comment on the accuracy of his method for comparing sites to the overall trend.

DF: The method is explained in my report and on my RTM page and was also used by Professor Richard Allsop in his 2013 report for the RAC Foundation. To calculate relative collision rates, the collision data in each cell in the TVSRP database was divided by the number of collisions of the same severity that had occurred in the same calendar year across the whole Thames Valley area, and was then multiplied by 100. The results were entered into a second database. By comparing relative collision rates in the second database, all general influences are automatically compensated for, including the 1999 serious-injury reporting adjustment.

The final flaw that requires addressing is the most damaging of all, namely the measurement of RTM. RTM is a known phenomenon which has been tackled in many previous reports. The effects have been quantified in previous award-winning research by Professor Richard Allsop[6] which indicated a reduction in KSI collisions of between 15 and 30% at mobile sites once RTM and other trend factors have been taken into account.

DF: The 2010 RAC Foundation report by Professor Richard Allsop is reference 3b in my report and it uses part of the FTP method. Data before speed cameras was obtained from 6 partnerships and split into two periods: a three year baseline period and a pre-baseline period. On the assumption that the pre-baseline period was before the SSP, the RAC report produced an "...*indication of the possible scale of the contribution of RTM*" and found that RTM might be the largest factor at speed camera sites. However, if the assumption was wrong, RTM would tend to have been under-estimated.

The RAC Foundation report demonstrates the need to measure RTM, hence why this was done in my report.

6 http://www.racfoundation.org/research/safety/effectiveness-of-speed-cameras

Mr Finney's approach to dealing with RTM is to arbitrarily ignore most of the 'before installation' data by claiming it is tainted by selection bias. No attempt is made to say how much bias is included within those figures, they are all just ignored. All academic analysis of RTM includes a calculation of the random fluctuation within the dataset so Mr Finney's methodology is highly questionable.

DF: No data was arbitrarily ignored in my report. All four time periods were used as follows:

1) Measurements of RTM made use of the PreSSP, SSP and ASBiC periods.
2) The final results made use of the PreSSP, ASBiC and first three years of mobile speed camera operations.

I believe the academic analysis to which RO refers, such as the government's four year evaluation[7]), have used the EB (Empirical Bayes) method. The EB method has 3 major problems:

1) The EB method requires extra data not routinely collected (TVSRP did not collect this data)
2) The EB method is so complex that very few people understand it
3) The EB method can only produce estimates of RTM effects, it cannot fully exclude them (or measure them)

My FTP method has none of these problems. By simply excluding the data that influenced site-selection, RTM effects are fully excluded from the final results, without indications or estimates that might be prone to error.

7: http://speedcamerareport.co.uk/03_government_reports.htm

By limiting the RTM effect to a very short period of 2 ½ years, Mr Finney undermines the concept's validity altogether. It is surely no coincidence that this portion of the dataset most strongly indicates the positive influence camera deployment has had on collision rates leaving the report open to accusations of post-hoc manipulation.

DF: Figure 8.1 in my report shows KSI collision data direct from the TVSRP database and we agree this data is correct. The fact that the entire reduction in KSI collisions at these sites had already occurred a full year before mobile speed cameras started operating means that the speed cameras simply could not have caused any part of that reduction.

CONCLUSION

Truly independent analysis requires the forming of a hypothesis, devising the methodology and then analysing the results. By adapting a theory to match the results Mr Finney has shown a lack of independent approach and demonstrates clear bias against speed enforcement.

DF: My FTP method to measure and fully exclude RTM was on the RTM page of my website in 2008 and TVSRP sent me their database in 2009. Maybe RO didn't read my RTM page, or perhaps he did but doesn't remember.

Even if the artificially manufactured allowance for RTM is valid, the basic misunderstanding of how site selection was carried out in Thames Valley between 2000 and 2007 renders the results an irrelevant waste of time.

The failings outlined earlier undermine the conclusion:

"Table 8.2 shows the changes that occurred due to RTM at the mobile speed camera sites in Thames Valley prior to speed cameras being deployed."

It doesn't show any such thing. It shows a difference between two arbitrary time periods based on a very small sample with an inherent reporting issue. The subsequent conclusion that speed cameras increased crash rates is totally unreliable for the same reasons.

Independent analysis of published statistics is to be welcomed and Mr Finney's perseverance is to be commended. It is unfortunate that in his pursuit of *the truth* he has followed a path that leads to an analytical cul-de-sac with the only option to turn back and start again. Analysis the effectiveness of a group of interventions is not a complicated matter and Mr Finney's research seems to only complicate and twist the results to suit his own agenda.

There are many published reports into the effectiveness of speed camera enforcement and Mr Finney would do well to study these, their methodologies and analytical techniques before re-attempting to analyse the Thames Valley dataset.

DF: Results in my report are consistent with official reports. Regarding KSI collisions at speed camera sites:

1) RTM was indicated to be larger than any other factor (Professor Allsop's 2010 report)
2) RTM was estimated to be larger than all other factors combined (four year evaluation)
3) RTM was measured and found to be larger than the entire reduction that had occurred (my report)

There is surely more agreement than dispute when comparing my report to official reports. RTM is acknowledged to be the largest factor at speed camera sites and, the more accurate the analysis, the greater the level of RTM that is found. When RTM was finally measured in my report, all official estimates were found to have been under-estimates.

There have been developments since my 2012 report. In 2013, Professor Allsop published his second report for the RAC Foundation and RO commented on the RAC report: "*The level of statistical analysis … really lends weight to the findings*" and also: "*I have yet to see a report of similar quality from any source…*". At the time he wrote that, RO was not aware that Professor Allsop had actually used the FTP method that I had developed and used in my report[8].

Since then, my report has been submitted and accepted into the Road safety Knowledge Centre and both the DfT and the RAC Foundation have agreed that my FTP method does assess speed camera effectiveness accurately[9].

Of course, all of the issues discussed above could be dealt with simply by running speed cameras within scientific trials.

https://speedcamerareport.co.uk/scientific-trials/

8: http://www.roadsafetygb.org.uk/news/2870.html
9: http://www.roadsafetygb.org.uk/news/3234.html

Richard Owen

Operations Director
Road Safety Analysis Limited
Formerly Operations Manager of the Thames Valley Safer Roads Partnership (2003-2011)