The Placebo Effect and Red-Light Cameras

Dale Gedcke

June 8, 2008

When I began my exhaustive research into the available studies on the effectiveness of Red-Light Cameras (RLCs), I very quickly discovered that there are two dominant limitations in testing their effectiveness for reducing accidents:

1) The number of accidents reported before and after RLC installation is so small that the statistical uncertainty in those numbers obscures any trend.
2) There are so many confounding variables, that it is extremely difficult to extract the true effect of the RLCs without removing the effect of those confounding variables.

This article briefly explains problems 1) and 2), and shows how an analog of the placebo test used in the pharmaceutical industry is applied to RLC data to sequester the effects of the confounding variables. Finally, the "placebo effect" explains why those with a huge stake in the implementation of Red-LIght Cameras always seem to report favorable results that are at odds with the plethora of statistically sound studies.

The Random Uncertainty in Accident Numbers

If you wished to accurately predict the winner in the next Presidential Election, you could ask 20 people you know, how they plan to vote. Would you have any confidence based on how those 20 votes fell? You shouldn't, because the sample size (number of people asked) is extremely small, and the people sampled may not be representative of the nation's voting population. For example, if most of the friends you polled belonged to one party, the candidate from that party would likely be predicted to be the winner. Professional pollsters understand that the sample size has to be very large to reduce the significance of the random error, and the population polled must be representative of the nation's voters. When the two candidates are evenly matched, statistical theory[1] shows that one must poll at least 2,500 people to be 95% certain that the margin of error is less than 2 percentage points in the prediction of 50% for candidate A and 50% for candidate B. In other words, that poll leads to a 95% certainty that the results are between 48% A vs. 52% B and 52% A vs. 48% B.

Accidents are rare, random events. For example, angle collisions at a busy intersection controlled by traffic lights are typically less than 15 accidents for 1 million vehicles travelling through that intersection[2, 3]. Consequently, there is a random variation in the number of accidents recorded at a particular intersection. The minimum variation, derived from the Poisson Probability Distribution, can be shown to be[1, 4]

$$\sigma = \sqrt{N} \qquad (1)$$

where N is the number of recorded accidents. The simple explanation for the standard deviation, σ, is this: If the measurement is repeated a very large number of times, and the average of all the measured values is calculated as the mean, then 68% of the measured values will fall within $\pm\sigma$ of the mean[1, 4]. In other words, σ is an estimate of the random dispersion of the measured values about the mean.

A more useful concept for comparing accident rates is to compute the percent uncertainty in the number of accidents. This is given by[4]

$$\sigma\% = \frac{\sigma \times 100\%}{N} = \frac{100\% \times \sqrt{N}}{N} = \frac{100\%}{\sqrt{N}} \qquad (2)$$

Table 1 shows the percent uncertainty in the predicted mean number of accidents as a function of the measured number of accidents, N. Clearly, one needs at least 1,000 recorded accidents to reduce the uncertainty to circa 3%.

Table 1. How the Percent Standard Deviation Varies with N

N	1	10	100	1,000	10,000	100,000	1,000,000
σ%	100	31.6	10.0	3.16	1.00	0.316	0.100

Rigorous statistical tests typically require that a new measured value differ from the mean by more than $\pm 2\sigma$ to be considered characteristic of a truly different mean. This is the so-called 95% confidence limit[1]. That is, if the measured values are actually from the same distribution, 95% of the measured values will lie within $\pm 2\sigma$ of the mean.

In a simple before and after test of 30 Red-Light-Camera intersections, one may find the total number of accidents for the year before camera installation is $N_B = 842$. If the total number of accidents in the year after RLC installation is $N_A = 758$, that difference is a 10% decrease. Is that difference statistically significant, or does it simply represent a random fluctuation in the number of accidents?

When computing the comparison ratio

$$R = \frac{N_A}{N_B} \qquad (3)$$

The percent standard deviation in the ratio, $\sigma\%_R$, is calculated from[4]

$$\sigma\%_R = \sqrt{(\sigma\%_A)^2 + (\sigma\%_B)^2} \qquad (4)$$

where $\sigma\%_A$ is the percent standard deviation in N_A, and $\sigma\%_B$ is the percent standard deviation in N_B. For the example, $\sigma\%_R$ is calculated as

$$\sigma\%_R = \sqrt{\left(\frac{100}{\sqrt{N_A}}\right)^2 + \left(\frac{100}{\sqrt{N_B}}\right)^2}$$

$$= 5.0\% \qquad (5)$$

Consequently, the $\pm 2\sigma$ value for the symmetric 95% confidence limits would be ±10%. The number of accidents after RLC installation was 10% lower than the number of accidents before installation, and that change is exactly at the 95% confidence limit. Therefore the change is barely statistically significant.

Whenever one looks for a change in the accident rate from before to after RLC installation, a statistical test for significance, like the one above, must be applied before any valid conclusion can be drawn. In the next section it will be discovered that other factors can increase the uncertainty beyond what has been considered with the basic numbers in the above example.

Confounding Variables Require a Placebo Test

Main issue number 2 involves the confounding variables. When examining the accident rates before and after RLC installation, one finds that the cameras are not the only factor influencing the accident rate. Changes in the following confounding variables can also affect the before and after accident rates.

Table 2. Confounding Variables that Affect the Accident Rate

- traffic flow rate
- traffic patterns
- road improvements
- construction projects
- weather
- intersection characteristics
- speed
- time of day
- day of the week
- month of the year
- vehicle safety features
- driver characteristics
- truck traffic
- yellow duration
- all-red duration

Any one of these variables could experience a large change in one direction, causing a significant corresponding change in the accident rate. That unidirectional change could mislead one to conclude that the RLCs were responsible for the change. How does one eliminate the confusion? The most common method is to identify "comparison" or "control" intersections, and divide the change in accident rate at the RLC intersections by the before/after change in the accident rates at the comparison intersections[3]. The comparison intersections need to be similar in all characteristics to the RLC intersections. Preferably the comparison intersections are no closer than one mile to the RLC intersections, so that they are not influenced by the RLC intersections. Dividing the before/after accident ratio for the RLC intersections by the before/after accident ratio for the comparison intersections should eliminate any systematic change that is common to both types of intersections. Generally, one should choose at least twice as many comparison intersections as RLC intersections. This ensures that the statistical uncertainty from the comparison intersections degrades the overall statistical uncertainty in the answer by no more than a factor of 1.22.

This procedure of using comparison intersections to cancel systematic changes in confounding variables is similar to the placebo test used when evaluating new drugs. The clinical test population is divided in half. One half gets treated with the new drug, and the other half gets the placebo, which is often a sugar pill. This is normally a double-blind test to avoid the possibility of humans biasing the test. Neither the patients, the pill dispensers, nor the Doctors evaluating the patient's response know which pill was administered to whom. Only the computer knows when it is analyzing the results from both groups of patients. The placebo test rejects symptoms that were common to both the new drug and the placebo, making it possible to narrow in on the results that can be ascribed to the new drug.

In the simplest before and after test for the effects of RLCs on accident rates, the comparison intersections must be used, and the results tested for statistical significance. Without at least this minimum process, the results of the RLC analysis will be unreliable and inconclusive.

More Uncertainty from the Confounding Factors

Some of the confounding factors can be monitored and measured. For those that are measured, the statistician can run a linear regression calculation to estimate the contribution of that variable to the before/after accident rates and remove their contribution. Often, a more complicated regression model is applied[2, 3, 5].

For the confounding variables that have not been monitored, there is another lingering problem. Not only can a unidirectional change in one of these variables augment or detract from the perceived effect of the RLCs, but random, bidirectional changes in the confounding variables can increase the uncertainty of the before/after measurement. The standard deviation in equation (1) presumes a constant factor is driving the accidents. If that driving factor is randomly varying up and down, the variations in the measured number of accidents will increase beyond what is described in equation (1). Statisticians account for this effect by examining the variations over a sequence of shorter periods of time, and by using that information to derive an empirical estimate of the actual standard deviation. That empirical standard deviation is used in the statistical significance test. Needless to say, the accuracy of the before/after comparison is degraded by the unmonitored confounding variables. Thus the example from equation (5) is usually optimistic.

Unless the organization reporting the effect of RLCs on accident rates has accounted for the variance introduced by the confounding variables, the conclusions about statistical significance are seriously suspect.

Psychology and the Placebo Effect

Why do enthusiastic proponents of Red-Light Cameras, who have a stake in their application, describe the benefits in such certain and glowing terms, in contradiction to what the unbiased statistical studies show? The overly optimistic perspective may be a result of the placebo effect.

The placebo effect is well known in medicine. One gives a patient a placebo pill and tells them it is a powerful new drug that will cure their symptoms. In a large fraction of the patients, the disease symptoms disappear, and the patient reports being cured. It is not clear what physical mechanism causes this cure, but psychology is certainly involved. In the other direction, one can give a patient a placebo and tell them about all the negative side effects they may experience. Some of the patients will report having those side effects. The existence of hypochondriacs is also well known. These are people who hear of symptoms of a disease, and convince themselves they have those symptoms. Clearly, motivation can bias conclusions.

How does the placebo effect relate to Red-Light Cameras? The camera companies target the most vulnerable customer, the Police Chief. They point out that the cameras will increase the number of citations for red-light running at least 10-fold, without adding a single patrol car. Furthermore they offer to absorb the cost of installation in exchange for collecting approximately 70% of each $50 ticket that is issued. They show the Police Department and the municipality that 30% of the ticket revenue will be a substantial income for the city. The camera company appeals to the fondest wishes of the Police Chief and the City's government. That sets up the scenario for a biased perception of the benefits of the Red-Light Cameras.

It should come as no surprise that the salesmen for the camera companies pitch all the positive benefits and suppress the negatives. That is the normal function of a salesman or a marketeer. The sales presentation of the camera company concentrates on the number of tickets that can be automatically cranked out, the revenue generated, and their selected best numbers on accident reductions. Rarely will you hear reference to an unbiased statistical study on accident results. Comparisons to non-RLC intersections never enter the conversation, and statistical significance tests are avoided like profane language. There is no doubt that the camera companies have superb technology for monitoring intersections, and for deriving a wealth of information from what happens at that intersection. And they are right about the number of red-light running tickets that will be automatically issued. Compared to accidents at an intersection, the number of tickets written with cameras is orders of magnitude larger[2, 3]. So, the percent statistical uncertainty on the citation rate is very small. But, because the accident rate is so low, it is easy to paint the safety argument an optimistic color, because the true situation is obscured. Given the opportunity to sell the eager customer on a desirable outcome, the camera company will do just that.

But why do their customers perpetuate a biased and optimistic perception? There are two reasons: a) lack of expertise in statistical analysis, and b) the need to sell success.

Statistical Expertise: Even among scientists and engineers, there are few people who are fluent in statistical analysis of random variables. Expecting to find comfortable familiarity with statistical uncertainty among non-scientists is unrealistic. Most of the customers involved in deciding to purchase Red-Light Camera systems find statistical analysis to be a foreign language. They are expert in other disciplines that are more frequently vital to their tasks. Consequently, the plethora of sound statistical studies that are available on the topic get ignored.

Selling Success: Everyone has watched politicians, salesmen, and even business managers sell the positive aspects of their program while discounting any negatives. Once one is heavily committed to making a program successful in the long term, he will do what it takes to convince others that the program will continue to deliver success. In cases where the data to substantiate success is complex and offers opportunities to extract the positive and ignore the negative, the tendency will be to accentuate the positive and eliminate the negative. For Red-light Cameras, the data on accident statistics is obscure enough that it is possible to select numbers that seem to indicate success. Note that the enthusiastic proponents of Red-Light Cameras, do not mention ratioing to non-RLC comparison intersections, and never acknowledge applying tests for statistical significance. Often, the reports of success are anecdotal.

Summary

In bullet form, here is what it all boils down to:

- The small number of accidents per year and the large number of confounding variables obscure the conclusions that can be drawn from simple before and after accident reports for Red-Light Camera intersections.

- Simple before/after accident reports that do not use non-RLC comparison intersections, and do not apply statistical tests for significance offer meaningless information.

- Statistically sound studies show:

  - Red-light cameras **do** issue voluminous red-light running tickets, and the statistics on the 50% drop off in volume over two years is highly accurate.

  - 70% to 80% of red-light running happens in the first second of red[3, 6]. RLCs are most effective in discouraging first-second red-light runners. But, cross-traffic starting up on a delayed green avoids accidents with these first-second violators.

- The severe angle collisions are caused by red-light runners well after that first second of red[3, 6]. These are drivers who are distracted or irrationally reckless. RLCs are relatively ineffective in discouraging these later red-light runners.

- Red-Light cameras have a statistically non-significant effect on angle collisions and fatal crashes, but increase rear-end collisions up to 42% and total crashes up to 40%[6].

For further information see the PowerPoint presentation "Red-Light Cameras: The Good, the Bad, and the Uncertain." at www.TNLiberty.org.

References

1. Philip R. Bevington and D. Keith Robinson, *Data Reduction and Error Analysis for the Physical Sciences*, WCB/McGraw-Hill, ISBN 0-07-911243-9, 1992.

2. Nattaporn Yaungyai, *Evaluation Update of Red Light Camera Programming in Fairfax County, Virginia*; Master's Thesis, Virginia Polytechnic Institute and State University, April 2004.

3. Mark Burkey and Kofi Obeng, *A Detailed Investigation of Crash Risk Reduction Resulting from Red Light Cameras in Small Urban Areas,* North Carolina Agricultural & Technical State University, Greensboro, NC, July 2004.

4. Ron Jenkins, R. W. Gould and Dale Gedcke, Quantitative X-Ray Spectrometry, Marcel Dekker Inc., New York, 1981, Chap. 4.

5. Nicholas J. Garber *et al*., *The Impact of Red Light Cameras (Photo-Red Enforcement) on Crashes in Virginia,* Virginia Transportation Research Council, Charlottesville, VA, June 2007.

6. Barbara Langland-Orban et al., Red Light Running Cameras: Would Crashes, Injuries and Automobile Insurance Rates Increase If They Are Used in Florida? Florida Public Health Review, 2008; 5:1-7, and 5: 47-52.