Submitted to the Journal of the American Statistical Association

USING LOG-LINEAR MODELS TO

ASSESS STAGGERED POLICY CHANGES

Thomas R. Willemain, Ph.D.

Nelson S. Hartunian, Ph.D.

Paul L. Zador, Ph.D.

March 26, 1982

Abstract

This paper considers the analysis of series of Poisson counts generated in multiple jurisdictions. At staggered times in some of the jurisdictions, similar policy changes are made, altering the Poisson intensity. For instance, over a period of several years, roughly half the states repealed or weakened their laws mandating helmet use by motorcyclists. We describe an iterative method that facilitates the log-linear modeling of such situations. Using Monte Carlo simulation, we document conditions under which the method provides unbiased and consistent estimates of the effect of the policy change.

Keywords: log-linear models, Monte-Carlo simulation, highway safety, Poisson process, motorcycle helmets, evaluation

AUTHORS' FOOTNOTE

Thomas R. Willemain is Associate Professor of Public Policy, John F. Kennedy School of Government, Harvard University, Cambridge, MA and Senior Vice President, SMART, HARTUNIAN & WILLEMAIN, INC., Belmont, MA. Nelson S. Hartunian is Senior Vice President, SMART, HARTUNIAN & WILLEMAIN, INC. Paul L. Zador is Senior Statistician, Insurance Institute for Highway Safety, Washington, D.C.

The authors acknowledge the assistance of Charles N. Smart, Brian O'Neill, Michael Stoto, William DuMouchel, Sharon Hartunian and Linda Barry.

This work was supported by a contract from the Insurance Institute for Highway Safety for SMART, HARTUNIAN & WILLEMAIN, INC. Neither the Insurance Institute for Highway Safety nor those whose help we acknowledged share responsibility for the results and conclusions of this work.

1. Introduction

A common situation in policy evaluation involves staggered (non-simultaneous) replications of policy changes in some but not all jurisdictions. For instance, in 1975 all but three states had laws requiring that motorcyclists wear safety helmets, but in 1976 congressional elimination of federal financial penalties for states without such laws led to their repeal in about half the states over the next several years [U.S. Department of Transportation 1980]. We report here on a new method of analyzing the effects of such changes when they are well modeled as multiplicative changes in Poisson rates.

Our work was motivated by the need to assess the impact of the repeal of motorcycle helmet laws on specific demographic groups. Since the basic data were monthly fatality counts and since we anticipated that all effects would be multiplicative, it was natural to consider using methods for the log-linear analysis of contingency tables [Bishop, Fienberg and Holland 1975]. However, in order to make separate estimates of the impact of repeal on various demographic groups, we were forced to work with very sparse count data. We could compensate for the sparcity caused by demographic disaggregation by reaggregating counts across space and time, but the non-simultaneity of helmet law repeals, together with the multiplicative nature of effects, meant that the aggregated counts would not be in a form compatible with log-linear analysis. We were able to overcome this limitation by using an iterative method of adjustment that, in essence, first "cancelled out" the effect of helmet law repeal in each state that changed its law, then aggregated the resulting "adjusted" counts of deaths into tables fit by log-linear models [Hartunian et al 1982]. A detailed description of this methodology is given in Section 2.

Alternative approaches had some obvious drawbacks in our situation. Watson, Zador and Wilks [1980] had used a matched-regression method to estimate the effect of motorcycle helmet law repeal on a statewide basis. In their analysis, each state that repealed its law was matched with several others that did not, but were otherwise similar. Regression analysis was used to summarize the pre-repeal relationship between the square root of monthly deaths in the state that repealed its laws and the square root of the aggregate monthly death toll in the matched states. Then the pre-repeal relationship was extrapolated to the post-repeal period, with discrepancies attributed to the effect of law repeal. Unfortunately, the extremely sparse counts associated with particular demographic groups (averaging on the order of 1/2 or fewer counts per state per month) precluded this kind of state-by-state analysis in our work.

The method we chose to use, involving log-linear fits to tables of adjusted counts, provided a way to aggregate data across states without requiring extensive computational capabilities. For example, direct maximum likelihood estimation would have required the solution of a system of at least 54 simultaneous nonlinear equations.

This paper reports the results of Monte Carlo simulations designed to assess the bias and precision of estimates obtained by our methodology.

2. Description of Methodology

In this section, we provide a formal description of our method of fitting log-linear models to tables of adjusted counts.

2.1. Adjustment

Consider a non-homogeneous Poisson process with intensity a function of jurisdiction and time. For purposes of exposition, we will consider jurisdictions to be states and reckon time in months. We regard the Poisson intensity as a multiplicative function of an initial intensity specific to each state, a temporal factor (such as trend plus seasonality) common to all states, and a multiplicative effect of policy change. It is this latter effect that we seek to estimate.

Formally, let

I_{it} = Poisson intensity in state i during month t

I_{io} = initial intensity in state i

F(t) = temporal function representing trend plus seasonality

X_{it} = 1 if policy has been changed in state i by month t
0 otherwise

E = magnitude of effect of policy change in log scale.

Our model of the Poisson process is log-linear

$$I_{it} = I_{io} \exp[F(t)] \exp[EX_{it}] \qquad (2.1)$$

$$\log I_{it} = \log I_{io} + F(t) + EX_{it} \;. \qquad (2.2)$$

Let A represent the set of all states that do not change their policies ("comparison states"). The aggregation of counts across comparison states produces a Poisson process with intensity

$$I_{At} = \sum_{i \in A} I_{it} = \exp[F(t)] \sum_{i \in A} I_{io} \qquad (2.3)$$

provided counts are independent across states. Thus the counts in separate comparison states can be aggregated into one large equivalent series of Poisson counts that is log-linear because the term $\sum_{i \in A} I_{io}$ can be treated as a single parameter.

Unfortunately, this is not the case for counts in states that change their policies ("change states"). Let B represent the set of all change states and again assume independence across states. The aggregated series has intensity

$$I_{Bt} = \sum_{i \in B} I_{it} = \exp[F(t)] \sum_{i \in B} I_{io} \exp[EX_{it}]. \qquad (2.4)$$

Here the parameter of interest, E, is ensconced within the summation and the model of counts is no longer log-linear.

To overcome this difficulty, imagine performing a process of "adjustment". If the value of E were known a priori, we could consider the adjusted Poisson process with rate

$$\tilde{I}_{it} = I_{it} \cdot \exp[-EX_{it}] = I_{io} \exp[F(t)]. \qquad (2.5)$$

This adjusted process is log-linear for both comparison and change states, and any aggregation across states will produce a process that remains log-linear. Hence, given the value of the effect of the policy change E, the remaining nuisance parameters (i.e. the I_{io} and the parameters in F(t)) could be determined by log-linear analysis of contingency tables. Of course, the value of E is not given -- it is the key unknown -- but we can estimate it by an iterative process.

2.2. Iteration

In practice, the adjustment method works iteratively in conjunction with log-linear fitting. Figure 1 gives a representation of the iterative process. First, a trial adjustment factor E' is used to convert the observed counts into adjusted counts in change states during months after the policy change. For instance, if we expect the policy change to increase the counts, we try a negative value of E' to offset the increase and produce a set of smaller adjusted counts. Second, the adjusted counts are aggregated across states and months to form a contingency table with a serviceable number of counts in each cell. Third, this table is fit by a log-linear model using the iterative proportional fitting algorithm of Haberman [1972]. Fourth, a measure of goodness of fit is used to assess how well the nuisance parameters explain the variation remaining after adjustment for the (presumed) effect of the policy change. Because the adjustment process changes the cell totals in the contingency table and the traditional chi-square and log-likelihood ratio statistics are sensitive to the cell totals, we could not use these standard measures of goodness of fit to compare different values of the adjustment factor. Instead, we resorted to the heuristic solution of using as our criterion of fit the "normalized chi-square", i.e. chi-square divided by the total number of (adjusted) counts. The final step in the process involved searching for that value of the adjustment factor that produced the smallest value of normalized chi-square.

3. Monte Carlo Simulations

In this section, we provide a description of the simulations used to assess the estimates produced by the methodology in Section 2. We designed the simulations to explore the influence of a number of important background variables. We used in all instances reasonable facsimiles of the Department of Transportation Fatal Accident Reporting System (FARS) data on motorcycle fatalities [U.S. Department of Transportation, no date] that motivated our original analysis. However, we did not attempt to reproduce exactly all the peculiarities of the FARS data. All computations were performed on a Hewlett-Packard HP-9835A minicomputer.

3.1. Artificial Random Deviates

We generated Poisson counts in a standard way by adding together as many exponentially distributed intervals with appropriate mean as would fit into a month [Larson and Odoni 1981]. The exponential intervals were formed by transforming deviates uniformly distributed over the interval (0.0, 1.0). We generated these uniform deviates using a method of Knuth's [1969] especially adapted for the HP9835A minicomputer. Bauer [1981] reported that this random number generator has passed the spectral, digits and runs tests. We independently confirmed the results of the digits and runs tests and were able to verify that the cycle length of the generator exceeds 2E7. Finally, we examined the resulting Poisson variates both for distribution and serial independence using the chi-square test and found no statistically significant discrepancies.

As explained in Section 3.2.4, some simulations added normally distributed stochastic "shocks" to the underlying Poisson intensities. We used Knuth's polar method [8] to generate the normal deviates.

3.2. Simulated Data Sets

3.2.1. Temporal Factors

Our analysis of the FARS data on motorcycle fatalities revealed pronounced season effects and annual trends that were well-described by multiplicative models. Figure 2 illustrates by plotting the common logarithm of the total motorcycle death count in the United States. Analysis of the data in Figure 2 led us to use an approximate figure of 8E-3 for the monthly increase in the logarithm of deaths and to represent the seasonal factor as a sine wave of unit amplitude. Some of the simulations allowed for differences in trend and seasonality across comparison and change states.

3.2.2. Comparison and Change States

Our analysis of the FARS data used 4 regional groupings of states to help control for influences other than the change in helmet use laws. Across the 4 regions, the number of comparison states varied from 3 to 11, and the number of change states ranged from 5 to 8. For the simulations, we considered both a single region comprised of 5 comparison and 5 change states and two such regions. In the actual case of helmet law repeal, FARS contained 72 months of data, including 18 months before the first repeal. In the simulations, there were also 72 months of data, with repeals in the change states independently and uniformly distributed over months 13 to 72.

3.2.3. Initial Poisson Intensities

Our analysis of the FARS data was executed repetitively for seven demographic groups. These groups averaged from about 0.1 to about 3.0 deaths per month per state. Our simulations used values from 0.25 to 2.0 for the initial Poisson intensity .

3.2.4. Systemic Shocks

Besides "borrowing strength" for the estimation of parameters common to both groups of states, another reason to use comparison states in the analysis of policy changes is to protect against confounding caused by those external influences referred to as "shocks" in the time-series literature and as "history" in the literature on quasi-experimental designs [Campbell and Stanley 1963]. If such "shocks" are systemic in that they affect both the comparison and change states, they can be distinguished from the effects of policy change. Examples of systemic shocks salient for the motorcycle problem include weather phenomena and changes in the price or availability of gasoline. While we could imagine many ways to model the occurrence of systemic shocks, we chose to represent them by adding to the Poisson intensity of every state the same stochastic term drawn from a normal distribution. Across months, these shocks were independent and identically distributed. Some of our simulations involved no systemic shocks; others used shocks with zero mean and standard deviation 0.25.

3.2.5. Effects of Policy Change

Our analysis of the FARS data produced estimates of the effect of helmet law repeal ranging from E=0.052 for males aged 30+ to E=0.805 for females aged 0-19. Our simulations used values of E between -0.5 and +0.5. These values span a range of effects likely to be encountered in real problems.

3.3 Log-linear Models

3.3.1. Variables

The simulated Poisson counts in each state and month were aggregated into contingency tables having 3 dimensions: "group", "year" and "season". Group referred to the type of state, either comparison or change. Year actually referred to one of 3 successive 24 month periods. Season referred to one of 2 sets of months within each 12 month cycle. For instance, within one 24 month long "year", the first season might comprise months 1 to 6 and 13 to 18, while the remaining months would comprise the second season. One set of simulations tested the effect of changing seasonal definitions to achieve more or less contrast between the seasons.

3.3.2. Specification

In our analysis of the FAPS data, we divided the United States into 4 regions to standardize for possible regional differences in trend and seasonality. Ideally, such regional matching would make it unnecessary to include interaction terms in the specification of the log-linear model, in which case we need include only the one-way effects of group, year and season to model the adjusted counts.

However, it can happen that the spatial pattern of policy change will not be well mixed (as when policy in one state is influenced by its neighbors), and the change and comparison states may display differential seasonality and/or trend. We believed this to be the case in our analysis of helmet law repeal: our "Western" region had most of its comparison state deaths in California, whereas many of its change state deaths occurred in states with more severe winters like Colorado, Idaho, Montana, Oregon and Washington. Thus we permitted an interaction between group and season in the log-linear model used in our simulations, and some simulations stipulated a 2:1 differential in the magnitude of the seasonal effect.

Similar reasoning would indicate the value of including in the model specification an interaction between group and year, in case the trends in comparison and change states should be different. Unfortunately, we found it impossible to obtain reliable estimates of the effect of policy change when we used our methodology on simulated data sets constructed with differential trends. Apparently this problem can be traced to the extreme collinearity exhibited by the group-year term and the policy changes themselves, which occur only in change states and only in later months. Accordingly, we did not include a group-year interaction in the specification. We did, however, assess the quality of estimates made when the interaction term was not specified but the data did in fact contain differential trends.

For the simulation of a single region, the log-linear model specification was

$$\log n^{syg}_{ijk} = u_o + u^{s}_{i} + u^{y}_{j} + u^{g}_{k} + u^{sg}_{ik} \qquad (3.1)$$

where s=season $1 \le i \le 2$
y=year $1 \le j \le 3$
g=group $1 \le k \le 2$

n^{syg}_{ijk} =number of deaths in season i, year j and group k.

3.4. Experimental Design

3.4.1 Primary Analysis

Our primary analysis considered a single region and consisted of a factorial design with 6 factors:

- seasonal definitions (low v. high contrast)
- differential seasonality across groups (same v. different)

- differential trend across groups (same v. greater in comparison states v. greater in change states)
- systemic shock (present v. absent)
- initial Poisson intensity (low v. high)
- effect of policy change (none v. positive v. negative).

Table 1 provides details of the parameterization of the factorial design. The given factors and levels produced a total of 144 experimental conditions. For each condition we obtained 9 replications, which we used to determine the bias and precision of our estimates. We computed the bias as the difference between the mean of the 9 estimates and their true value. We assessed precision in terms of the sample standard deviation of the 9 estimates.

3.4.2. Secondary Analysis

Our analysis of a single region uncovered the fact that the estimates of experimental effect were inconsistent. This finding prompted us to perform a secondary analysis in two phases. In the first phase, we selected for analysis 24 pairs of values of Poisson intensity and size of policy effect in order to better document the inconsistency. For each pair we generated 9 repetitions and computed the standard error (i.e., the sample standard deviation of the 9 estimates) as a measure of precision.

In the second phase, we analyzed two identical regions. We speculated that the inconsistency of the estimates might be associated with small numbers of degrees of freedom and might disappear if we analyzed two regions rather than one. The model we fit for the two regions was

$$\log n_{ijkl}^{sygr} = u_o + u_i^s + u_j^y + u_k^g + u_{ik}^{sg} + u_l^r \qquad (3.2)$$

which is similar to (3.1) except for the addition of the region term. Model (3.2) had 16 degrees of freedom, compared to 5 degrees of freedom for (3.1).

We analyzed the 2 identical regions using both the original 24 pairs of parameter values and 6 additional pairs. For the 24 pairs, we reused the random deviates generated in the analysis of a single region by combining successive pairs of the 9 simulated data sets, obtaining 4 repetitions (and setting aside the 9th data set). Thus in our analysis of 2 regions the standard error was computed from 4 estimates. In all cases we used high contrast seasons, identical trend and seasonality in both types of states and no systemic shock.

4. Results

4.1. Primary Analysis

4.1.1 Bias

None of the factors except differential trend produced systematic biases in the estimates. We analyzed the 144 results using a regression equation containing dummy variables for each of the 6 factors and all possible two-way interaction terms for these 4 factors: season definition, differential seasonality, presence of systemic shock and initial intensity. The only significant regression coefficients were those for the 2 dummy variables representing the 3 levels of the trend factor. We achieved virtually identical performance from a simplified specification containing only the trend factor (t-values in parentheses):

$$\text{Bias} = \underset{(.67)}{.009} + \underset{(14.11)}{.255}\ \text{TREND+} - \underset{(-13.39)}{.242}\ \text{TREND-} \qquad (4.1)$$

$$R^2 = .84,\ \text{S.E.E.} = .009$$

where

TREND+ = 1 if trend is greater in change states, 0 otherwise

TREND- = 1 if trend is greater in comparison states, 0 otherwise.

Thus, the evidence is strong that, in the absence of differential trend, our methodology produces unbiased estimates of the effect of policy change.

The bias created by differential trend can be understood with the aid of the simplified representation in Figure 3. While the Poisson intensity in comparison states increases steadily (from A to B), the intensity in change states increases at a greater rate (from A to C). The policy change further increases the trend (from C to D). Ideally, the adjustment process would return the observed pattern in the change states (ACD) to its path in lieu of policy change (ACE). However, since the log-linear model has no term to account for differential trends, the adjustment process assumes this function and over-compensates (ACF) in an effort to equalize the average trend in the two groups. The result in Figure 3 is an overestimate of the impact of the policy change. Similar analysis for other combinations of trend and effect reveals that the sign of the bias follows the sign of the trend differential. The correspondence is reflected in the differences in the signs of the dummy variables in the regression analysis reported above.

4.1.2 Precision

We examined the precision of the estimates by regressing their standard errors in the factorial design against dummy variables for the factors and their 2-way interactions. As with the analysis of bias, only a few terms were statistically significant, and the essence of the result can be captured in a less complicated specification. The only factors that influenced precision were the initial Poisson intensity and the absolute value of the policy effect:

$$\underset{}{\text{STD ERROR}} = \underset{(26.8)}{.309} + \underset{(20.6)}{.291\,\text{EFF}} \underset{(-10.0)}{-\,.163\,\text{INT}} + \underset{(6.0)}{.120\,\text{EFF} \times \text{INT}} \qquad (4.2)$$

$R^2 = .91$ S.E.E. $= .056$ where

EFF $= 1$ if $|D| = .5$, 0 if $D = 0$

INT $= 1$ if $I_{ic} = 2.0$, 0 if $I_{ic} = .25$.

Since only the values of I and |E| had significant influence on the standard errors, we performed the secondary analysis mentioned in Section 3.4.2 to better assess their quantitative relationship. Our first observation was that the results seemed not to depend on the sign of the policy effect, so we used its absolute value in our analysis. The 24 data points are shown in Table 2.

Scanning down the columns, note that the standard errors decrease not to zero but rather to levels slightly above the absolute values of the policy effects. Thus the standard errors are of the same size as the effects, and it is never possible to obtain precise estimates, even with large counts.

We subtracted the size of the policy effects from the columns of Table 2 and studied the resulting residuals. We found a complex interaction between the initial Poisson intensity and the size of the effect. Increasing intensity reduced the residuals to a positive limit that depended on the size of the effect. For a given intensity, the residuals showed a minimum when the policy effect was roughly equal to the inverse square root of the intensity. We incorporated these observations into the specification of an ordinary least squares regression relationship:

$$\text{STD ERROR} = \underset{(1.5)}{.041} + \underset{(37.2)}{1.014}|E| + \underset{(6.9)}{.059}\left(|E| - 1/\sqrt{I_{io}}\right)^2 / \sqrt{I_{io}} \qquad (4.3)$$

$$R^2 = .99 \qquad \text{S.E.E.} = .074.$$

These results document the inconsistency of the estimates in the case of a single region.

Inconsistency was not a problem in the case of two regions, the data set for which is given in Table 3. The ordinary least-square fit corresponding to (4.3) was

$$\text{STD ERROR} = \underset{(2.9)}{.063} + \underset{(1.1)}{.025}|E| + \underset{(3.3)}{.022}\left(|E| - 1/\sqrt{I_{io}}\right)^2 / \sqrt{I_{io}} \qquad (4.4)$$

$$R^2 = .29 \qquad \text{S.E.E.} = .006.$$

We tried several other specifications but in all cases the results were the same: the coefficient associated with the size of the policy effect was non-significant. The best fit to the data involved no dependence on the size of the effect and the usual inverse square root dependence on intensity

$$\text{STD ERROR} = \underset{(.5)}{.009} + \underset{(6.2)}{.094}/\sqrt{I_{io}} \qquad (4.5)$$

$$R^2 = .58 \qquad \text{S.E.E.} = .005.$$

These results document the consistency of the estimates in the case of two regions.

5. Summary and Conclusions

It frequently happens that similar policy changes are made in different jurisdictions at different times. The example that motivated our work was the repeal of state laws mandating helmet use by motorcyclists. When the data are Poisson counts with a multiplicative structure and the effect of the policy change is a multiplicative alteration of the Poisson intensity, it is natural to use log-linear models to represent the data. However, when the counts in each jurisdiction are sparse and one attempts to compensate by aggregating across jurisdictions, the log-linearity of the data is lost. This occurs because of the different times at which the policies were changed (see equation (2.4)).

One way to restore the log-linearity of the data is to "adjust" or pre-multiply the counts by a factor that just offsets the effect of the policy change. Of course, to properly adjust the series of counts one needs to know the size of the effect that one seeks to estimate. Fortunately, this problem can be solved by an iterative process. The iterations involve a cycle of adjustment followed by log-linear fitting. The cycle ends when the adjusted table of counts is best fit by the non-policy related variables.

We examined the properties of this methodology using Monte Carlo simulation methods. The synthetic data sets were patterned after the counts of motorcyclist fatalities in the United States. We reported two main empirical findings. First, the method of adjusted log-linear analysis provides unbiased estimates so long as there are no differential trends in jurisdictions that do and do not change their policies. Second, the method provides consistent estimates whenever there are sufficient degrees of freedom remaining after the log-linear model has been estimated.

Given the pertinence and power of log-linear models, it is useful to have a method that permits their application to the common problem of assessing policy changes staggered across jurisdictions. We have described such a method and illustrated its properties in an empirical way. Future work might develop theoretical assessments of the estimates provided by the method and make comparisons with alternative methods.

FIGURES AND TABLES

FACTOR	LEVEL	PARAMETER	PARAMETER VALUE
Definition of season	Low contrast	Months in season 1	1 to 6, 13 to 18
	High contrast	"	1 to 2, 9 to 14, 21 to 24
Seasonality	Same	Sine wave amplitude	1.0 in both groups
	Different	"	1.0 in change, 0.5 in comparison
Trend	Same	Log of monthly increase	.008 in both groups
	Change greater	"	.008 in change, .004 in comparison
	Comparison greater	"	.004 in change, .008 in comparison
Systemic shock	Absent	std. dev. of Normal	0.0
	Present	"	0.25
Initial intensity	Low	Deaths per month per state	0.50
	High	"	2.00
Effect of policy	None	Increment to log of intensity	0.0
	Increase	"	0.5
	Decrease	"	-0.5

Table 1: Parameters of factorial design

ABSOLUTE VALUE OF EFFECT OF POLICY CHANGE

Value of Initial Poisson Intensity	0	1/4	1/3	1/2	3/4	1	5/3
1/4	.59	---	.51	---	---	1.14	1.78
1/2	.45	---	.43	.61 .57	---	---	1.79
3/4	.11	---	.40	---	---	1.06	1.77
1	---	---	---	---	---	----	---
2	.04	---	---	.54 .55	---	1.06	---
4	---	---	.36	---	---	---	1.77
8	.06	---	.36	.54	---	1.06	1.77
12	---	---	---	---	---	---	---

TABLE 2: Standard Errors of Estimates From Simulations of One Region.

ABSOLUTE VALUE OF EFFECT OF POLICY CHANGE

Value of Initial Poisson Intensity	0	1/4	1/3	1/2	3/4	1	5/3
1/4	.19	---	.22	.24	---	.23	.10
1/2	.05	---	.10	.15	---	---	.21
3/4	.12	---	.13	---	---	.19	.10
1	---	.11	---	.06	.26	---	.07
2	.04	---	---	.07	---	.05	---
4	---	---	.08	.06	---	---	.05
8	.04	---	.04	.04	---	.02	.04
12	---	---	---	.03	---	.05	---

TABLE 3: Standard Errors of Estimates From Simulations of Two Regions.

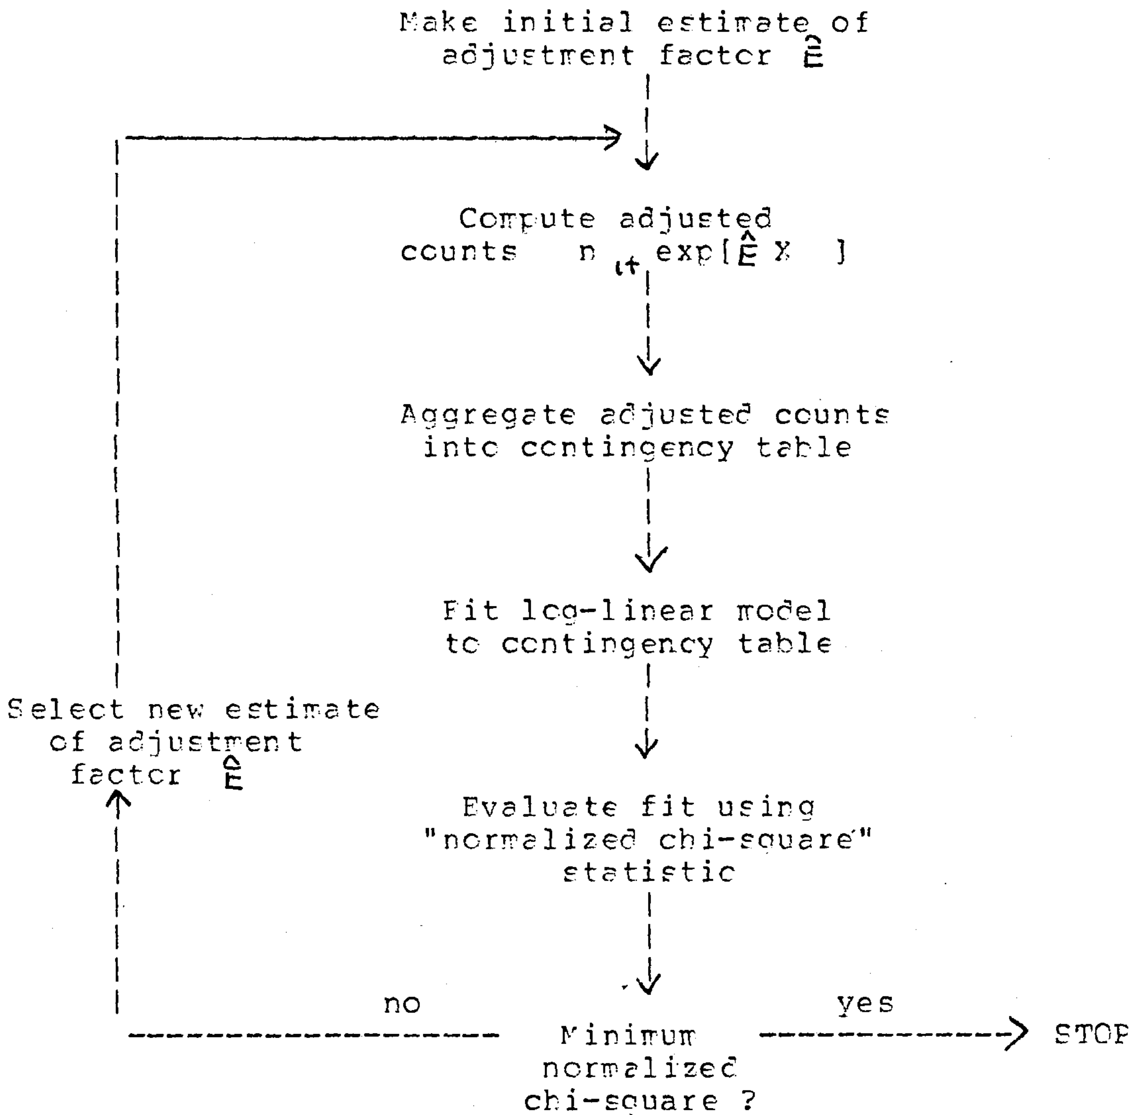

Figure 1: Iterative procedure for estimating the effect of policy change

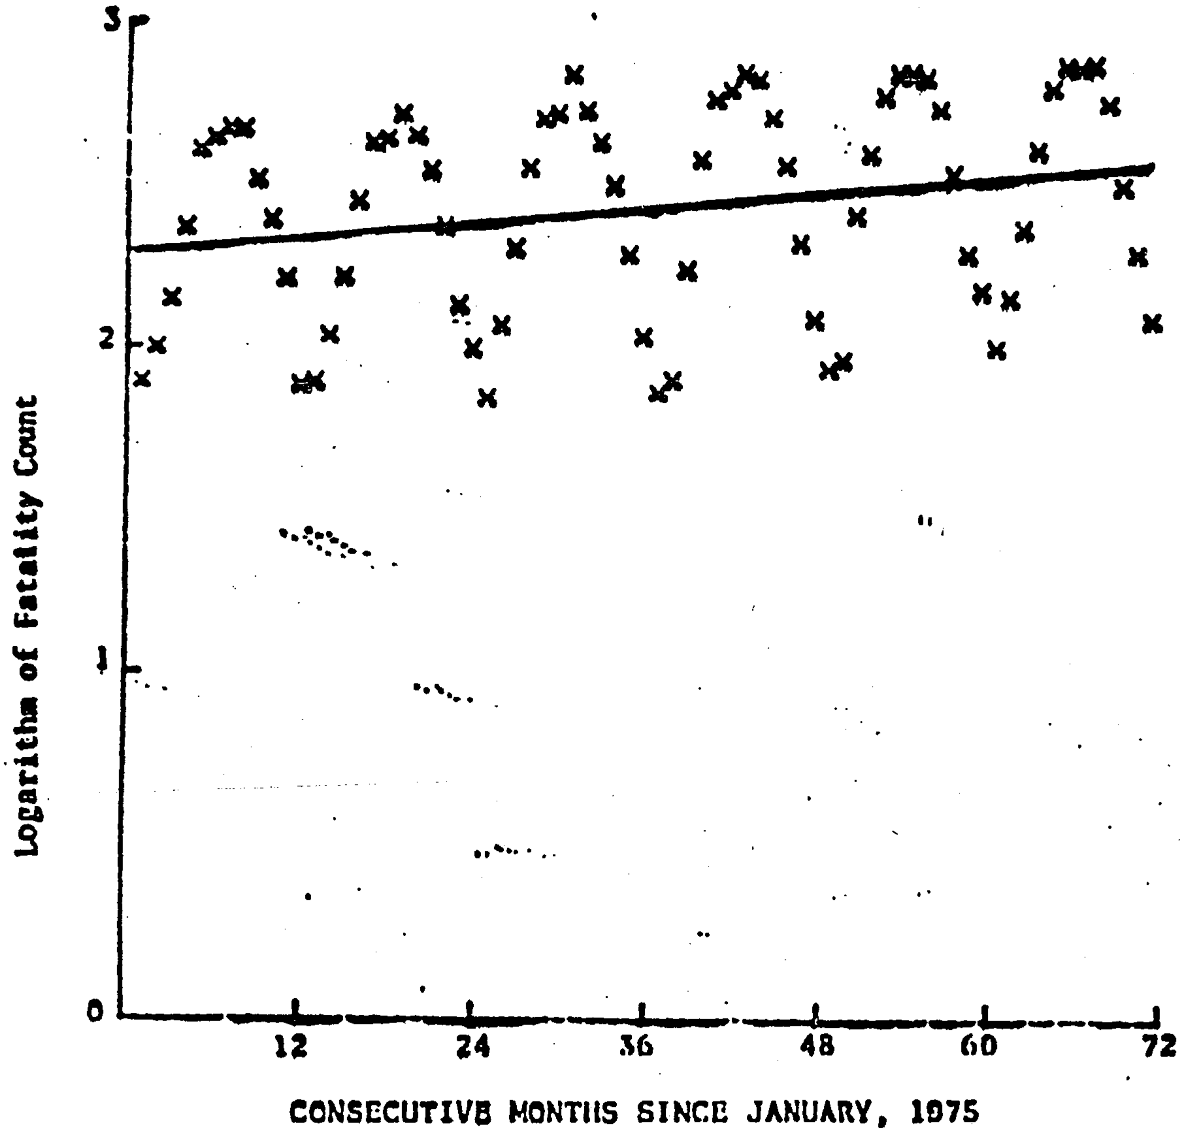

Source: Fatal Accident Reporting System, U.S. Dept. of Transportation.

Figure 2: Illustrating Multiplicative trend and seasonality in U.S. motorcycle fatality counts.

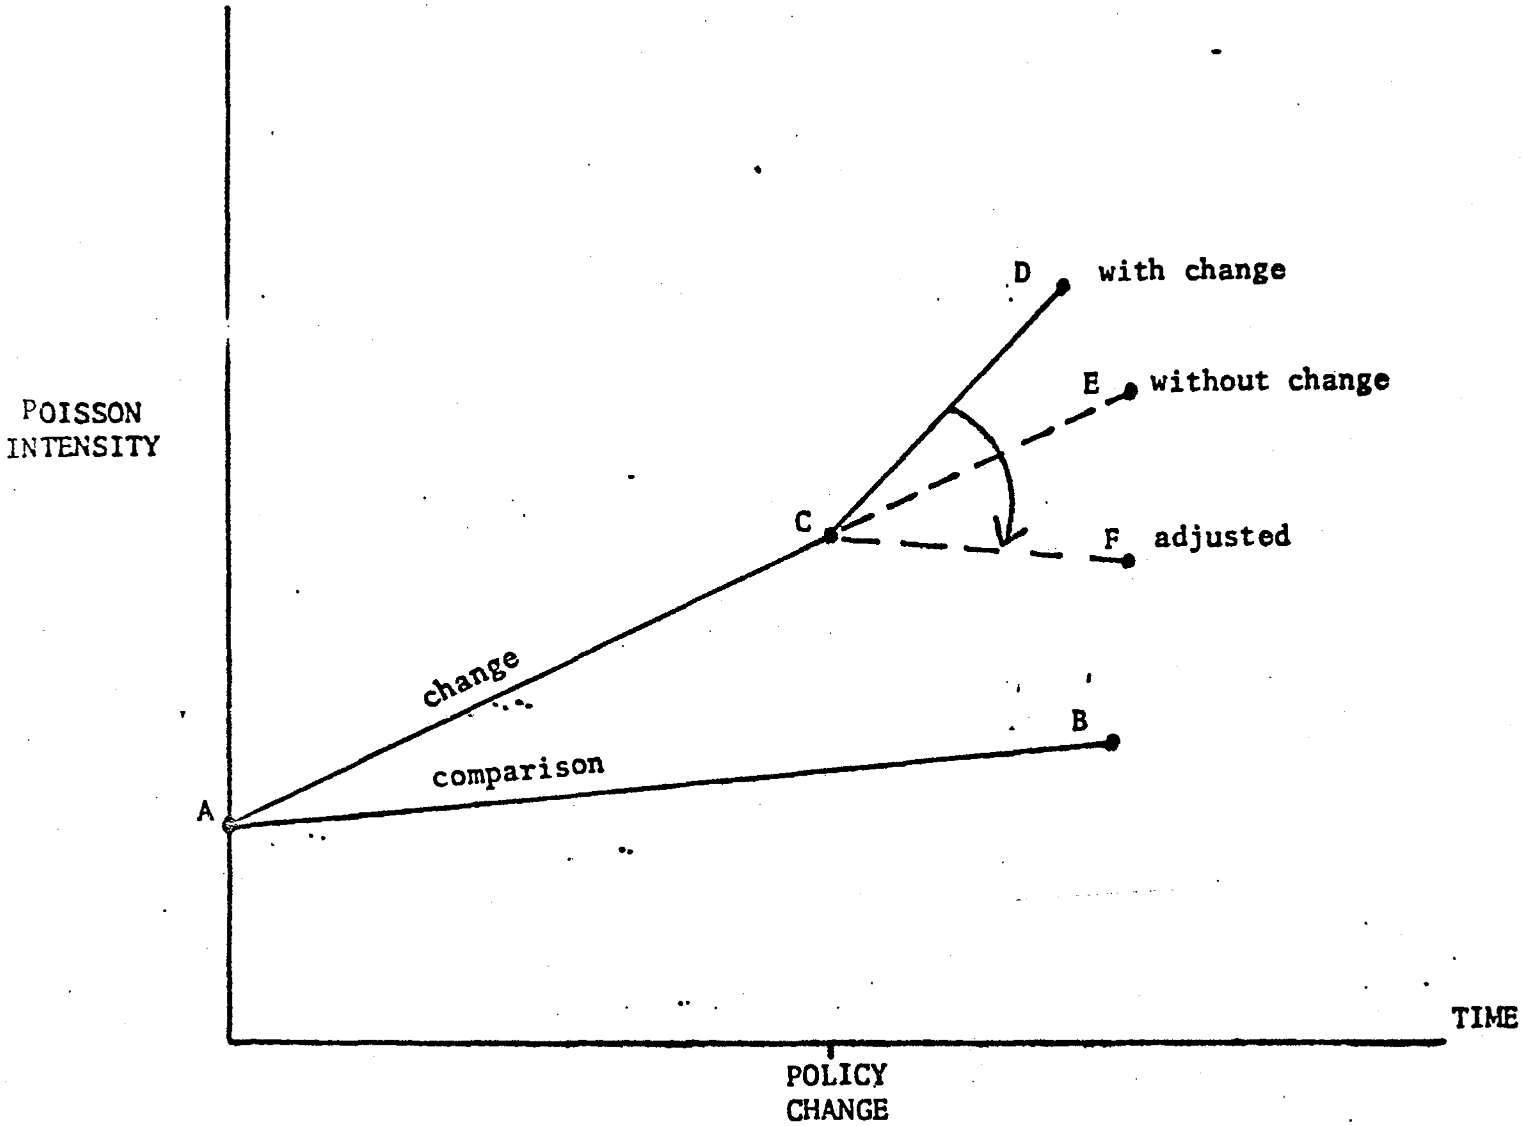

Figure 3: Illustrating how differential trend creates bias in estimate of the effect of policy change

REFERENCES

1. Bauer BM: Ordered Facts About Random Numbers. Keyboard (A Publication from Hewlett Packard), Sept/Oct 1981; pp. 10-11.

2. Bishop YM, Fienberg SE and Holland PW: Discrete Multivariate Analysis: Theory and Practice. Cambridge, MA: MIT Press, 1975.

3. Campbell DT and Stanley JC: Experimental and Quasi-Experimental Designs for Research. Chicago, IL: Rand McNally, 1963.

4. Haberman SJ: Log-linear Fit for Contingency Tables. Applied Statistics 1972; 21: 218-225.

5. Hartunian NS, Smart CN, Willemain TR and Zador PL: The Economics of Safety Deregulation: An Analysis of Excess Deaths and Costs Due to Motorcycle Helmet Law Repeal in the United States (submitted for publication), March 1982.

6. Knuth DE: The Art of Computer Programming, Vol. 2, Seminumerical Algorithms. Reading, MA: Addison-Wesley, 1969.

7. Larson RC and Odoni AR: Urban Operations Research. Englewood Cliffs, N.J.: Prentice-Hall, 1981.

8. U.S. Department of Transportation, National Highway Traffic Safety Administration: A Report to Congress on the Effect of Helmet Use Law Repeal -- the Case for Helmet Use. Washington, D.C.: DOT #HS-805-312, April 1980.

9. U.S. Department of Transportation, National Highway Traffic Safety Administration: Fatal Accident Reporting System: Unpublished data, 1975-1980.

10. Watson GS, Zador PL and Wilks A: The Repeal of Helmet Use Laws and Increased Motorcyclist Mortality in the United States, 1975-1978. American Journal of Public Health 1980; 70: 579-585.