PROFESSOR DR. HELMUT SCHNEIDER
CHRISTIAN RASCHKE

HSRG AT LSU

ANALYSIS OF LOUISIANA CRASH DATA

EFFECTIVENESS OF RUMBLES STRIPS IN REDUCING RUN-OFF-THE ROAD CRASHES ON LOUISIANA'S INTERSTATES

LOUISIANA STATE UNIVERSITY
ISDS DEPARTMENT
PATRICK TAYLOR HALL 3199
BATON ROUGE

Analysis of Louisiana Crash Data:

Effectiveness of Rumbles Strips in reducing run-off-the road crashes on Louisiana's Interstates

Summary

This report presents the results of an analysis of the effectiveness of rumble strips in preventing crashes on Louisiana Interstate Highways. We outline the Empirical Bayesian statistical method that was used to conduct a Before/After Study and we provide an index of effectiveness that reflects the performance of the rumble strips for various subsamples.

Data

Population of Rumble-Strip-Treated Road Segments

The Louisiana Department of Transportation and Development provided data on 36 road sections which had rumble strips installed between 1999 and 2007. Data include records for Interstates I-10, I-12, I-20, I-49, I-55 and I-610. Each segment is identified by the begin mile post and the end milepost of the treatment section; work order dates and final inspection dates are provided. The length of the treatment sections ranges from 0.34 miles to 16.39 miles with average of 5.98 miles.

The time between the work order date and the inspection date varies from 26 days up to 1057 days, with an average of 276 days between work order and final inspection. Four of the segments had no final inspection date. The large variation of time between work order date and inspection date poses a problem for choosing a time frame for selecting the crashes. In an ideal situation one would like to select crash data right before the rumble stripes were installed and crash data from the time right after the rumble stripes were installed. These data are not available and not knowing the exact date when the rumble stripes were installed complicates the analysis. It seems unlikely that the installation of the rumble stripes took more than a month. Thus the long times between work order and inspection date may be due to recording issues.

For this study the following method was used to select a time frame: 1) If the year of the work order and the year of the inspection coincide, then that year was chosen as the Effective Year; 2) If the year of the work order date and the year of the inspection date do not coincide, then the Effective Year was set to the year of work order date if the work order was placed before the month of October, and the Effective Year was set to the year of the work order date *plus one* if the work order was place during or after the month of October. We then use these calculated Effective Years in the before/after study by examining crashes from the year before and the year after the Effective Year.

Crash Data

The crash data from 1999 to 2007 were used to analyze the effectiveness of the rumble stripes in reducing run-off-road crashes. There are several issues that had to be dealt with when collecting crash data on the road segments:

1. Intersection crashes had to be excluded because of the specific coding of intersection crashes. Until 2007 all crashes within 100 feet of an intersection were coded to the route with the lower number. Thus if a crash occurred on an underpass of an interstate this crash was coded to the milepost of the interstate. Also, off and on-ramp crashes were coded to the interstate.
2. Rumble strip sections that were effectively installed in 1999 and in 2007 had to be excluded from the analysis because we needed at least one year prior and after the effective installation in order to conduct the before/after study.
3. There was a slight overlap in the treatment section as provided by the DOTD. For instance, the spreadsheet shows that Interstate 10 had rumble strips installed up to mile marker 10.22 in 2002, and then in 2005 another set was installed starting at mile marker 10.16 up to marker 20.12. Since this overlap is so slight (it is in the decimal points) we eliminated the overlap by subtracting half of the absolute difference from the higher number, and adding the same to the lower number. This makes the numbers continuous without any overlap. Since the overlap is likely caused by recording issues and is so small, our results will not be affected.

We analyze several different subsamples of overall crash counts to differentiate the effect of rumble strips on different types of accidents.

The first set includes accidents with any number of vehicles involved on the relevant sections. We examine this entire sample, but also restrict the sample to accidents that 1) occurred at night, 2) are of the run-off-the-road type, 3) involve fatigued, inattentive, and apparently asleep drivers, 4) occurred at night and involve drivers as in point (3), and 5) are of the run-off-the-road type and involve drivers as in point(3). For this analysis we consider accidents to have occurred at night when the recorded crash time is between 9:00 pm and 5:00 am.

The next set includes accidents that involve only a single vehicle. The same subsamples as described above are taken and we analyze these subsamples separately. Finally, the last set includes accidents which involve injuries, and we again restrict our sample as described above.

Methodology

We use as before/after study to investigate the effect of rumble strips on accident counts. The idea straightforward: Compare the accident counts from the post-treatment period to the accident counts we would have expected in that period if the treatment had not occurred. A naïve estimate of this expected accident in the after period count may be the observed account from the after period. But, when trying to estimate the effects of a road safety measure it is not enough to simply compare accident counts from before and after the implementation of the measure. This is due to several reasons.

First, the characteristics of the road segments in question may change over time. For example the traffic density may change resulting in a change of observed accidents that is unrelated to the safety measure in question. The bias here can go either way: one might underestimate or overestimate the effect of the safety measure depending on the type of change that occurred on the treated road section.

The second issue is called ‘regression to the mean’ and is a well known problem in the road safety literature. Among others, regression to the mean occurs due to the selection of treatment sites. For example, the decision to install rumble strips along a particular road section is not made at random, but sites are selected due to their accident histories. We cannot determine for certain that the change in sites’ accident histories is the result of a purely systematic process. Therefore it is likely that, following large deviations from the mean accident counts for a particular road section, the next observation will deviate less far. Not accounting for this mean reversion will bias results to make the safety measure seem more effective than it really is.

Empirical Bayes method

In order to deal with the issues describes above, we follow a standard procedure from the road safety literature, called the 'Empirical Bayesian Method'. This method fits a nonlinear model to the data and then uses the predicted values from the model in a weighted average with the actually observed accident counts to produce the expected value of crash counts. The method used in this study follows closely Persaud et al (2004), who study the effectiveness of centerline rumble strips on rural two-lane roads. Also see Hauer (2001, 2002) for a good discussion of the empirical Bayes method. The following is a brief discussion of the methodology used in this study.

The idea is to compare the actual observed accident count after the installation of rumbles strips, K^A to the expected count we would have observer had the strips not be installed, $E[\widehat{K}^A]$. The Poisson model is an intuitive candidate to use for count data problems such as estimating accident counts. However, since equidispersion is not a reasonable assumption in our context, a nonlinear model with an underlying Negative Binomial distribution is estimated. The Negative Binomial distribution can be interpreted as a compound distribution of the Poisson distribution with a Gamma distributed dispersion parameter (Hilbe 2007). Using negative binomial regression to model crash counts is therefore quite intuitive, and is also standard in the literature. This will allow us to capture a dispersion parameter, call it φ, to be used in the construction of an Effectiveness Index. We fit the Negative Binomial model according to the straightforward specification

$$E[K_i|\boldsymbol{X},\boldsymbol{H},\boldsymbol{T}] = \widehat{K}_i = \exp(\boldsymbol{x}_i'\boldsymbol{\beta} + H_i\gamma + T_i\delta) + \varepsilon \qquad (1)$$

where $\boldsymbol{x}_i$ is a vector of characteristics of subsection i such as the average daily traffic on that subsection, and the subsection length, and includes an intercept term; H is a dummy for the highway number, which is included to capture some of the unobserved heterogeneity between sections; T is a set of dichotomous indicators for the year of observation. The above is our specification of what is sometimes called a safety performance function, or SPF.

Since the average daily traffic varies within each treatment section and simply averaging the values results in a poor fir of our model as well as a low number of observations for the estimation, we create subsections within each treatment section. The subsections are identical in their average daily traffic. We use these subsections to fit the negative binomial model, then aggregate the predictions back to the treatment section level.

In estimation the specification in equation (1) we obtain a predicted value $\widehat{K}_i$ for all periods and all subsections. The empirical Bayes part of the study is to construct a weighted average of the prediction and the actual observed count, and then apply a factor of the sum of the SPF predictions from after the treatment over the sum of the SPF predictions before the treatment. This yields

$$E\left[\widehat{K}^A\right] = g\{\omega \cdot K_i^B + \ (1-\omega) \cdot \widehat{K}_i\}$$

where the weight ω is derived from the estimation of the SPF, specifically the predicted values and the (for each model constant) dispersion parameter, such that

$$\omega_i = \frac{\widehat{K}_i}{\widehat{K}_i + 1/\varphi}$$

To minimize concerns regarding latent changes in section characteristics over time that are not picked up by our dummies, we concentrate on the year immediately before, and immediately after the rumble strip installation for a particular section. At this point we aggregate the predictions back to the section level by adding the measures across the subsections contained within a treatment section.

In order to obtain the Index of Effectiveness, θ, we again follow Persaud et al (2004) and apply the following formulas for the index, and for the standard deviations. Let the subscript 'sum' denote that the measures have been aggregated across all sections in order to obtain an index of overall effectiveness.

$$\theta = \frac{K_{sum}^A / E\left[\widehat{K}_{sum}^A\right]}{1 + [var(\widehat{K}_{sum}^A)/E\left[\widehat{K}_{sum}^A\right]^2]} \qquad (2)$$

$$Std\ Dev\ (\theta) = \frac{\theta \cdot \sqrt{[var(K_{sum}^{A})/(K_{sum}^{A})^2] + \left[var(\widehat{K}_{sum}^{A})/E[\widehat{K}_{sum}^{A}]^2\right]}}{1 + \left[var(\widehat{K}_{sum}^{A})/E[\widehat{K}_{sum}^{A}]^2\right]} \tag{3}$$

The interpretation of the index is straightforward: $100(1-\theta)$ is the percentage change in crashes of the observed count and the expectation without the rumble strips. For example, if $\theta = 0.85$, then we conclude that following the installation of rumble strips crashes have been reduced by 15%.

Results

Any number of vehicles involved in crashes

The following Table 1 is a summary of results from the SPF estimation exercise. All models include controls for Highway, as well as year of observation dummy. Though sometimes significant, coefficients on these controls are not reported to save space.

Table 1: SPF Estimation Results (Any number of vehicles involved)

	Any	*Night*	*ROR*	*Driver Condition*	*Driver Condition & Night*	*Driver Condition & ROR*
Const	2.6050*** (0.1932)	1.0875*** (0.2843)	2.2051*** (0.3228)	1.8908*** (0.3183)	0.6469** (0.3132)	1.7559*** (0.3294)
ADT	0.0062*** (0.0013)	0.0085*** (0.0019)	-0.0083*** (0.0022)	0.0125*** (0.0432)	0.0065*** (0.0020)	-0.0086*** (0.0023)
Length	0.3427*** (0.0248)	0.3332*** (0.0363)	0.3929*** (0.0425)	0.3938*** (0.0432)	0.3612*** (0.0405)	0.3957*** (0.0437)
Dispersion	1.1484	2.3449	3.1207	3.2647	2.7093	3.1796
Deviance	1.1307	1.0003	0.8964	1.0552	0.8320	0.8127

*** denotes statistical significance at the 99% level;
** denotes statistical significance at the 95% level

All coefficients reported are highly statistically significant, which is congruent with intuition. Again, note that control variables' coefficients were omitted in the presentation. The value reported for *Deviance* is a criterion for assessing goodness of fit. Values wildly away from unity would indicate a potential model misspecification. It is interesting to note that the sign of the coefficient on Average Daily Traffic flips whenever we condition on run-off-the-road accidents. This indicates that more traffic usually increase the predicted number of crashes in a particular section; however, run-off the-road accidents are more likely on less densely traveled road. This is again intuitive because empty roads contribute greatly to drivers' inattention and 'Highway Hypnosis'. Note that when conditioning only on the state of mind of the driver, all types of accidents are considered. This category includes inattentiveness and being distracted for any reason, not only because of boredom and monotony on the road. Therefore the positive coefficient here is consistent with the fact that in general inattentive drivers are more likely to cause crashes with other vehicles.

The following table displays the resulting Effectiveness Indices for all types of crashes. Again, these crashes may involve any number of vehicles. Results for single vehicle crashes only are presented in the next section.

Table 2: Effectiveness Indices (Any number of vehicles involved)

	Any	*Night*	*ROR*	*Driver Condition*	*Driver Condition & Night*	*Driver Condition & ROR*
Effectiveness Index, θ	0.8909	0.7271	0.8036	0.7995	0.7671	0.8703
Std Deviation	0.0033	0.0198	0.0208	0.0046	0.0419	0.0345
Percent Accident Reduction	10.91%	27.30%	19.65%	20.05%	23.34%	13.02%

The pattern of effectiveness of the rumble strips can be readily seen from Table 2. While there is an improvement in accident counts for any subsample, we note that the strips are more effective at nighttime and for run-off-the road crashes. This is intuitive because the strips are designed to prevent exactly those types of run-off-the-road crashes, and crashes that occur due to driver fatigue and inattentiveness (i.e. those that tend to occur during nighttime hours.)

Single Vehicle Crashes

This section presents results when restricting our sample to accidents that only involve one vehicle. Intuition suggests that rumble strips might be especially effective in preventing those types of accidents. Table 3 shows the estimation results for the negative binomial model of the safety performance function.

Table 3: SPF Estimation Results (Single Vehicle Accidents)

	Any	*Night*	*ROR*	*Driver Condition*	*Driver Condition & Night*	*Driver Condition & ROR*
Const	1.9830***	0.8297***	1.8365***	1.6048***	0.6319*	1.4713***
	(0.2910)	(0.2810)	(0.3199)	(0.3259)	(0.3275)	(0.3294)
ADT	0.0015	0.0054***	-0.0083***	-0.0014	0.0044*	-0.0080***
	(0.0019)	(0.0019)	(0.0023)	(0.0021)	(0.0023)	(0.0024)
Length	0.3706***	0.3076***	0.3962***	0.3877***	0.3314***	0.3923***
	(0.0374)	(0.0353)	(0.0421)	(0.0427)	(0.0417)	(0.0433)
Dispersion	2.5744	2.0522	2.9958	3.1921	2.4470	3.0066
Deviance	1.0383	0.9158	0.8573	0.8786	0.6939	0.7855

*** denotes statistical significance at the 99% level;
* denotes statistical significance at the 90% level

Again, most of the coefficients are highly significant and sign of Average Daily Traffic flips when the sample is restricted to run-off-the-road accidents. The fit of the model as indicated by the Deviance is very good to reasonable.

Next we calculated the Effectiveness Indices for single vehicle accidents. Results are presented below in Table 4.

Table 4: Effectiveness Indices (Single Vehicle Accidents)

	Any	*Night*	*ROR*	*Driver Condition*	*Driver Condition & Night*	*Driver Condition & ROR*
Effectiveness Index, θ	0.8078	0.7175	0.8804	0.8231	0.7976	0.9011
Std Deviation	0.0129	0.0424	0.0309	0.0253	0.0828	0.0462
Percent Accident Reduction	19.22%	28.31%	12.00%	17.71%	20.47%	9.98%

Rumble strips seem to be more effective in preventing overall single vehicle accidents when compared to the results in Table 2. Also, the Effectiveness Index on night time crashes is slightly higher. Further restricting the sample results in significant improvements after installation of rumble strips, yet lower point estimates of effectiveness when compared to the over crash counts of any number of vehicles involved.

Crashes Involving Injuries

It is also interesting to investigate the effects of the rumble strips on accidents that involve injuries. We still expect a positive effect; yet, if injuries are viewed as a proxy for the severity of an accident, then we can examine whether those accidents that are being prevented by rumble strips tend to generally be more severe accidents, or whether they tend to be not so severe (at least not involving injury.) While a separate data field for the severity of an accident is available, the coding is not very consistent. Therefore, proxying severity with injuries is a viable alternative.

The following is the table of the Effectiveness Indices when only considering accidents that involve injuries. We did not model the total number of injuries that occurred because this depends on many other unobservable vehicle, as well as driver and passenger characteristics. Instead, we re-estimated the model using only accidents in which at least one injury occurred. This method also allows us to shed light on the question of severity as proposed above. Effectiveness Indices are presented in the following Table 5. SPF estimation results are omitted to save space; they did not produce any surprising results.

Table 5: Effectiveness Indices (Injury Crashes)

	Any	*Night*	*ROR*	*Driver Condition*	*Driver Condition & Night*	*Driver Condition & ROR*
Effectiveness Index, θ	0.9358	0.8399	1.0047	0.8739	0.8131	0.9650
Std Deviation	0.0113	0.0495	0.0511	0.0166	0.0831	0.0809
Percent Accident Reduction	6.42%	16.01%	-0.47%‡	12.61%	18.69%	3.77%‡

‡ Note that this is statistically indistinguishable from zero.

Overall, there is some improvement in crashes involving injuries after the rumble strips are installed. However, the value for run-off-the-road crashes, as well as run-off-the-road crashes involving a fatigued or inattentive driver is not statistically different from zero. This may be for a variety of reasons, about which we can only speculate. One certain possibility is again that the coding for these types of crashes may not always be correct. A general limitation of the result in Table 5 is that the data on injuries may have flaws as well. For instance, there are certain accidents which are coded as 'no injury', yet there is an indicator that denotes someone was transported by ambulance. The direction of the bias is uncertain in general; however, if we assume that over time the coding has improved as officers become more familiar with the reporting format, then Table 5 would underestimate the true effect of the rumble strips.

The decreased magnitude and partial insignificance of the effectiveness index suggests that the rumble strip do not necessarily prevent more severe types of accidents. Rather, most of the improvement in crash counts as reported in Table 2 and Table 4 may be contributed to more minor accidents.

Fortunately there are not enough fatal crashes in Louisiana to make an analysis of fatal accidents possible using our approach. The prevalence of zeros in the data makes it impossible to satisfactorily fit a safety performance function.

Conclusions

We have used the Empirical Bayes method to examine the effectiveness of rumble strips in the state of Louisiana. We estimated a safety performance function of the negative binomial form, and constructed an Index of Effectiveness that reflect the decrease in crash counts from before and after the installation of rumble strips.

The data suggest that rumble strips are highly effective in preventing accidents. The highest level of improvement occurs in nighttime crashes and in crashes where the driver was fatigued or inattentive. There is some difference between single vehicle crashes and crashes potentially involving more than one vehicle, where the strips tend to be overall slightly more effective in preventing single vehicle crashes. Examining accidents that involved injury, the results are less strong. While there is an overall decrease in crash counts, the effect is lesser than when compared to all accidents regardless of injuries.

Bibliography

Hauer, Ezra. "Estimating Safety by the Empirical Bayes Method." Tutorial, University of Toronto, Department of Civil Engineering. Toronto, CA, 2001.

Hauer, Ezra. *Observational Before-After Studies in road Safety.* Oxford: Pergamon, 2002.

Hilbe, Joseph M. *Negative Binomial Regression, 1ed.* Cambridge University Press, 2007

Persaud, Bhagwant N, Richard A Retting, and Craig A Lyon. "Crash Reduction following installation of centerline rumble strips on rural two-lane roads." *Accident Analysis and Prevention*, 2004: 1073-1079.