An analysis of the effects of speed limit enforcement cameras with differentiation by road type and catchment area

Stephane Hess
Centre for Transport Studies
Imperial College London
Exhibition Road
SW7 2AZ London
United Kingdom
Tel: +44(0)20 7594-6105
Fax: +44(0)20 7594-6102
stephane.hess@imperial.ac.uk

Submission Date: August 1st, 2003

Word Count: 6535 + 2 Tables = 7035

ABSTRACT

In this paper, we present a detailed statistical analysis of the effects of speed limit enforcement cameras on injury accident numbers. The approach used is constructed in such a way that it is possible to differentiate not only between the effects of the camera and the effects of trend and seasonality, but to actually produce estimates that are independent of any other overall time-dependent effects. Crucially, the estimates produced are also net of the effects of regression to the mean. In order to allow for the simultaneous treatment of the different levels of severity, weights are used that reflect the frequency of the different types of accidents. This approach is then used on a data set for all injury accidents in Cambridgeshire between 1990 and 2002, which also contains data from 49 camera-sites. In order to quantify the range of effectiveness of the cameras, estimates of the changes in accident numbers are produced for different distances from the camera-site. The analysis shows that overall, in the immediate vicinity of the camera-site, the installation of a camera can be expected to lead to decreases in weighted injury accident numbers by an astounding 45.74%. Lower, but still significant decreases are observed in the wider surrounding area. Finally, to gain further insight into the differences in performance on different road types, the sites are grouped together according to road category. This analysis shows that the biggest reduction in accident numbers can be obtained on roads with higher incidence of speeding offences.

INTRODUCTION

Over the past decade, the installation of speed limit enforcement cameras (SLECs) as a deterrent against speeding has become one of the most widely used tools in the fight not only against speeding, but also against the ever-increasing number of severe accidents caused by this offense. Indeed, the link between speed and accidents is well established (e.g. *1*), as is the link between increases in speed and increases in the probability of serious and fatal injury (*2*). As it is also known that reductions in speed can lead to important reductions in accidents and the average level of severity (*3*, *4* and *5*), SLECs do indeed have the potential to lead to important reductions in road casualty numbers, through reducing the number and severity of speeding offenses.

  Even though there have been a large number of studies showing the effects of SLECs in terms of reducing speeding offenses (e.g. *6*, *7*, *8, 9*), there has only been a relatively low number of studies looking into the effects that SLECs have on accident rates and casualty numbers. Although most of these research projects report important drops in accident and casualty numbers (e.g. 9, *10, 11*, *12*, *13*), the few really informative such reports are by far outnumbered by reports of a more descriptive nature. In fact, a significant portion of such research does simply come in the form of annual performance reports published by the local highway authorities. Apart from a few notable exceptions, the analyses presented in such reports do not generally rely on the use of statistically sound methods and often ignore the influence that other time-dependent factors, like seasonality and trend, can have on accident figures. Similarly, the effects of regression to the mean are often simply ignored. There is also a distinct lack of precise estimates of the wider area (halo) effects of the cameras, although it should be acknowledged that such effects are even more site-specific than the estimated reduction in the immediate surrounding area of the cameras.

  In this paper, we further develop the approach used in one of the more methodological research projects looking into the effects on injury accident numbers (*9*). This approach gives an estimate of the reduction in the number of injury accidents following the installation of a SLEC, net of the effects of trend, seasonality and regression to the mean. Using data from 42 camera sites, (*9*) estimated the average effect of the installation of a SLEC to be a drop in accident numbers by 31.26%. Due to the limitations of the approach, the analysis could however not differentiate between the different levels of severity (slight, serious and fatal injuries), nor did (*9*) give any precise estimate of the range of effects or catchment area of a SLEC. The analysis presented in this paper uses 7 additional camera sites compared to the original study (raising the number to 49); very few previous studies have used data from such a high number of sites in a proper statistical analysis of the effects of SLECs on accident levels.

  The aim of this paper is to extend on the research conducted by (*9*) through using this larger data-set, and through adapting the method such that different weights are assigned to different levels of severity. Furthermore, the approach is refined in such a way that the resulting estimates of the camera effects are independent of any overall time-dependent factors, not just of trend and seasonality. This approach is then used to give an estimate of the effects of SLECs for various ranges around the cameras, such as to quantify the performance of the cameras in terms of their catchment area.

  The remainder of this paper is organized as follows. In the next section, we describe the data set used, and conduct a preliminary analysis on the data. The third section introduces the methodology used in the analysis, and the fourth section presents the results produced in the analysis. Finally, the fifth section compares the results from the current analysis to those produced in comparable previous studies.

DATA DESCRIPTION

As was the case in the study conducted by (*9*), the data used in the present analysis is again taken from the database of injury accidents recorded in Cambridgeshire, one of the southeastern shire counties of the United Kingdom. The data-set contains all injury-accidents recorded in Cambridgeshire between 1990 and 2002 and thus includes two additional years' data when compared to the original study, as well as data from Peterborough. The total number of injury-accidents used is 44,376, compared to 31,042 in the work undertaken by (*9*). Again, only injury accidents were considered and no difference was made between accidents with different numbers of casualties; the severity of an accident is defined as the highest level of severity of any of the casualties resulting from the accident.

The number of camera sites used has also increased from 42 to 49. A further 17 camera sites were not included in the study as less than a full year of data was available for the time following the installation of the camera; this could have undermined the robustness of the results in terms of independence from the effects of regression to the mean. Under UK legislation, and following the rules of the "hypothecation" scheme (c.f. *9, 10, 11*), all SLECs are highly visible, signing is used within a radius of 1 kilometer of the camera site and the camera locations are widely publicized. The aim of SLEC programs is not to punish drivers for speeding, but to prevent them from doing so through use of the "fear-factor"; hence the most successful cameras are those that never catch any driver.

Preliminary data analysis

In this paper, we aim to produce estimates of the effects of the SLECs, classified by distance from the camera, also referred to as catchment area of the camera. Four different catchment areas were used for each camera, comprising accidents within a distance of 250, 500, 1000 and 2000 meters respectively from the camera. Data on traffic accidents and camera locations was geo-coded and then overlaid on road network data using geographic information systems (GIS). This was then used to allocate accidents to the different camera sites. Accidents were only associated with a given camera site if a sensible assumption could be made regarding a link between the camera and the accident site (i.e. even though an accident on a nearby parallel road may fall into the geographical catchment area of the camera, no link between the two can be assumed in the absence of any side roads connecting the two main roads).

Especially for the two wider catchment areas, 1000 and 2000 meters, there was some overlapping of the catchment areas for cameras positioned close to each other. There was thus a problem of how to allocate the accidents in the intersecting regions to the respective cameras. One possibility would have been to assign weights to the accidents, such as to assign a portion of each such accident to the different camera sites. This approach was rejected, mainly because of the problems of finding appropriate weights. In the case of a low level of overlapping, such that the number of accidents comprised in the intersecting region is insignificant compared to the overall size of the catchment area, it was decided to simply ignore the overlapping, and to assign the accidents to each of the different cameras concerned. A separate analysis showed that, overall, the risk of biased results was in this case very small and any bias would generally lead to underestimating of the effects of the cameras rather than overestimating (due to higher data levels). A different approach was however taken in the case where several cameras are positioned very close to each other (rather than simply having slightly overlapping catchment areas), as is for example the case when different cameras are used for different approaches on the same road (for example in the case where two one-directional cameras are used for the two

directions on a motorway). In this case, special care was taken to assign the accidents to the correct cameras. The difference between this case and that of some minor overlapping of the catchment area is simply that in this case, most of the geographical catchment area is shared between the two cameras, whereas in the earlier example, only a minor portion of the catchment area was shared.

Given the differences in the impacts that accidents of different levels of severity have on the safety record of a given accident site (as well as on society), it is clearly desirable to treat the three different levels of severity separately. There is however little point in repeating the widely used approach of separately calculating the effects for the different levels of severity. Indeed, serious injury and fatal accidents are so rare that for most sites, they are unlikely to occur in both the period before installation of the camera and the period after installation of the camera (at least for fatal accidents); this would thus regularly give a reduction to zero or increase from zero in the number of these accidents. In order for all accidents to be treated simultaneously while also reflecting the differences in their levels of severity, it is necessary to assign different weights to the different types of accidents. No weights were used in the approach conducted by (9) due to the use of specific time series models with which the use of weighted data led to inconclusive results. The problem was now to find weights that would reflect the differences in severity without unnecessarily distorting the time series nature of the accident data. A first attempt was made with weights reflecting the differences in the cost of accidents. The costs used are those defined as cost prevention benefits generally used in the United Kingdom; these account for costs to the Health service, as well as loss of income and cost of damage to property. The respective figures for slight-, serious- and fatal-injury accidents are £15,380, £154,080 and £1,323,880, leading to weights of 1, 10.02 and 86.08 respectively. These weights were however found to be impractical as the resulting distortion to the data-set was too significant (the time series nature of the data set was disrupted as the weights led to significant peaks in the data). Rather than using these weights, it was decided to use weights that reflect the likelihood of the different types of accidents. The numbers of slight-, serious- and fatal-injury accidents recorded between 1990 and 2002 are 35,582, 7,950 and 844 respectively. From these figures, the respective weights of 1, 5.58 and 41.46 were calculated. This approach thus produces weights according to the frequencies of the different types of accidents, as opposed to explicitly taking into account the differences in severity between the different types of accidents. Such approaches accounting for the differences in severity are for example given by (14) and by (15). A review of studies looking into the weighting of accidents is given by (16). In the present analysis, it was found that the use of weights reflecting the frequency of accidents was the most appropriate approach, as they were consistent with the desire of treating the different levels of severity simultaneously. The above given weights were thus used in the remainder of the analysis.

METHODOLOGY

The approach used to calculate an estimate of the effects of the installation of the cameras is broadly similar to that used by (9), but differs from it on several points which make it more accurate. In the following four sections, we describe the methods used to
- differentiate between the effects of the cameras and the effects of regression to the mean,
- remove trend and seasonality from the data-set,
- calculate the effects of the camera installation,
- and average the results over the different camera sites used in the analysis.

Effects of regression to the mean

Regression to the mean is the statistical phenomenon of a time series returning to its natural mean level. This notion is used in very diverse areas of statistical analysis, and can for example be used to (at least partially) explain the fluctuation observed over time in stock-market indices. In the present context, regression to the mean can be observed in the case where, after a stretch of months with very poor performance, the level of accidents observed in the following months returns to the expected levels for those months. In the case where a SLEC is installed following a stretch of very bad months, and where this is followed by a drop in accident numbers, it can thus be difficult to differentiate between the effects of regression to the mean and the actual effects of the installation of the SLEC. Even though the installation of a camera is not generally such a short-term decision (certain criteria must be met regarding the accident record over a three-year period), the effects of regression to the mean clearly need to be taken into account. Indeed, a major mistake made in some previous studies (of a less mathematical nature) was to simply compare the accident count for the three months before the installation of the camera to that for the three months following the installation of the camera. Given the possible existence of regression to the mean, this approach is clearly not appropriate.

A brief overview of existing techniques for detection of regression to the mean given by (*9*) shows that these are not appropriate in the current application. A new method is thus developed by (*9*), based on using long-term mean levels before and after the installation of the camera. The use of such long-term mean levels means that any peaks and slopes caused by regression to the mean are evened out. Indeed, the long-term mean level before installation of the camera can be regarded as being the stable level of accident numbers prior to the installation of the camera, while, for a sufficiently long posterior period, the mean level after installation can be treated as the stable level following the introduction of the camera. The use of these mean levels also lessens the risk of the results being biased by the good performance in the first few months following the installation of the camera. Long-term posterior levels are not generally available, but special care should be taken to at least only include sites with a posterior period of sufficient length to safely assume that any mean level can be treated as a stable state. It was decided to set this limit to twelve months, such that only cameras installed prior to January 2002 were used in the present analysis. A further requirement with this approach is that, in order for the difference between the two mean levels to show only the effects of the cameras, the mean levels need to be free of any other time-dependent effects, such as seasonal variation and effects of trend. A special approach would need to be developed to remove the effects of these factors from the data-set.

Removal of time-dependent components

Besides the use of weights for different levels of severity, the main difference between the approach used in the present analysis and that used by (*9*) is the way in which time-dependent effects are removed from the data. The original study focused solely on trend and seasonality and used a time series analysis to derive coefficients for the seasonal and trend components. While the use of these coefficients can be used to smooth the data, it cannot be used to completely remove all time-dependent factors from the data. Indeed, the fact that the seasonal factors used are the same for all years combined with the use of a linear trend component means that there will be a significant amount of random variation remaining in the data after smoothing. The main advantage of using such coefficients generated by a time series analysis is that they can be used

to forecast accident levels for later time periods, this however was not necessary in the present study.

The approach used in this paper aims to remove not only trend and seasonality from the data-set, but to account for all variation (including random) that leads to a deviation of a certain month's accident count from the mean monthly accident count over the period of observation. This approach does even account for the different lengths of months in the calendar. Furthermore, the three different levels of severity are treated separately, reflecting differences in the seasonal patterns for different levels of severity. This was not possible in the approach used by (9), as no time series model could be fitted to the data sets for fatal and serious injury accidents, due to the low monthly counts in these two categories. The new method does thus have several significant advantages over the approach used by (9).

Let X_{si} be the accident count for month i and severity level s. The multiplicative coefficients θ_{si} to be used to remove the time-dependent components can now be calculated very easily by

$$\theta_{si} = \frac{\overline{X}_s}{X_{si}},$$

where $\overline{X}_s$ is the mean monthly accident count for severity level s. There are no problems with using this method for the different levels of severity, as it does not attempt to explicitly model the behavior of the time series and it does thus not require a certain minimum level of data for each month (which was necessary with the other approach in order to discard random fluctuation). For months with no fatalities, the above formula cannot be used, but the multiplier can in fact simply be set to any value, as any multiplication will leave the count unchanged at zero.

All accidents, including those associated with camera sites, are included in the calculation of coefficients; the low number of camera sites (compared to the overall number of sites) means that the inclusion of these sites in the calculation does not lead to a significant bias in the coefficients. Given the fact that little weight is given to individual accident sites in this calculation, the method does not account for any time-dependent factors that are specific to single sites, but considers only the overall variation. This is crucial for the later identification of the actual effects of the cameras.

For a given accident site, with accident count (of severity level s) in month i given by Y_{si}, the time-dependent components can be removed by calculating

$$Z_{si} = \theta_{si} Y_{si}$$

This process is then repeated for all months. If the resulting sequence of data points shows a significant time-dependent trend (independent of random fluctuation), then the time-dependent variation for the given site is different from the overall variation over sites.

Calculation of effects

After removing the time-dependent components from the data, the monthly accident-counts for the different levels of severity are combined, after multiplication by the respective weights. The next step is to calculate the mean monthly accidents levels for the periods before and after installation of the camera. A problem arises because of the different installation dates of the cameras. Indeed, cameras are installed at varying days of the month, such that it is difficult to find a common start time for the after-installation period. While (9) use the first month following the installation of the camera as the start-date, in this paper, we start the after-installation period with the month comprising the installation of the camera. While this may lead to some

underestimating of the effects of the camera (assuming a decrease in accidents) due to the inclusion of some before-time, this underestimating is less likely to lead to wrong recommendations than is any overestimating of the effects. Some effect preceding the actual installation of the camera can also be expected, given the use of awareness campaigns prior to the installation. Furthermore, the fact that 38 of the 49 cameras used were installed on or before the 15[th] day of the month supports this decision.

The prior and posterior monthly mean levels can now be used to estimate the effect of the installation of the camera. Indeed, if, after removal of the overall time-dependent effects, there is a significant difference between the long-term mean levels before and after installation of the cameras, there has been a significant change in (weighted) accident numbers, net of the effects of regression to the mean as well as overall time-dependent factors. Assuming that no other site-specific event (with possible effects on accident rates) took place on the given camera site during the period of the analysis, any changes in accident numbers can be interpreted as being caused by the installation of the SLEC. This can be seen as a significant assumption given for example the possibility of significant differences in traffic levels (compared to overall changes), but it is necessary to make this assumption in order to use the method. As no significant changes in traffic level over time have been reported for the sites under investigation (other than those observed on other sites, i.e. the regional trend), and as no other reasons for major changes in accident levels (excluding SLEC-related signalling and awareness campaigns) seem to exist at these sites (never mind to coincide exactly in time with the installation of the cameras), the differences in the long-term mean levels, net of time-dependent effects, can indeed be seen as robust estimates of the effects of the installation of the cameras.

Averaging over sites

An issue treated in great detail by (*9*) is that of what approach should be used when averaging the effects of the camera installations over different sites (as opposed to using single-site effects). An important point to consider in this calculation is that the different sites have different characteristics, not only in terms of installation dates, but also in terms of accident levels. The differences in installation months have been accounted for by the removal of all overall time-dependent components and the calculation of monthly accident levels before and after installation of the cameras. However, the differences in the level of accident counts over sites have yet to be taken into account. These differences can be very significant; as an example, when using the 250 meter radius, the weighted monthly accident count (after removal of time-dependent components) at the *largest* site (site with highest accident count) is 241 times more important than the corresponding count at the *smallest* site. While this is an extreme case, and the rest of the differences are far less important, such differences clearly need to be taken into account.

The problem with sites with very low accident counts is that even small unit changes lead to important percentage changes, while significant unit changes are needed at *larger* sites to even lead to a modest percentage change. This means that simply calculating the ratio of changes for individual sites and averaging these ratios over sites will not give a robust estimate of the overall ratio. (*9*) and (*17*) experiment extensively with using different weights for different sites in the calculation of the overall effects of cameras, but show that the best approximation of the effects can in fact be made by calculating the total monthly accident levels over all sites (after applying severity weights and removing time-dependent components) and to calculate the rate of change over all sites from these two totals. This is essentially equivalent to a ratio of totals (which is the same as a ratio of means) while the earlier approach uses a mean of ratios, given by the average

of changes over sites (*9*). The ratio of totals approach is clearly more appropriate than the simple averaging of changes over sites given by the mean of ratios. As it is also more applicable than any complicated weighting schemes, it was decided to use the ratio of totals approach in the calculation of the overall effects in the present analysis. It should be noted that this approach does in fact indirectly assign different weights to different sites, as sites with higher accident counts carry more weight in the calculation of the ratio. One disadvantage of this method is that there is no practical way of producing a confidence interval for the effects of the camera installation. While the calculation of such confidence limits is possible when using the mean of ratios approach, the resulting values are of little practical use as the calculation is hugely influenced by the presence of a few outliers.

RESULTS

The approach described in the previous section was used to calculate the monthly accident levels for the individual sites, before and after installation of the SLECs, net of the effects of any overall time-dependent factors for the four different ranges (250, 500, 1000 and 2000 meters). These values were then used to calculate overall changes in monthly accident levels at camera sites, by summing the prior and posterior mean accident levels over all sites, and taking the ratio of these totals. The resulting effects are net of the effects of regression to mean, trend, seasonality and other overall time-dependent factors, and are representative of the overall effects of cameras, by indirect weighting according to the importance of the different sites. The experiment was conducted for the four different ranges and the results are summarized in table 1.

The results show that in the 250 meter range, the average effect of the installation of a SLEC is a drop in (weighted) injury-accident numbers by an astounding 45.74%, corresponding figures for the 500, 1000 and 2000 meter ranges are reductions by 41.30%, 31.62% and 20.86% respectively, where the estimate for the reduction in the 1000 meter range is very close to that given by (*9*). Although, as expected, the effects decrease with distance to the cameras, in an area spreading 2 kilometers to all sides of the camera, the reduction is still almost a half of the reduction observed within the immediate vicinity (250 meter radius), thus indirectly refuting the general claims that the (highly visible) cameras make people slow down abruptly for the camera before accelerating again. Such a speed-changing behavior would almost inevitably lead to a higher accident count in the surrounding areas. Also, given that the signing of SLECs extends only to 1000 meters to all sides of the cameras, the fact that there is still a significant effect outside this area suggests that SLECs do indeed succeed in making people drive more carefully in the wider network, as opposed to simply in zones where SLECs are known to be in operation.

As a further extension of the results, the camera sites were split into different groups. First, the sites were divided by road type, and the overall effects were calculated for A-roads (major roads, 38 cameras) and for non-A-roads (minor roads, 11 cameras). Secondly, the sites were divided again, this time by location, such that effects were calculated for cameras situated in urban areas (34 cameras) and cameras situated outside urban areas (i.e. on trunk roads, giving 15 cameras). Finally, to look into the effects in multi-camera areas, the sites were split according to whether they are sited in multi-camera areas or whether they are stand-alone cameras. Cameras were selected into the multi-camera group if they were positioned within 1000 meters of another camera, thus not simply if there was any overlapping of their respective 1000 meter ranges. This way, 19 out of the 49 cameras were placed into the multi-camera group, with the remaining 30 being treated as stand-alone cameras. As mentioned previously, special care was in this case taken to attribute accidents to the right camera sites. The results from these three

separate analyses are summarized in table 2, showing the percentage variations for the different ranges.

The results show that findings from the overall analysis regarding the decrease of effects with increasing distance to the camera site are repeated for A-roads, trunk roads and stand-alone cameras groups; but not for the other three groups, where the best performance is always obtained in the 500 meter range. In comparing the performance in the different pairs of groups, we note that cameras on A-roads perform better than cameras on non-A-roads, except in the 2000 meter range. Similarly, cameras on trunk roads outperform cameras on urban roads except in the 2000 meter range. For both pairs, the differences in the 2000 meter range are however far less significant than the differences in the other ranges. Furthermore, it should be noted that the results in the 2000 meter range analyses are in general less robust than those in the other ranges, as there is an increased risk of other factors (outside the cameras) playing a role. The results in these two pairs of groups are broadly consistent with earlier results (e.g. *6*, *7*, *8*, *9*) noting that the use of SLECs can lead to reductions in excessive speeding. As such excessive speeding is well known to be directly linked to road accidents, and as the scope for reducing the amount of excessive speeding is far more significant on A-Roads and on trunk roads than on non-A-Roads and on urban roads, it is no surprise that the reduction in accidents is similarly more important on the former types of roads.

Finally, a comparison between the performance of multi-area cameras and stand-alone cameras shows that the latter outperform the former except for the 2000 meter range. This does however not immediately indicate that the use of multiple cameras will lead to a less significant reduction than the use of stand-alone cameras. Indeed, due to the overlapping of catchment areas, the individual catchment areas of the single cameras decrease in size, thus also decreasing the scope for reductions in accidents. However, it can equally well be noted that by reducing the size of the catchment area in this way, the numbers of prior and posterior accidents should be expected to decrease by equal proportions, leaving the ratio between the two numbers unaffected. The results from this part of the analysis are thus not conclusive, although it can be observed that both types of camera do indeed lead to very significant reductions in the number of injury accidents.

COMPARISON OF RESULTS WITH OTHER STUDIES

It is of interest to briefly compare the results produced in the present analysis with results reported by other authors. Studies conducted in the United Kingdom have shown that the installation of a SLEC can result in significant reductions in accident numbers; a study by (*11*) shows that the installation of a SLEC leads to a decrease in casualties (rather than accidents) by an average 28%, while an analysis by (*10)* estimates the reduction in KSI (killed and seriously injured) in the immediate vicinity and wider surrounding area of camera sites to be 47% and 18% respectively. In broad terms, these results are confirmed in the present analysis, which in some cases shows even higher reductions than previous estimates. Finally, in a comparison with the results produced in the original analysis conducted by (*9*), the results from the present analysis show a very similar reduction in the 1000 meter range, even though a higher number of sites were used, in addition to the use of severity weighting and a more refined approach for removing time-dependent components. The results from the present analysis thus validate those produced by (*9*).

Given the very significant reductions in accident numbers in the immediate surroundings of the camera (e.g. 250 and 500 meter ranges), this analysis also confirms the results by (*18*) who

reject the hypothesis that the installation of SLECs can lead to increases in accidents in the immediate vicinity of the cameras by leading to abrupt braking manoeuvres.

Results from around the world show similarly positive effects by the cameras; as an example, (*7*) report decreases in collisions by 36% and reductions in fatalities by 70%, while (*13)* report reductions in casualty crashes (in low alcohol hours) by up to 30%.

Overall, this comparison thus shows that the results from the present analysis confirm and extend on earlier results that report significant effects of SLECs; while also adding some additional statistical evidence regarding the effectiveness of SLECs in terms of accident reductions in the wider surrounding area. Finally, it should be noted that a separate analysis has shown that the use of severity weighting does not in fact lead to significant changes in the estimated effects it does however make the analysis more consistent with the intuitive perception of the different types of accidents.

SUMMARY AND CONCLUSIONS

During the research carried out for this paper, it became very obvious again that, in order to ensure the robustness of results produced by an analysis of the effects of SLECs, these effects need to be separated from other time-dependent effects, like trend and seasonality. This separation is performed by calculating weight factors for each month, reflecting the relative overall injury accident frequency in Cambridgeshire in that month compared to the mean frequency over all months. These weights can be used to rescale the data for individual sites, thus removing any overall time-dependent effects from the data for the given sites.

Besides being able to identify overall time-dependent effects, any comparison should also be able to differentiate between the effects of the camera and the effects of regression to the mean. In order to ensure the greatest possible independence from the effects of regression to the mean, our analysis compares long-term prior and posterior *stable* mean levels, thus reducing the effects of peaks and slopes found in the data.

In order to account for the differences in terms both of social impacts and accident frequency, higher weights are assigned to serious-injury and fatal accidents.

Finally, the calculation of the average effect over sites is performed by calculating the total monthly number of accidents before and after installation over all sites, net of the effects of any time-dependent components, and by calculating the ratio of these two totals. This approach indirectly gives more weight to sites with higher levels of data (which are less subject to random variation).

The analysis shows that, in the immediate vicinity of an SLEC, a very significant reduction in weighted accident numbers by 45.74% is achieved. Although, as expected, the effects decrease with distance from the camera, the effect inside a 2000 meter radius from the camera is still important, showing a reduction by 20.86%. Finally, a division of the camera sites into different groups shows that the most important reductions are achieved on road types where speeding offences (and related accidents) are more common.

Overall, the present analysis has thus reiterated previous results showing that the installation of SLECs as a deterrent (rather than a trap) can lead to very significant reductions in injury accidents, with the obvious benefits this has for society. This suggests that highly visible SLECs, in conjunction with adequate signing and awareness campaigns, can play an important part in the reduction of road-accident casualty numbers.

ACKNOWLEDGMENTS

REFERENCES

1. Bowie, N. and Walz, M., "Data Analysis of the Speed-Related Crash Issue", Auto and Traffic Safety, Vol. 1, No. 2, pp. 31-38, 1994.

2. Hobbs C.A. and Mills P.J., "Injury probability for car occupants in frontal and side impacts", TRL Report 1124, TRL, Crowthorne, 1984.

3. Taylor M., Lynam D., Baruya A., "The effect of drivers' speed on the frequency of road accidents", TRL Report 421, TRL, Crowthorne, 2000.

4. Finch D.J., Kompfner P., Lockwood C.R. and Maycock G., "Speed, speed limits and accidents", TRL Report 58, TRL, Crowthorne, 1994.

5. Transportation Research Board, "Managing Speed", Special report 254, Transport Research Board, National Research Council, National Academy Press, Washington, 1998.

6. Retting, R., "Evaluation of Speed Camera Enforcement in the District of Columbia", paper presented at the 82nd Annual Meeting of the Transportation Research Board, Washington DC, 2003.

7. Anderson, R. and Edgar, A.W., "Mobile Speed Cameras in the ACT – Slashing Speeds and Cutting Crashes", paper presented at the Road Safety Research, Policing and Education Conference, Melbourne, 2001.

8. City of San Jose, "NASCOP: An Evaluation of the Photo-Radar Speed Enforcement Program", San Jose, CA, City of San Jose, 2001.

9. Hess, S. and Polak, J.W., "An Analysis of the Effects of Speed Limit Enforcement Cameras on Accident Rates", paper presented at the 82nd Annual Meeting of the Transportation Research Board, Washington DC, 2003.

10. PA Consulting, "Cost Recovery System for Traffic Safety Cameras - First Year Report", DTLR Road Safety Division, 2001.

11. Hooke, A., Knox, J. and Portas, D., "Cost benefit analysis of traffic light and speed cameras", Police Research Series Paper 20, Home Office Police Research Group, London, 1996.

12. Bourne, M. and Cooke, R., "Victoria's Speed Camera Program", in Clarke, R.V. (ed.) Crime Prevention Studies Vol.1. Criminal Justice Press, New York, 1993.

13. Cameron, M., Cavallo, A., and Gilbert, A., "Evaluation of the Speed Camera Program in Victoria 1990-1991", Report 42, Monash University Accident Research Centre, Victoria, Australia, 1992.

14. Deacon, J.A., Zegger C.V. and Deen R.C., "Identification of hazardous Rural Highway Locations", *Transportation Research Record* 543, TRB, National Research Council, Washington D.C., pp. 16-33, 1975.

15. Gharaybeh, F.A., "Identification of Accident-Prone Locations in Greater Amman", *Transportation Research Record* 1318, TRB, National Research Council, Washington D.C., pp. 70-74, 1991.

16. Quddus, M.A., "A technique to rank hazardous intersections", paper presented at the 6[th] international conference on Application of Advanced Technologies in Transportation (AATT) Engineering, Singapore, 2000.

17. Hess, S., "A statistical analysis of the effects of safety cameras on traffic accident rates in Cambridgeshire", MPhil Thesis, Statistical Laboratory, Cambridge University, 2002.

18. London Accident Analysis Unit, "West London Speed Camera Demonstration Project, July 1997", Highways Agency, London, 1997.

TABLE OF CONTENTS

INDEX OF TABLES

TABLE 1. Average reductions (by range) in accident numbers over all sites following introduction of a speed camera

Range (meters)	Total Prior Monthly Count	Total Posterior Monthly Count	Percentage change
250	23.372	12.682	-45.74%
500	50.267	29.507	-41.30%
1000	94.990	64.951	-31.62%
2000	236.319	187.021	-20.86%

TABLE 2. Reductions in accident numbers using grouping by positioning of camera

	250 draws	500 draws	1000 draws	2000 draws
A-Roads	-55.33%	-41.30%	-31.62%	-20.86%
Non-A-Roads	-10.04%	-44.33%	-35.70%	-19.78%
Urban Roads	-28.60%	-31.83%	-19.71%	-23.65%
Trunk Roads	-72.25%	-33.95%	-25.45%	-21.02%
Multi-area	-30.99%	-57.92%	-50.15%	-20.02%
Stand-alone	-56.00%	-35.18%	-23.18%	-21.54%