The Pennsylvania State University

The Graduate School

College of Engineering

SAFETY EFFECTS OF PAVEMENT MARKING RETROREFLECTIVITY: AN APPLICATION OF CAUSAL BAYESIAN NETWORKS

A Thesis in

Civil Engineering

by

Vishesh Karwa

Submitted in Partial Fulfillment
of the Requirements
for the Degree of

Master of Science

August 2009

The thesis of Vishesh Karwa was reviewed and approved* by the following:

Eric T. Donnell
Assistant Professor of Civil Engineering
Thesis Advisor

Martin T. Pietrucha
Professor of Civil Engineering

Aleksandra B. Slavkovic
Assistant Professor of Statistics

William D. Burgos
Professor of Environmental Engineering
Professor in Charge of Graduate Programs

*Signatures are on file in the Graduate School

Abstract

Pavement markings provide a wide variety of information to drivers. For pavement markings to be seen by drivers at night, they must be retroreflective. Like many other roadway materials, pavement markings deteriorate over time and lose their ability to retroreflect headlamp illumination. There is general consensus among practitioners that this reduced performance may be a causative agent in the rate and severity of nighttime crashes, although previous research has not yet quantified the relationship. This work examines the relationship between safety and pavement marking retroreflectivity using data from selected highways in North Carolina and attempts to quantify the strength of the relationship between pavement marking retroreflectivity and safety. Safety is defined as the propensity of a segment to experience a crash during a given period of time. The framework of Causal Bayesian Networks is used to model the probabilistic dependence between pavement marking retroreflectivity and safety. Results indicate that the probability of a target crash increases by almost 87 percent with a decrease in retroreflectivity levels from above 350 mcd/m^2/lux to below 207.6 mcd/m^2/lux, for two- and four-lane roads.

Table of Contents

List of Figures

List of Tables

Acknowledgements

First and foremost I offer my sincerest gratitude to my adviser, Dr. Eric Donnell, who has supported me throughout my thesis with his guidance, patience and knowledge, at the same time, allowing me the room to work in my own. Be it guidance on general discussions, academic and research matters, or even selecting the best doughnut in town, I could always count on him. I attribute the successful completion of my Masters degree to his support, encouragement and effort and without him this thesis, too, would not have been completed or written. I simply could not have wished for a friendlier adviser.

I am also grateful to my committee members Dr. Martin Pietrucha and Dr. Aleksandra Slavkovic for their support. While doing work on a narrow problem using an abstract framework, it is easy to get lost in the nitty-gritty details, and lose the general picture. Insights into the broader picture, from Dr. Pietrucha, brought me back on line and helped make sense of what I was trying to accomplish. Refreshing ponderings with him on life, universe and everything helped me become a better person and brought excitement to the sometimes-monotonous research life. Dr. Slavkovic was my "statistical" sounding board, making sure that I was walking down the right path of analysis. It was in Dr. Slavkovic's class where I first got the idea of Bayesian Networks for safety analysis.

Being the last-minute-guy that I am (in)famous for, I offer my appreciation to the support from all my committee members to accommodate my requests in short notices of time; specially, to Dr. Pietrucha, for going out of the way for his ensuring presence in my defense.

I thank all my friends at PTI: Scott for all the movies, Lekshmi for a being patient ear to all my crazy ideas (actually, this one is for everyone!) and for all the brain storming sessions, Adam for entertaining us with "Muppet" songs and new nick names, Ken for always bearing the blame for my late arrival, Wen Hu for letting us make fun of her name, Ivette for making the lab environment lively, and Sarah and Donna for helping me out with my unusual and frequent computer needs.

I want to thank Judy, for ensuring that I am able to graduate on time. Without her, I would be a "paper-work" mess.

Lastly, but definitely not the least, I want to thank my brother, sister and father, from whom, in spite of being away for such a long time, I always receive unconditional love, care and support.

Chapter 1 Introduction

1.1 Background

An estimated 2.6 million people suffered some kind of transportation-related injury in 2006. About 99 percent of these injuries resulted from highway crashes (Transportation Statistics Annual Report, 2007). Transportation safety management aims to identify causes of such crashes, develop countermeasures to mitigate crashes, and evaluate the effectiveness of a safety countermeasure. Most of this analysis is carried out by developing statistical models that examine the correlation between traffic crashes and their potential causal factors. That correlation is not causation is a well-known maxim, yet it is believed that causal conclusions can be drawn from studying associations in carefully designed experiments or from observational studies by careful correction of bias.

Experimental study is a tool that is not always accessible to a transportation safety researcher. The types of data available in transportation safety studies are primarily observational, which makes it difficult to consistently estimate the causal effects of countermeasures. The aim of this thesis is to examine the applicability of causal inference methods to transportation safety studies. This is done by estimating the causal effect of pavement marking retroreflectivity on nighttime target crashes (target crashes are defined in Chapter 6) using a Causal Bayesian Networks framework.

1.2 Problem Statement

The risk of a fatal traffic crash increases significantly at night. Statistics indicate that the fatality rate doubles during the period of time between 9 p.m. and 6 a.m. (National Highway Traffic Safety Administration: Traffic Safety Facts 2001). This leads to the following question, "What engineering countermeasure(s) can be applied to highways and streets to reduce the ratio of night-to-day crashes?" Because the ratio of night-to-day traffic fatalities is 2:1, the specific causes for this unbalanced ratio must be found. It is hypothesized that low pavement marking visibility may be one cause of the increased rate of nighttime crashes. The objective of this thesis is to examine this hypothesis.

Pavement markings provide a wide variety of information to drivers. They are unique in terms of traffic control devices because drivers do not have to shift their attention away from the roadway to receive continuous information. The four main purposes for using road markings and symbols, as identified by Elvik and Vaa (2004) are to:

1. Direct traffic by indicating the path of the traffic on a road in relation to the surroundings;
2. Control traffic, for example, by reserving certain parts of the road for certain traffic groups (e.g., public transport), and by allowing or prohibiting overtaking and lane changing;
3. Warn road users about specific or hazardous conditions related to the road alignment; and
4. Supplement and reinforce information on traffic signs.

Pavement markings become one of the key methods of conveying this information to drivers at night, so their proper placement and maintenance are critical for safe driving

(Carlson et al., 2005). For pavement markings to be seen by drivers at night, they must be retroreflective. Retroreflectivity is a measure of an object's ability to reflect light back towards a light source along the same axis from which it strikes the object. In the case of retroreflective pavement markings, incoming light from vehicle headlamps is reflected back towards the headlamps and, more importantly, the driver's eyes. The retroreflective property of pavement markings is what makes the pavement markings visible to drivers at night. Pavement markings are made retroreflective by embedding glass beads in the marking material (sometimes called the binder material) as shown in Figure 1-1 (Hawkins et al., 2000). Rather than scattering light, as the pavement marking material would do without the glass beads, the beads refract the incoming light in such a way that it is returned back towards the vehicle's headlamps and driver's eyes. The most common measurement of retroreflectivity is the Coefficient of Retroreflected Luminance (R_L), which can be described as "the ratio of the luminance of a projected surface of retroreflective material to the normal illuminance at the surface on a plane normal to the incident light" (Austin and Schultz, 2002). Retroreflectivity measurements can be used to assess the efficiency of pavement markings in terms of their nighttime visibility on highways and streets.

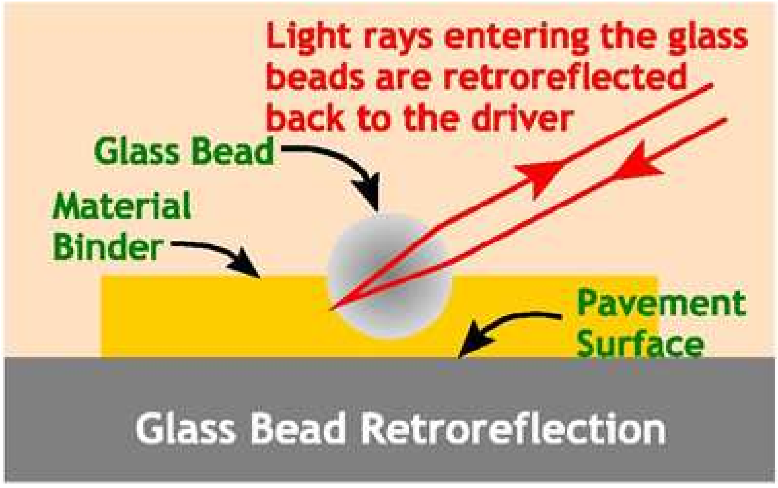

Figure 1-1: Pavement Marking Retroreflectivity Principles (Hawkins et al, 2000)

Pavement markings, like many other roadway materials, deteriorate over time. As pavement markings deteriorate, they lose their ability to retroreflect headlamp illumination. As a result, retroreflective measurements of pavement markings decrease over time. There is general agreement among the engineering community that this reduced performance may be a causative agent in the rate and severity of nighttime crashes, although previous research has not yet quantified the relationship.

While the *Manual on Uniform Traffic Control Devices* (MUTCD) requires that pavement markings be illuminated or retroreflective, it contains no minimum maintained retroreflective requirements (MUTCD, 2003). In 1992, Congress mandated that such standards for signs and pavement markings be developed, and research to develop these standards has been ongoing. It is anticipated that the minimum levels of pavement marking retroreflectivity proposed by DeBallion et al. (2007) will be included in future editions of the MUTCD (see Table 1, page 11); however, it is not yet clear how minimum levels of pavement marking retroreflectivity relate to safety. This thesis examines the relationship between safety and pavement marking retroreflectivity by applying the Causal Bayesian Network framework to model crash probabilities on highways in North Carolina.

.

Chapter 2 Overview

Before describing the extant literature and the analysis methods, it is crucial to understand the broad picture of the problem and the nature of its solution. As noted in the objectives, the purpose of this study is to determine the relationship between safety and pavement marking retroreflectivity. A study of the safety effects of pavement marking retroreflectivity is complicated due to the nature of pavement markings and the availability of data. The fact that pavement marking retroreflectivity (PMR) deteriorates over time, varies significantly by location, and can vary due to exposure to different environmental conditions, among other factors (that may also influence safety performance), makes safety analyses difficult. There are several issues associated with carrying out such a study -- they can be broadly put into the category of data and methodology, although these two aspects are inherently linked with each other and one guides the selection of the other.

2.1 Data

In any empirical study of this kind, the dataset is the most important part. Crash data and PMR data come from separate sources. Crashes are assigned to road segments, but the information about the PMR on that road segment at the time of a crash is generally not available. Instead, PMR data are collected at different time intervals (sometimes equally-spaced intervals, sometimes not) and different geographic locations. These data are generally spatio-temporally disjoint from the locations of crashes. To carry out a study of the effects of PMR on safety, these two databases from disparate sources must be combined. The problem is classical in the merging of geo-temporal databases. Figure 2-1 illustrates the situation. The X and Y axes denote the location of the segment and the Z axis denotes the time period. PMR levels must be determined at the location and time of occurrence of a crash, but the data for PMR are available at non-overlapping time periods.

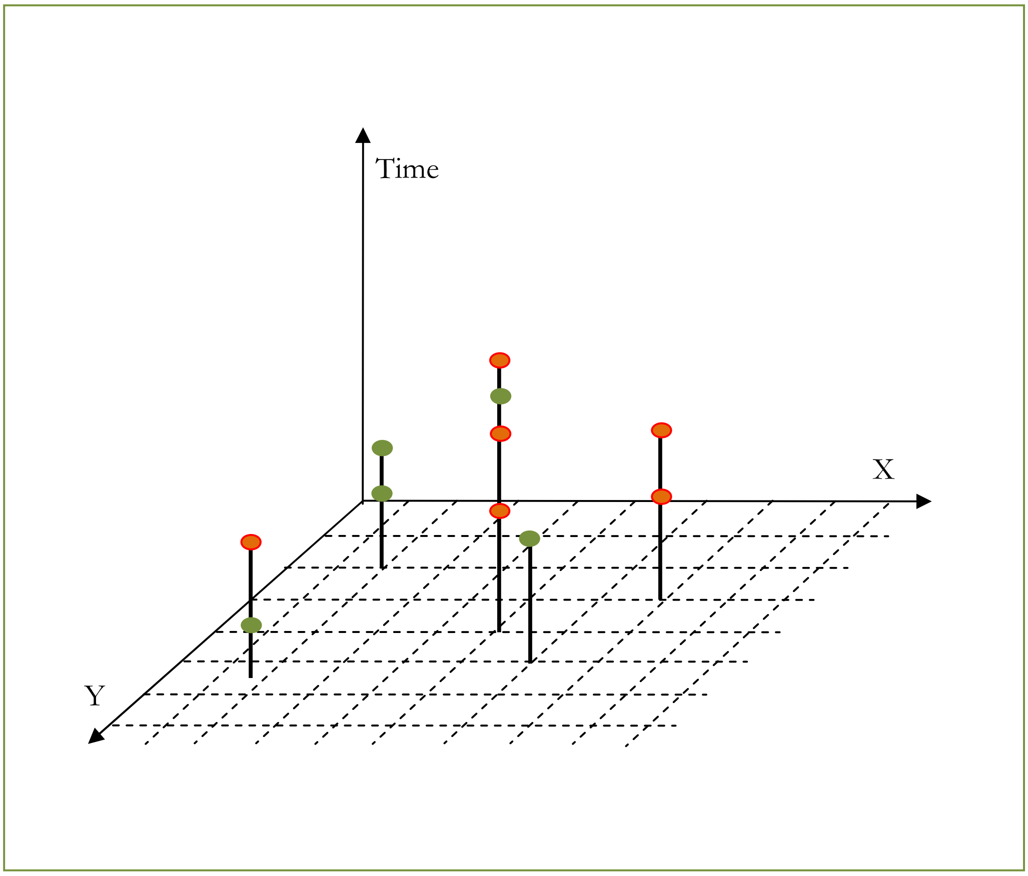

Figure 2-1: Illustration of the problem in combining spatio-temporally disjoint databases of accidents and PMR levels. The green dots represent the segments where PMR data are available and the red dots represent the segments where crashes have occurred.

Simple assignment or interpolation of measured PMR cannot solve the problem since a segment may contain a combination of line types. It is not clear how drivers use the information conveyed by yellow and white pavement markings during driving. Do the drivers use the information from both lines at the same time to navigate through the segment? Or, are the white edgelines sufficient for navigation and the yellow lines used only to understand passing regulations? Moreover, the type of crashes that are influenced by different markings on the road is not known with certainty. For instance, it is believed that better visibility of right edgelines may be helpful in reducing single- vehicle, run-off-the-road to the right crashes on a curve whereas the visibility of yellow centerlines may be helpful in reducing head-on or side swipe collisions on undivided roadways. Similarly, improved visibility of yellow edgelines on divided highways may be helpful in reducing single-vehicle, run-off-the-road crashes to the left of the travel lanes (i.e., median encroachments). Some of these questions are in the realm of human factors and may not be answered by an empirical study of this kind, but insights into the nature of the answers could be obtained. These issues are further discussed in section 3.3; the purpose of introducing them here is to highlight the problems associated with simple assignment of PMR levels to road segments.

The most common solution to the above problem has been to seek insights into the pavement marking degradation phenomenon and to develop models of this degradation process. These models are then used to predict what the retroreflectivity would have been on a segment where a crash has occurred. The process of PMR degradation is also of interest to gain knowledge about the service life of a pavement marking. Service life is defined as the amount of time required for a pavement marking to reach the end of its useful lifetime. The criterion for deciding the end of the useful life of a marking is decided based on the minimum level of retroreflectivity that is required to be maintained for safe driving. As noted previously, estimating this level is not easy and is a subject of considerable research and debate. The current study focuses on using pavement marking degradation models as a tool to combine the crash and PMR databases. A myriad of methods are available to solve this problem using both techniques, and a literature review of the methods is warranted.

2.2 Methodology Review

The question of causation requires a different framework of analysis when compared to predictive or exploratory work. Estimation of causal effects from empirical data is subject to a number of assumptions, some of which cannot generally be tested from observational data. A review of methods used for causal inference in the statistical literature in general and in the safety literature in particular, is required to understand different approaches used for causal modeling.

A review of the methods used to estimate pavement marking degradation is also needed to combine the crash and the PMR databases. A review of the methods used to explore the relationship between safety and PMR levels is also required. This step depends on the definition of safety which could be defined to include the number of crashes, and/or the severity of crashes. The types of crashes included in the study are also important. In this study safety is defined as the change in target crash probability due to a change in PMR levels. The formal definition of safety and target crashes is given in Chapter 7.

Chapter 3 Literature Review

3.1 Nighttime Visibility and Pavement Marking Retroreflectivity

The visibility of a pavement marking at night is dependent on retroreflectivity, which represents the portion of light from a vehicle's headlampes reflected back toward the eye of the driver of that same vehicle as discussed by Migletz et al. (2004). Retroreflection means that the light is reflected back at the same angle that it is projected. If the light from headlights was to be perfectly retroreflective, it would not reach the driver's eye's which are above the headlights. Since retroreflection is imperfect, some of the light reaches the driver's eyes, which are above the headlights, increasing the contrast between the delineator and the low-reflectance pavement background. The higher the percentage of light that is retroreflected, the greater the contrast and the further away the delineator will be seen. Both the retroreflectivity value (represented by the coefficient of retroreflected luminance, R_L, in millicandelas per square meter per lux (mcd/m^2/lux)) and the degree of contrast between the pavement marking and adjacent pavement surface are important to the visibility of a pavement marking (Migletz et al, 2004). Results from a study by Loetterle et al. (2004) indicated that the minimum acceptable level of retroreflectivity ranges from 80 to 120 mcd/m^2/lux. Recent research by DeBallion et al. (2007) have proposed minimum threshold retroreflectivity levels shown in Table 1; however, these levels have not yet been adopted by the FHWA for inclusion in the Manual on Uniform Traffic Control Devices (2003).

Table 1 Proposed Minimum Pavement Marking Retroreflectivity Levels by Deballion et al. (2007).

Roadway Marking Configuration	Without RRPM's			With
	≤ 50 mph	55-65 mph	≥ 70 mph	RRPM's
Fully marked roadways (with centerline, lane lines, and/or edgeline as needed)	40	60	90	40
Roadways with centerlines only	90	250	575	50

Notes:
All retoreflectivity values are reported in units of mcd/m²/lux.
RRPM's are retroreflective raised pavement markers.

The retroreflectivity of pavement markings is an important factor when determining driver detection distances at night. Parker and Meja (2003) compared objective measures (retroreflectivity measurements) with subjective evaluations (ratings by evaluators), and found a high correlation between retroreflectivity and visibility ratings. According to Parker and Meja (2003), the relationship between the objective and subjective measures is non-linear, as shown in Figure 3-1.

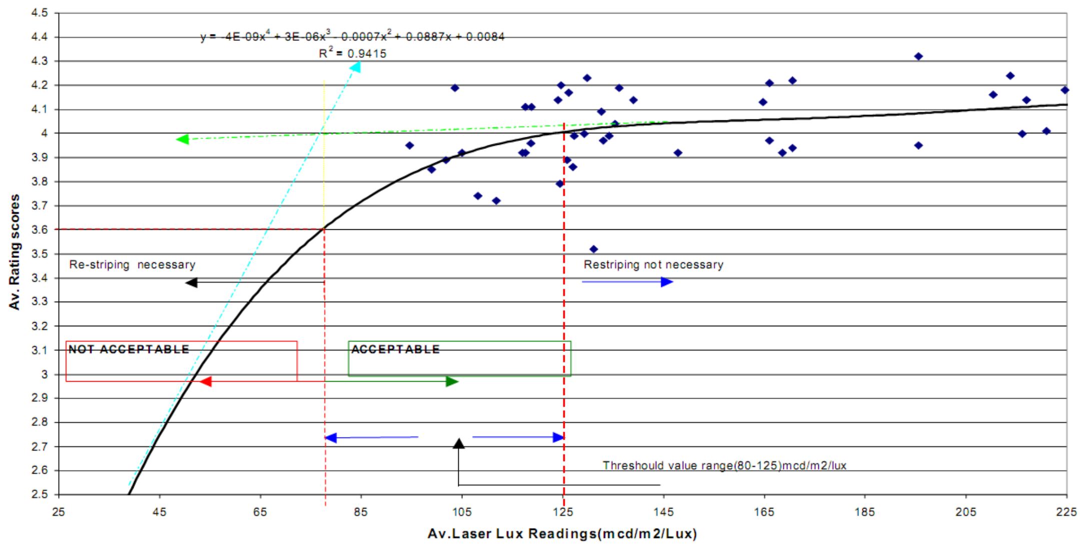

Figure 3-1: Non Linear Relationship between PMR and ratings (Parker and Meja, 2003)

A non-linear relationship between retroreflectivity and visibility ratings has also been established by Loetterle et al. (2004), as shown in Figure 3-2 using a similar 5-point rating scale.

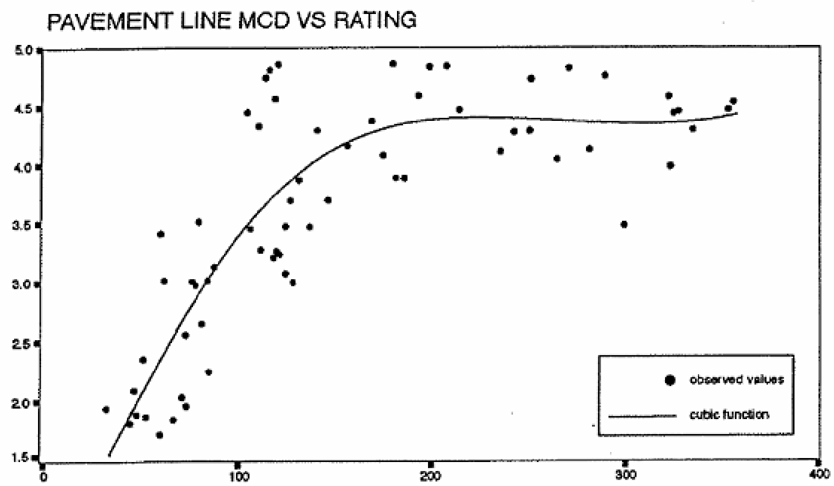

Figure 3-2: Non Linear relationship between PMR and ratings (Loetterle et al, 2004)

3.2 Pavement Marking Degradation Modeling

As noted previously, the purpose of using PMR degradation models in this study is to combine the crash and the PMR databases. A variety of pavement marking degradation models have been estimated using data from roadways in the U.S. Pavement marking retroreflectivity data can be collected using either handheld or mobile devices -- the current standard is to collect such data with a retroreflectometer that has a 30 m (98 ft) geometry (ASTM, 2005). Several modeling methods have been used to analyze retroreflectivity data based on the current measurement geometry, including time series, linear regression with first- or second-order time (age) effects, parametric survival analysis, and artificial neural networks.

Thamizharasan et al. (2003) found three basic patterns in terms of degradation of PMR levels over time as shown in Figure 3-3 below. In Figure 3-3A, the pavement marking retroreflectivity for newly placed markings increases initially. This is because after some amount of wear, the glass beads become exposed and reflect more light. The PMR then reaches a peak before starting to decline over time due to the loss of glass beads in the binder material. Figure 3-3B shows the decline in PMR levels of older markings (older than 300 days) and Figure 3-3C demonstrates how the retroreflectivity of pavement markings change due to remarking and snow plowing activities.

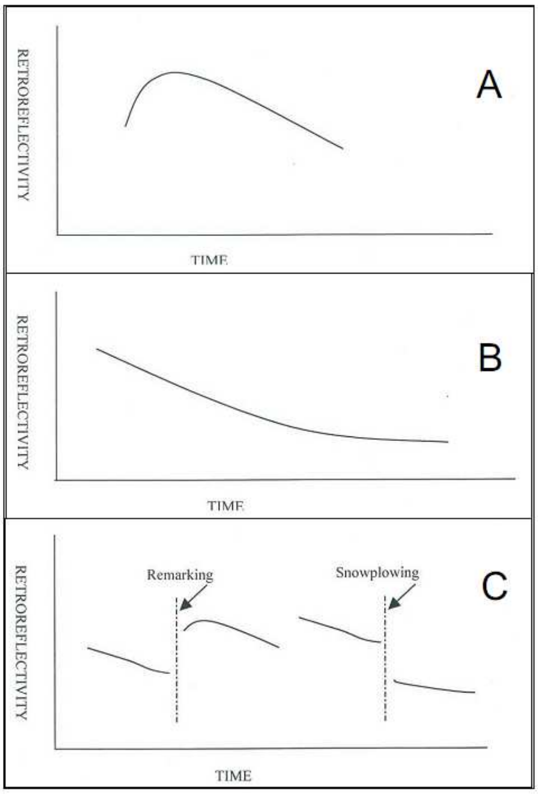

Figure 3-3: PMR degradation pattern of new (A), established (B), and snow plowing and remarking (C) (Thamizharasan et al, 2003)

Zhang and Wu (2006) used data from the Mississippi National Transportation Product Evaluation Program (NTPEP) test deck to estimate the service life of durable tapes, three-year waterborne paints, preformed thermoplastic and thermoplastic pavement markings. An autoregressive integrated moving average time series approach was recommended to model pavement marking retroreflectivity as a function of time (age in months). Using the first 18 months of retroreflectivity data, the models were used to forecast the retroreflectivity at 21 and 24 months. The difference in the measured and predicted values was less than 10 percent for all pavement marking types at 21 months and less than 15 percent for three of four pavement marking types at 24 months (predicted values were 31 percent different than measured values for durable tapes at 24 months). As noted by the authors, the models did not consider pavement surface type (asphalt or concrete), pavement marking color (yellow or white), or other factors that may be associated with pavement marking retroreflectivity degradation. Additionally, the data collected and used in the analysis were from transverse pavement markings applied at a test deck location and may not degrade in a similar manner as longitudinal pavement markings intended to delineate the travel lanes along a highway.

Bahar et al. (2006) developed pavement marking degradation models using inverse polynomial regression. Although the main objective of the research was to determine the relationship between pavement marking retroreflectivity and safety, separate models of pavement marking retroreflectivity were specified for various marking types, marking color, climatic region, and snow removal activities. These models were estimated using data from various NTPEP test decks across the U.S. The inverse polynomial model contained only age (in months) in the specification (age was included as a first- and second-order term) suggesting that pavement marking retroreflectivity degradation is non-linear. It was determined that pavement surface type and average daily traffic (ADT) volume were not associated with pavement marking retroreflectivity. The goodness-of-fit of the models were examined graphically using cumulative residuals (CURE) plots. A limitation of the data was that all pavement marking retroreflectivity measurements were taken on transverse pavement markings on NTPEP test decks.

Sarasua et al. (2003) used a combination of non-linear (polynomial) and linear regression methods to model the degradation of epoxy, thermoplastic, and tape pavement markings on South Carolina Interstates. A non-linear method was used to model the initial increase in pavement marking retroreflectivity until reaching a state of steady (linear) degradation. The initial increase in retroreflectivity after applying the pavement markings was likely due to the delayed exposure of glass beads. A linear regression model was then specified for the remainder of the pavement marking life cycle after reaching the state of steady degradation. Time (in days) was used as the explanatory variable in the models. The goodness-of-fit for the models ranged from 20 to 80 percent. It was determined that pavement marking type, pavement surface type, and the frequency of pavement surface maintenance activities were the variables associated with pavement marking retroreflectivity. Traffic volume and pavement marking application temperature and humidity were not associated with pavement marking retroreflectivity.

Kopf (2004) collected pavement marking retroreflectivity data using a mobile retroreflectometer along roadways with longitudinal markings in Washington State. Separate models of pavement marking retroreflectivity were specified for different regions of the state to control for environmental effects, and based on different pavement marking colors (yellow or white) and traffic volume levels. Best-fit trendlines were used to determine the degradation phenomenon of the waterborne and solvent-based pavement markings. The independent variable in all models was time (in days). In some instances, the best-fit trend line included a logarithmic or exponential transformation of the independent variable while in other cases the best-fit trend line was a first-order linear regression model. The goodness-of-fit (R^2), for the various models specified, ranged from 3 to 69 percent.

Migletz et al. (2001) considered first-order linear regression, second-order linear regression, and exponential decay models to estimate the degradation phenomenon of durable pavement markings collected in 19 states using a mobile retroreflectometer. Cumulative traffic passages (traffic volume multiplied by time) was the independent variable used in the models. It was determined that 67 percent of the pavement markings included in the sample exhibited a degradation phenomenon that was consistent with a first-order linear regression model and that 25 percent of markings exhibited an exponential decay. Only two percent of the pavement marking data exhibited a second-order degradation pattern while the remaining pavement markings could not be fit to the data. Lindly and Wijesundera (2003) used a similar approach to model the degradation pattern of profiled pavement markings in Alabama. Data were collected using a mobile retroreflectometer along longitudinal pavement markings. They concluded that a linear model and exponential decay model produced nearly identical goodness-of-fit to the pavement marking retroreflectivity data.

Sathyanarayanan et al. (2007) analyzed pavement marking retroreflectivity data from the Pennsylvania NTPEP test deck using parametric (Weibull) duration models. The modeling approach assumed that pavement markings reached the end of their useful service life when the retroreflectivity fell below 100 mcd/m^2/lux. Separate models were specified for pavement markings applied on asphalt and concrete pavements, white and yellow pavement markings, and at locations in the wheel path and skip line areas of the transverse NTPEP test deck markings. It was determined that pavement markings in the wheel path area of the NTPEP test deck exhibited greater degradation rates than the same markings measured along the skip line area. Furthermore, it was determined that white pavement markings have longer service lives than yellow pavement markings. The difference in degradation between markings applied on asphalt and concrete was nominal.

Karwa and Donnell (2009) proposed the use of layered feedforward artificial neural networks (ANN) to model pavement marking retroreflectivity. Separate degradation patterns were estimated for white edgeline, yellow edgeline, white skip line, and yellow centerline thermoplastic pavement markings in North Carolina. The results were compared to linear regression and panel data models. The authors concluded that the ANN provided a better fit to the pavement marking retroreflectivity data based on the coefficient of determination (R^2) and root-mean square error (RMSE).

Based on the existing literature, a variety of modeling methods have been used to predict the degradation phenomenon of pavement markings. The pavement marking type and color, pavement surface type, traffic volume, snow removal activities, and other spatial and temporal variables have all been shown by some researchers to be associated with pavement marking retroreflectivity; however, the variable most strongly associated with pavement marking degradation is time. The most commonly used modeling technique used to predict pavement marking retroreflectivity has been linear regression, but alternative modeling methods have been shown to better explain the variability in the retroreflectivity data.

3.3 Human Factors

Given the tendency for drivers to drive too fast at night under low-visibility conditions and "overdrive" low-beam headlights (Migletz et al., 1994; Zwahlen and Schnell, 2004), Migletz et al. (1994) defined the preview or visibility distance as the distance that the delineation provides the driver to see upcoming changes in the roadway. This distance must provide drivers with enough time to detect roadway features and changes in alignment ahead, and respond with steering and speed adjustments. The preview distance offered by pavement markings is particularly important when the view of the road ahead is limited and drivers are forced to depend on roadway and traffic information that is visible only from a short distance (Migletz and Graham, 2002). In adverse weather conditions or at night, this preview distance is dependent on the visibility of the pavement markings, which in turn is a function of the retroreflectivity level of the markings.

Many previous studies have attempted to determine driver requirements in terms of preview time for both short- and long-range guidance (Migletz and Graham, 2002). The consensus is that 2 seconds of preview time is required for short-range guidance in extreme situations. Examples of extreme situations include changes in alignment or when drivers may be required to respond quickly to perceived hazards. For long-range guidance, the general view is that a minimum of 3 seconds of preview time is required to allow for comfortable driving (Migletz and Graham, 2002). Zwahlen and Schnell (2004) investigated this concept further and recommended a preview time of 3.65 seconds (3.00 seconds of true preview time and 0.65 seconds of perception-reaction time) to accommodate drivers with a margin of safety and comfort. Consistent with the research conducted by Molino et al. (2004), Zwahlen and Schnell (2004) also showed that requirements for preview times could be substantially relaxed if raised pavement markers were used along the centerline of an undivided roadway or along lane lines on divided highways. In general, the concept of using a preview time implies a static view of the driving task rather than an adaptive one. Most driving simulator experiments, and even most field studies, assume a constant speed and use this constant speed as a base for preview time calculations. However, on the highway, drivers may change their speed as a function of visibility and road conditions and therefore may not maintain a constant speed.

Parker and Meja (2003) found that driver age has a significant impact on the visibility (which was quantified using detection distances) of pavement markings at night. The field study found that older drivers (≥ 55 years) tended to assign lower scores to pavement markings compared with younger drivers (<55 years). This observation was expected given that contrast sensitivity is likely to diminish with age. However, it has also been shown that older drivers may sometimes assign higher subjective visibility ratings to pavement markings simply because they are aware of their visual limitations and have lower expectations in regards to the brightness or visibility of the pavement markings (Loetterle et al. 2004).

Regardless, pavement markings, signage, and other road features may not provide adequate nighttime visibility for drivers of all ages. In some cases, drivers 65 years old and over may need as much as four times greater contrast to see as well as a 39-year-old driver (Migletz and Graham, 2002).

3.3.1 Driver Perception of Pavement Markings

Drivers scan the visual scene using an efficient set of eye movements called saccades and fixations (Shinar, 2007). The eyes move in a synchronized manner in a series of jumps which are saccades. The saccades are separated by short periods of stops called fixation. The information from the visual environment is extracted during the fixations (Velichkovsky et al., 2002).

A considerable amount of research has been devoted to describing driver eye scanning behavior and determining the process by which drivers fixate their eyes when driving (Macworth and Morandi, 1967; Victor, 2000). Figure 3-4 shows the points of driver's eye fixations (Klebelsberg, 1982) and Figure 3-5 represents relative duration of concentration in different areas (Lobanov, 1980).

The points in Figure 3-4 represent driver eye fixation for the following purposes:

1 and 6—observation of situation on the left and right adjacent traffic lanes;

2 and 8—control of vehicle position, relative to the left and right lane edges;

3—observation of pavement quality;

4—observation of a leading vehicle;

5—observation of road signs;

7—visual field center of gravity.

As shown in Figure 3-5, drivers spend around 14 percent (6.9 percent + 7.2 percent) of total time for conscious estimation of the vehicle's position on the roadway. On roadways with high traffic volumes, such estimation could take up to 20 percent of the time on tangent highway sections and up to 25 percent on horizontal curve sections (Klebelsberg, 1982; Lobanov, 1980).

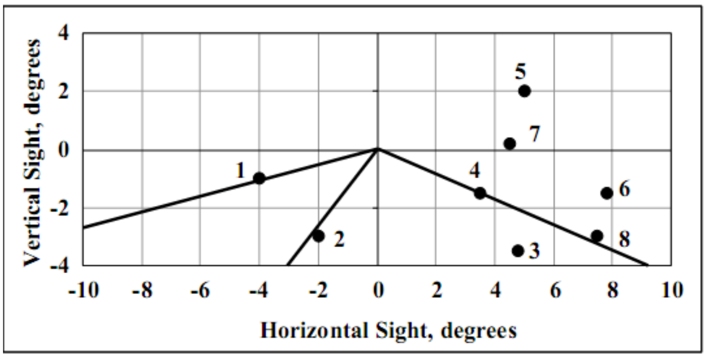

Figure 3-4: The points of driver's eye fixations (Klebelsberg, 1982)

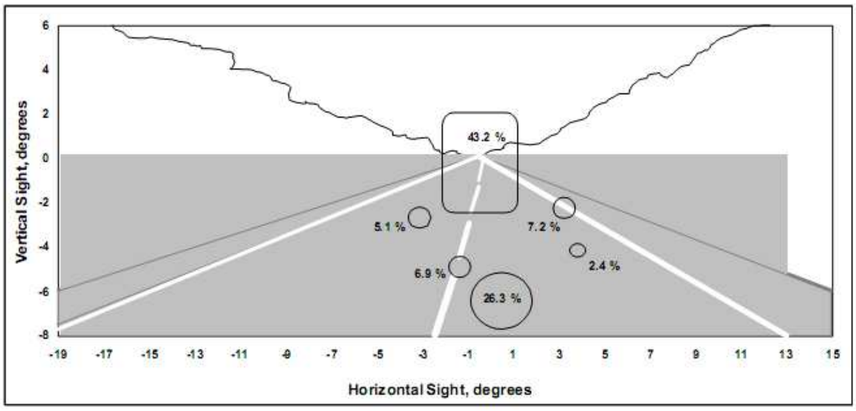

Figure 3-5: Distribution of Driver's Attention on Two-Lane Rural Roads (Lobanov, 1980)

3.3.2 Driver Behavior on Curves

Dewar and Olson (2002) state, "The study of vehicle operations on horizontal curves has shown that speeds reduce on curves and that encroachment onto the edgeline occurs on right curves and onto the centerline on left curves." Johnston (1983) has provided quantitative evidence that optimal driving behavior does not mean driving consistently within the center of the lane. Instead, what drivers actually do, and what, according to Johnston (1983), should be considered proper driving, is to 'cut the corner' without departing from the lane (Figure 3-6). In Figure 3-6, the radius of the curve driven is larger than the geometric curvature of the road.

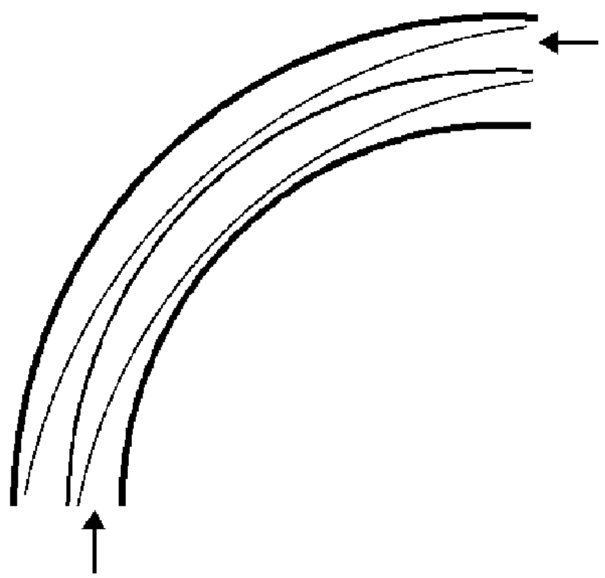

Figure 3-6: 90-degree curve showing the trajectory drivers actually drive (Johnston, 1983)

3.3.3 Edgelines and Speed

Research by Zwahlen and Schnell (1995) has shown that pavement marking retroreflectivity requirements increase exponentially with an increasing preview distance. Using a constant preview time of 3.65 seconds and 2.00 seconds for roadways without markers and with markers, respectively, Zwahlen and Schnell (1995) established the relationship between minimum retroreflectivity requirements and vehicle speeds for a 62-

year-old driver (accounting for 95 percent of the nighttime driver population in the United States), as illustrated in Figure 3-7 below.

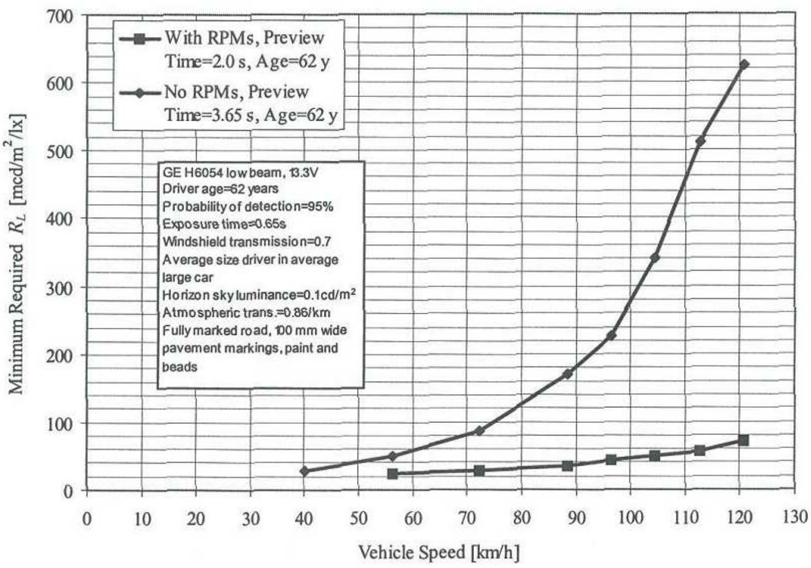

Figure 3-7: Retroreflectivity requirements by vehicle speed (Zwahlen and Schnell, 1995)

In general, studies that have recorded speed measures before and after the applications of markings and markers have reported a modest increase in operating speed immediately after installation. For reasons that are not clear, this speed change fades over time. For example, Willis et al. (1984) collected daytime speed measures of traffic before-and-after the installation of both continuous and broken edgelines, and observed an average 1.9-km/h increase in speed one month after installation. However, 11 months later, the speed change was only 0.5-km/h higher when compared to the before condition. A meta-analysis conducted by van Driel et al. (2004) shows that adding an edge line or a centerline to a road that was previously unmarked leads to an increase in vehicle speed and a shift of the lateral position towards the edge of the road. However, the authors did not find any increase in the speed when an edge line was added to a road that was already marked with a centerline or in the case where a center line was replaced with an edge line.

Based on the literature review, several human factors have been shown to affect the visibility of pavement markings as well as operational variables. In particular, the driver speed and lateral vehicle position depends on the visibility of PMR as well as the radius of curvature. To account for the changes in speed with changes in visibility, the concept of preview time is required. However, the actual combined effects of human factors and PMR levels on crashes are not known due to difficulties associated with data collection. The human factors literature review is used to illustrate how causal diagrams can be used to incorporate prior knowledge in the form of qualitative relations that are shown in Appendix D.

3.4 Safety Modeling of Pavement Marking Retroreflectivity

The existing literature related to the safety effects of pavement marking retroreflectivity contains a variety of analytical methods with different purposes. Analysis methods vary from simple before-after observational studies to empirical Bayes before-after observational studies to cross-sectional studies.

Masaeid and Sinha (1994) investigated the safety effects of pavement marking improvements applied to 100 randomly-sampled, rural, undivided highways in Indiana using a Bayesian before-after study methodology. The results indicated that, when improving the pavement markings, total crashes increased by 3.4 percent; however, the results were not statistically significant. When considering only those sites that had a higher than average crash frequency in the before period (36 sites in sample) – a statistically significant reduction of 13.5 percent in the total crash frequency was found. The authors noted that the most likely reason for this finding was that the hazardous locations had a higher than average number of crashes, and the reduction in total crashes was likely attributable to the improved visibility of pavement markings.

Abboud and Bowman (2002) attempted to develop a relationship between crashes and pavement marking retroreflectivity levels for paint and thermoplastic materials in Alabama. A critical crash rate for white paint and thermoplastic lines was calculated and compared to the crash rate on over 950 miles of two-lane and four-lane highways. The critical crash rate was defined as the maximum allowable crash rate that corresponds to the pavement marking retroreflectivity threshold beyond which re-striping is warranted. This rate was determined as 0.220 and 0.103 crashes per million-vehicle miles travel for white paint and thermoplastic lines, respectively. The results indicated that a retroreflectivity threshold value of 140 – 156 mcd/m^2/lx is the minimum acceptable value if the crash rate is to be maintained below the critical rate. A limitation of the study is that only white pavement markings were considered and the use of crash rates in the analysis assumes that the relationship between crashes and traffic volume is linear.

Lee et al. (1998) attempted to develop a correlation between pavement marking visibility and nighttime crashes using data from four different geographic areas in Michigan. Test sites contained longitudinal pavement markings with waterborne paint, thermoplastic, polyester, and tapes. The following crash characteristics were used to identify events that were associated with nighttime visibility crashes:

1. None, other, and unknown crash contributory factors;

2. Crash types such as miscellaneous, overturning vehicle, fixed object collision, other object collision, and head-on crash;
3. Dawn, dusk, and dark conditions; and,
4. Highway area type of non-intersection and non-interchange.

Linear regression lines showed that no statistically significant relationship existed between nighttime crash rates and pavement marking retroreflectivity.

Migletz et al. (2000) studied 55 locations in 19 states to evaluate various all-weather pavement marking types. Pavement marking re-striping occurred once the pavement marking retroreflectivity level reached its end-of-service life, defined as the time when the retroreflectivity level fell below a minimum level. A before-after evaluation, using both a paired sign evaluation and a yoked comparison, was used to determine the effects of pavement marking retroreflectivity on safety. The "after" period was defined as the time when pavement marking re-striping occurred. Based on the paired sign test, the results indicated that daytime crash rates, under dry pavement conditions, increased at 53 percent of the sites and decreased at 47 percent of the sites. A similar conclusion was reached for the nighttime, dry pavement condition. Daytime crash rates under wet pavement conditions increased at 40 percent of sites, decreased at 40 percent of sites, and remained the same at 20 percent of sites. Nighttime crash rates under wet pavement conditions increased at 45 percent of sites, decreased at 29 percent of sites, and remained the same at 25 percent of sites. It was not possible to draw statistically significant conclusions based on the paired sign test. The yoked comparison analysis found that, under dry pavement conditions, the crash frequency was expected to decrease from 1 to 20 percent. Under wet pavement conditions, the crash frequency was expected to be between a 5 percent decrease and a 40 percent increase. It is difficult to conclude, based on the findings from the all-weather pavement marking evaluation, that there is a consistent link between pavement marking retroreflectivity and safety under various time-of-day and weather conditions. Further, the degradation phenomenon of pavement markings was not considered in the evaluation.

Cottrell and Hanson (2001) conducted a before-after study of various pavement marking types in Virginia to determine when each type of material should be used on Interstate and other primary roadways. Roadway segments that were remarked using a different pavement marking material type were included in the "treatment" group while those that were remarked using the same material were included in the "comparison" group. Target crashes (i.e., sideswipe in the same direction and run-off-road) that typically occur as a result of poor delineation were used as the performance measure. Nighttime target crashes were considered "treatment" crashes and daytime crashes were considered "comparison" crashes. The mean number of target crashes occurring during nighttime conditions (i.e., treatment) was 3.175 crashes per year. The mean number of target crashes occurring during daytime conditions (i.e., comparison) was 4.525 crashes per year. The results of the analysis were inconclusive as some sites exhibited a crash increase (15 to 17 percent) while others exhibited a decrease in the crash frequency (18 to 78 percent).

Bahar, et al. (2006) evaluated the safety effects of pavement markings using retroreflectivity data from the National Transportation Product Evaluation Program (NTPEP, 2003). The expected number of monthly non-daylight, non-intersection crashes were modeled using eight years of HSIS data from California for multi-lane freeways, multi-lane highways, and two-lane highways. This estimate was then adjusted using seasonal and retroreflectivity multipliers. The data analysis indicated that there is no relationship between crash counts and pavement marking/marker retroreflectivity degradation.

In summary, several before-after observational studies have attempted to determine the safety effects of improved pavement marking visibility. These studies have generally been inconclusive and lacked control for the degradation phenomenon of pavement markings over time.

Chapter 4 Causal Inference

4.1 Goal of Causal Inference

In a causal inference, one reasons to the conclusion that something is, or is likely to be, the cause of something else. After such a conclusion has been reached, one tries to measure the effect size of the cause. Questions in transportation safety are of similar nature. For instance, what is the effect of increased PMR on safety? An effect size of 0 would imply the lack of a causal relationship. In many other disciplines, causal knowledge of this kind comes from experiments. The key aspect of an experiment is that one can deliberately change the value of only a few factors while either keeping all other factors constant, or neutralizing the effect of all other factors through randomization. In road safety, experimentation of this kind is rare and, in some cases, unethical. As such, knowledge has to be extracted from 'observational studies'; that is, by interpreting data that are retrospective. In road safety, observational data can be found in police crash reports, or in electronic roadway inventory and crash databases maintained by transportation agencies. Several limitations exist in using observational data, including that the cause of a crash is often subjective and a threshold (monetary, occupant injury or vehicle damage) is used to determine if a crash is reportable.

Irrespective of the type of data used for causal inference, the definition of cause and causal effect must be formalized to facilitate the use of statistical analysis methods.

4.2 What is a Cause?

The area of casual inference has received a lot of attention from philosophers, statisticians, computer scientists, and experts from many other fields. The debate on what is considered to be a cause is of philosophical origin, and a review of two competing theories is helpful for the purpose of this analysis.

Broadly speaking, the two approaches can be labeled as difference-making and causal process theories. The former rely on the guiding idea that causes must make a difference to their effects, in comparison with some appropriately chosen alternative. This difference-making can be explicated mathematically in a variety of ways. Probabilistic theories attempt to do this in terms of inequalities among conditional probabilities. A cause must raise or at least change the probability of its effect, conditional on a suitable set of background conditions. This attempt to define causation has affinities with association and the use of contingence information as a measure of causal strength (Dickinson and Shanks, 1995). Counterfactual theories explicate difference-making in terms of counterfactuals: A simple version might hold that a cause C causes an effect E if and only if it is true that both:

a) If C were to occur, then E would occur and,

b) If C were not to occur, then E would not occur.

Counterfactuals are often understood in terms of relationships among possible worlds. Roughly, a counterfactual like (a) is true if and only if there is a world in which C and E hold that is "closer" or "more similar" to the actual world than any possible world in which C holds and E does not hold.

The theory used in the present analyses, called the interventionist or the manipulative account of causation, is a version of the counterfactual theory. The counterfactuals in question describe what would happen to E under interventions (or idealized manipulations) on C. The interventionist theory does not require (although it permits) thinking in terms of counterfactuals or in terms of possible worlds.

It is important to distinguish between the so-called type causal claims that relate one type of event to another (e.g., "speeding causes crashes") with token or singular causal claims that relate particular events ("Jane ended up in a crash because she was speeding"). It must be pointed out that the statement "speeding causes crashes" does not mean that all cases of events in speeding will result in crashes, which matches intuition. The probabilistic definition of cause is mathematically consistent with our intuition; speeding increases the probability or chances of occurrence of a crash. If one were to believe that speeding has no effect on the chances of occurrence of a crash, then one would claim that speeding does not cause crashes.

There are several modeling frameworks for performing statistical causal inference. One of the most popular methods of statistical causal inference is the Potential Outcomes Framework proposed by Neyman, Rubin, and others (Neyman, 1923; Holland, 1986; Rubin, 1974, 2002, 2006; Rosenbaum 2002). This method uses the counterfactual definition of causation. The set of outcomes of application of treatment and control variables to the same unit are used as fundamental quantities to derive casual effects. Other approaches of causal inference are sufficient-component cause models (Rothman, 1976), structural equations models (Sobel, 2000; Kaufman and Kaufman, 2001), selection models and instrumental variables (Heckman et.al., 1998; Heckman and Vytlacil, 2003), and causal diagrams (Sprites et. al., 1993; Pearl, 1995; Pearl, 2000). For an overview of different causal inference methods, see Greenland and Brumback (2002).

Pearl (2000) recently complied and proposed a rigorous mathematical framework that spans both the counterfactual and manipulative definition of causation and has the ability to handle both token and type causal claims. In this method, causal information is represented by directed acyclic graphs. This work uses Pearl's Causal Diagram framework to estimate causal effects. Chapter 5 introduces Bayesian Networks and Pearl's Manipulative theory of Causation.

4.3 Statistical Causal Inference

The gold standard for statistical causal inference is a randomized study. In such studies units are assigned to the treatment and comparison groups randomly. This random assignment ensures probabilistic equivalence which means that any differences between the treatment and comparison groups is due to chance. In other words, after randomized assignment, the groups are equivalent in probability. If the randomization is ensured, it is known that the two groups (in probability) were the same before treatment and that any difference in the outcomes is due to the treatment, and most likely not due to differences between two groups due to the assignment process.

In an observational study, data are collected passively without any control over treatment assignment. Thus the treatment assignment could be (and mostly is) related to the outcome itself, creating a bias in estimating the causal effect. Thus, it is important to model the treatment assignment process when carrying out causal inference from observational data, and this is one of the cornerstones of methods that perform causal inference from observational data.

According to Rubin (2008), it is advisable to conceptualize the dataset obtained by passive observations as being generated by a randomized experiment where the rules of assigning treatment are not known and must be estimated from the data. Also, the assignment treatment is not random and depends on some pretreatment covariates. The information of this dependence is very important, and must be obtained by expert knowledge and cannot be tested using statistical data. Conceptualizing the observational study as a hypothetical experiment offers great insight to perform causal inference and removes all ambiguity in what is the treatment and what is the outcome It also helps to identify potential variables that could have affected the treatment assignment and the outcome variable (a loose definition of confounders), creating a bias in the estimate of causal effect. The hypothetical experiment to estimate the causal effect of PMR on safety is explained in Chapter 6.

Chapter 5 Causal Diagrams Framework

In this section, a brief introduction to the theory of probabilistic causal inference using Bayesian Networks is provided and the terminology and mathematical representation used in the analysis is explained. This method uses graphs to represent the causal relationship between different variables, qualitatively. The quantitative information is stored as a set of parameters of the graph. Causal effects are estimated by modifying the properties of the graph and performing inferences on the parameters of the graph.

5.1 Graph Terminology

A directed graph is a pair $G = (V, E)$ where V is a finite set of vertices and the set E of directed edges is a subset of VxV ordered pairs of distinct vertices. An edge from a vertex a directed onto an edge b indicates $(a, b) \in E$ and $(b, a) \notin E$. If there is an edge directed from a to b, a is said to be the parent of b, and b is said to be the child of a. The set of parents of b is denoted as $pa(b)$. A path between a and b is a sequence $a = a_0, a_1, a_2, \ldots b$ of distinct vertices such that $(a_i - 1, a_i) \in E$ or $(a_i, a_i - 1) \in E$ for all $i = 1,2, \ldots n$. A directed path from a to b is a sequence $a = a_0, a_1, a_2, \ldots b$ of distinct vertices such that $(a_i - 1, a_i) \in E$ and $(a_i, a_i - 1) \notin E$ for all $i = 1,2, \ldots n$. If there exists a directed path from a to b, a is said to be an ancestor of b and b is said to be the descendent of *a*. The directed path which begins and ends in the same vertex is said to be a cycle. If a directed graph has no cycles it is called a directed acyclic graph. An example DAG is shown in Figure 5-1

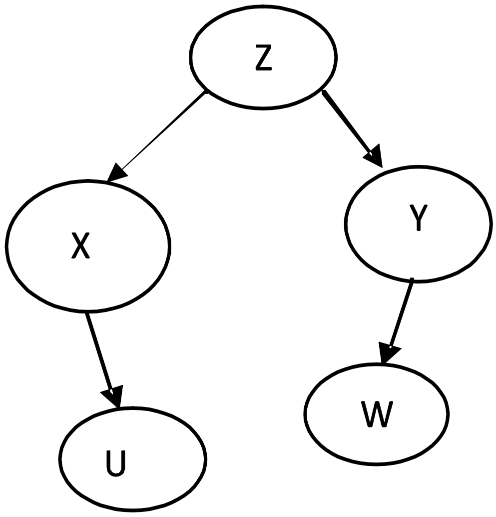

Figure 5-1: An example DAG

5.2 Bayesian Networks

A graph represents the qualitative relation between a set of variables. The strength of the relationship is represented by probability distributions connecting the variables. A Bayesian Network is a structure that combines both the qualitative and the quantitative relationship between variables of interest. What follows next is a formal definition of a Bayesian Network and some properties of a Bayesian Network useful for the analysis.

Let $f(x_1, \ldots, x_n)$ be a joint distribution of a set of variables $V = \{X_1, \ldots, X_n\}$ and $f(x_i|x_j)$ the conditional distribution of x_i given x_j. A Bayesian Network for the set of variables V consists of a directed acyclic graph G that encodes:

(1) A set of conditional independence assertions about the variables in V, and

(2) A set P of local probability distributions associated with each variable.

Together, these components define the joint probability distribution for V. The nodes in G are in one-to-one correspondence with the variables in V. The lack of possible edges in G encodes conditional independencies. The joint distribution of V is factorized recursively according to the graph G as the following equation:

$$f(x_1, \dots, x_n) = \prod_{i=1}^{n} f(x_i | pa(x_i)) \tag{1}$$

If a probability function P can be factorized by using equation (1) relative to a DAG G, G is said to be Markov compatible with P (Pearl, 2000).

If the joint distribution is factorized recursively according to the graph G, the conditional independencies implied by the factorization in (1) can be obtained from the graph G according to the following criteria due to Pearl (1988):

<u>D-Separation:</u> Let X, Y, Z be three disjoint subsets of vertices in a Bayesian Network G. Then Z is said to d-separate X from Y, if along every path between a vertex in X and a vertex in Y there exists a vertex W which satisfies one of the following three conditions:

1. W is in Z, and one arrow on the path points to w and the other arrow on the path merges from w,
2. W is in Z, and two arrows on the path emerge from w, and
3. Neither W nor any descendent of W is in Z, but two arrows on the path point to W.

If a path satisfies one of the conditions above, the path is said to be blocked by Z.

If sets X and Y are d-separated by Z in a DAG, then X is independent of Y conditional on Z in every distribution compatible with G. Conversely, if X and Y are not d-separated by Z in a DAG G, then X and Y are dependent conditional on G in at least one distribution compatible with G (Pearl, 2000)

It must be noted that Bayesian Networks merely describe the conditional independence structure of a set of variables and that they do not necessarily imply causation.

5.3 Causal Bayesian Networks

As noted previously, a Bayesian Network is simply an efficient way to represent the joint probability distribution of a set of variables V and the interpretation of conditional independence in the network does not necessarily imply causation. To be able to estimate causal effects from a Bayesian Network, the mechanism that generated the dependencies (or independencies) in the data should be modeled. A Causal Bayesian Network is a model for such casual mechanisms. It is a Bayesian Network in which the predecessors of a node are interpreted as directly causing the variable associated with that node. To be able to interpret a Bayesian Network as a causal network, the following assumption must be satisfied:

Markov Assumption: Given the values of a variable's immediate causes, the variable is independent of its earlier causes.

Having a causal interpretation facilitates predicting the effect of an intervention in the domain: setting the value of a variable in such a way that the manipulation itself does not affect the other variables.

The effect of intervention on the probability of a variable X is represented by the *do* operator (Pearl, 2000). Thus if X is a variable that can assume the state x, $do(X = x)$ represents the action of manipulation and forcing the variable X to take the value x without affecting other variables. The formal definition of a Causal Bayesian Network is given below.

Causal Bayesian Network: (Pearl, 2000) Let $P(V)$ be a probability distribution on a set V of variables, and let $P_x(v)$ denote the distribution resulting from the intervention $do(X = x)$ that sets a subset X of variables to constants x. Let P^* be the set of all interventional distributions $P_x(v)$, $X \sqsubseteq V$, including $P(v)$ which represents no intervention (i.e. $X = \Phi$). A DAG G is said to be a causal Bayesian Network compatible with P^* if and only if the following three conditions hold:

(1) $P_x(v)$ is Markov relative to G (i.e., every variable in G is independent of all its nondescendents (in G) conditional on its parents);

(2) $P_x(v_i) = 1$ for all v_i belongs to X whenever v_i is consistent with $X = x$;

(3) $P_x(v_i|pa_i) = P(v_i|pa_i)$ for all $v_i \notin X$ whenever pa_i is consistent with $X = x$.

It follows from the above definition that once we control the direct causes of a node v_i, no other interventions will affect the probability of v_i.

Causal Model: A Causal Model is a pair $CM = \langle S, \Theta \rangle$ consisting of a causal structure S and a set of parameters Θ compatible with S. The parameters Θ assign a function $x_i = f(pa_i, u_i)$ to each $X_i \in V$ and a probability measure $P(u_i)$ to each u_i.

In other words, a Causal Model is a Causal Bayesian Network with conditional probability distributions that are compatible with the causal mechanism encoded in the model structure. The problem of causal inference involves learning the causal structure and conditional probability distributions from data and estimating the causal effect under intervention by using the Casual Bayesian Network. There are several possible methods to perform these tasks and some of them depend on the nature of the problem. The method used for each of these steps in the present analysis is explained in Chapter 7.

Chapter 6 Description of the Data

6.1 Structure of the Data

To appreciate the difficulty associated with carrying out causal inference in transportation safety studies, one needs to know the mechanism in which crash data are generated and the structure of the database. First and foremost, crashes are rare events. In comparison to the amount of vehicular flow observed on a particular section of the road, the number of crashes observed on the same road section is almost zero; the probability measures are, therefore, very low and close to zero. This creates the problem of identifying probabilistic dependence of two events. Secondly, the factors causing crashes are different in different cases, and sometimes the chain of events causing a crash is specific only to a particular crash. This limitation makes generalization difficult and hence identifying causes that are specific to a single event is also difficult. Thirdly, causal relations are supposed to be reliable only when obtained from randomized experiments, but such experiments are not a part of the toolbox of a transportation safety researcher, due to ethical and financial reasons. A safety researcher can, instead, use the instrument of observational studies. Elucidation of causal relationships from observational studies requires several assumptions which are reviewed in sections 7.5 and 7.6

To understand how crash data are generated, an insight into the hierarchy of the fundamental units used for any safety analysis is needed. Figure 6-1 shows the hierarchical nature of transportation databases.

Figure 6-1: Hierarchy of Crash Databases

At the lowest level of the hierarchy are the individual drivers and their vehicles. Drivers use different roadway types for various trip purposes. These roadways can be divided into sections that have homogeneous geometric design features (horizontal curve, vertical curve, cross-section elements, etc.) called segments. This adds two more levels to the hierarchy, including: the segments at the second level and roadways at the third level. The hierarchy can be continued in this way to include geographically separated units like regions (or districts) within a state, or the entire state.

It is also important to note the mechanism in which crash data are collected. Generally, a crash event is followed by a police investigation. Transportation agencies use the reports of these investigations to create electronic records of the crash event. Detailed information related to the type of crash, its location, possible contributory factors, and driver and roadway characteristics, are among the many elements included in the data. Unfortunately, similar information is not available for vehicles and drivers that are not involved in a reportable crash. To overcome this problem, most traffic crash frequency analyses are carried out at the roadway segment level. Individual drivers are removed from the hierarchical structure of the transportation safety analysis framework due to the enormity of effort and near impossibility of collecting such data. This modeling framework eliminates the need to include any information specific to an individual driver in the database.

6.2 Design of Study

The fundamental unit of operation is a homogeneous road segment (homogeneous here refers to having uniform geometric characteristics such as number of lanes, lane width, shoulder width, grade, etc). Examination of the effect of increased pavement marking retroreflectivity on the safety of a road segment is the primary purpose of the present study. Safety of a road segment is defined as a dichotomous variable, taking the value 1 if there was at least one target crash in the segment during the treatment (or control) application period and 0 if there was no target crash in the segment. The safety of a segment is stochastic and each segment has a fixed probability of having a safety outcome of 1, which is assumed to be an inherent property of the road segment. Thus, safety of a segment is a Bernoulli random variable with the probability of success (where success is defined as occurrence of at least one crash within a fixed period of time) varying for each segment. A countermeasure may increase, decrease or leave unchanged, the safety of a road segment.

As per Rubin (2008) and Maldonando and Greenland (2002), the research question is conceptualized as a hypothetical experiment to make the problem statement clear. Consider a population of homogeneous road segments. The effect of increased pavement marking retroreflectivity (PMR) on the safety (crash vs. non-crash state) of a road segment is to be estimated. The treatment variable is PMR and has three levels: low, medium and high (the exact range of PMR levels for each class is specified in the description of the data section below). The effect of changing the PMR level from medium to low, from high to low and from high to medium on the safety of a segment are to be estimated. As per the definition of safety, the outcome of the experiment is stochastic and each unit (road segment) has a different probability of experiencing a crash. To measure the effect of changing the PMR levels from medium to low, the following hypothetical experiment would be performed. A random sample of homogeneous road segments would be divided into two groups. One of the groups (group1) would be administered the "low" treatment, i.e. pavement markings would be painted with retroreflectivity falling within the levels defined by the class "low" and the number of crashes would be recorded for a fixed period of time. Similarly, group2 would be administered the "medium" treatment. Let R1 be the proportion of segments with at least one crash in group1 and R2 be the proportion of segments with at least one crash in group2, for the given time period. The average causal effect (ACE) of changing PMR levels from medium to low is defined by the ratio of R1 to R2. The average causal effect for the other changes in the PMR levels can be defined in a similar way. Since such a hypothetical experiment is not feasible due to ethical and financial reasons, an observational study is used as an alternative.

6.3 Description of the data

Crash and pavement marking data were collected retrospectively from three districts in North Carolina for a period of 2.5 years. The pavement marking retroreflectivity data were measured by Flint Trading Company using a mobile retroreflectometer with a 30-meter geometry. These data were collected on two-lane and multi-lane highways in North Carolina. All pavement markings were thermoplastics. Since retroreflectivity estimates were not measured at the exact time and place of occurrence of the crash, a neural network model was estimated to interpolate the values of retroreflectivity on the segments where crashes were observed. The details of the model are provided in Karwa and Donnell (2009, forthcoming).

The roadway inventory and crash event data were obtained from the Highway Safety Information System (HSIS) data files, maintained by the Federal Highway Administration (FHWA). These data were collected for the 19 roadway sections where retroreflectivity data were collected. A roadway segment in the HSIS files is homogeneous in certain roadway characteristics. There were 192 total segments that corresponded to the 19 sections of roadway where pavement marking retroreflectivity estimates were computed based on the degradation model. Table 2 shows the locations where roadway inventory, crash, and pavement marking retroreflectivity data could be linked.

Crashes that satisfy the following criteria, referred to as target crashes, were used in the analysis:

- No weather contributing circumstances
- Occurred during dusk, dawn, or at night
- Non-work zone area

- No alcohol involvement
- Dry roadway surface conditions
- No roadway contributing circumstances
- Ran-off-the-road crashes
- Fixed object crashes (off-road)
- Opposite- or same-direction sideswipe crashes

Table 2 Roadways with Pavement Marking and Crash Data Available for Safety Analysis

County	Route	District	Begin MP	End MP	Number of Lanes	Total Length (miles)	Nighttime Target Crashes
Bertie	US13	1	0.00	11.07	4	11.07	8
Gates	US13	1	0.00	14.78	2	14.78	14
Northampton	US158	1	12.35	24.04	2	11.69	4
Washington	US64	1	10.54	19.67	2	9.13	2
Durham	I-85	5	7.88	14.19	4	6.31	5
Durham	US15	5	3.66	6.56	4	2.90	9
Durham	NC98	5	0.00	11.06	2	9.44[a]	15
Durham	NC157	5	0.70	3.98	2	3.28	2
Granville	I-85	5	0.00	23.73	4	1.80[b]	5
Person	US158	5	0.00	22.36	2	16.22	5
Vance	I-85	5	0.00	14.47	4	12.47[c]	46
Vance	US158	5	0.00	8.96	2	5.94[d]	1
Wake	I-40	5	6.47	20.19	4/6/8	13.72	67
Wake	NC98	5	0.00	4.55	2	4.55	1
Warren	I-85	5	0.00	9.88	4	9.88	39
Warren	US158	5	12.38	22.93	2	10.55	0
Catawba	I-40	12	13.13	19.67	4	6.54	10
Iredell	I-40	12	0.00	22.76	4	22.76	63
Iredell	I-77	12	14.75	23.75	4	9.00	17
		Total				182.03	313

[a] Roadway inventory and crash data were not available between mileposts 0.17 & 1.79.
[b] Roadway inventory and crash data were not available between mileposts 0.76 & 22.69 & 22.73.
[c] Roadway inventory and crash data were not available between mileposts 3.96 & 5.96.
[d] Roadway inventory and crash data were not available between mileposts 3.77 & 6.79.

This was done to ensure the absence of confounders on the probability of a crash in a segment and due to lack of appropriate covariates. For instance, if it was clear according to the police crash report that a particular crash occurred due to driving under the influence of drugs or alcohol, such a data point would not make a segment unsafe. Such a crash would have been deemed to occur as a result of human error, and hence excluded from the current analyses. Similarly, crashes in which weather was a contributing factor (such as heavy snow, or icy road conditions) were also excluded from the analyses. Weather conditions, human errors, etc. are stochastic factors causing crashes and not an inherent property of the segment; hence, any crash due to such conditions would fall into the error term of the observed safety of a road segment. This issue is further discussed in light of the assumptions involved in the learning approach in section 7.6.

The treatment variable, pavement marking retroreflectivity, is discretized into three bins of low, medium and high. Several different ranges of values were used for discretizing to examine the effect of discretization policy. An example discretization setting into 3 bins is given below:

$$low\ if\ 139 < retroreflectivity \leq 200$$

$$medium\ if\ 200 < retoreflectivity \leq 280$$

$$high\ if\ 280 < retroreflectivity \leq 446.53$$

Based on the above policy, out of the total sample size, about 36 percent of the segments had low levels of PMR, about 46.5 percent had medium levels, and the remaining segments had high levels of PMR. The assignment of PMR levels is clearly not random. It must also be noted that the data are sparse due to rarity of crashes. At least one crash occurred on only 6 percent of the total number of segments during the study period. Apart from the data on pavement marking retroreflectivity and the crash counts per month, data on 13 other covariates were collected. Information on the attribute of a segment, such as the shoulder widths; number of lanes; presence of a median; level of access control; posted speed limit; traffic flow characteristics, such as monthly traffic volumes (hereafter referred to as ADT); percentage of trucks; location-related variables, such as the geographic district in which the segment is located; the urban or rural setting of the segment location (which is defined by the population density of the location); and, the terrain type, were collected. The definition and summary statistics of the all the variables in the dataset are provided in Table 3.

A segment within a fixed time period is considered to be different from the same segment at any other time period. This period of time was set equal to one month. This period was selected to ensure homogeneity of treatment and other characteristics in a segment. The level of pavement marking retroreflectivity was assumed constant within a one month period. Also, the average daily traffic (ADT) was also assumed constant within this period.

The analysis is restricted to the class of directed acyclic graphs with discrete variables called discrete Bayesian Networks (hereby referred to as just Bayesian Networks). These restrictions are made to enable the use of efficient algorithms for the generation of and inference in Bayesian Networks. Bayesian Networks with continuous variables are possible, but require certain additional assumptions that are not feasible in the current analysis. For instance, one of the assumptions requires the variables to be Gaussian distributed. Also, many algorithms cannot handle mixed Bayesian Networks (mixed here refers to the combination of continuous and discrete variables) that have continuous parents of discrete children. In the present analysis, it is possible for such combinations to occur. For example, there could be a continuous parent of the target crash indicator, or the PMR levels need not be Gaussian. Thus, a straightforward solution is to discretize the entire dataset and perform the analysis using discrete Bayesian Networks. Discretization is performed using subject matter knowledge in safety and the binning boundaries to reflect common cut-offs used in the functional classification of roads.

Table 3 Descriptive Statistics of Data

Variable	**Definition**	**Mean**	**Std. Dev.**	**Min**	**Max**
Left Shoulder Width	Outer shoulder width in feet on the left side of the travelled way	2.84	4.23	0	13
Right Shoulder Width	Outer shoulder width in feet on right side of the traveled way	9.38	4.02	0	14
Adjusted AADT	Annual average daily traffic adjusted for the period of a month, veh/day	30,383	27,580	1615	114,400
Pct. Trucks	Percentage of the ADT that consists of heavy vehicles	14.77	8.54	0	83
Avg. Retroreflectivity	Average value of Pavement Marking Retroreflectivity of all markings on a segment (PMR)	227.36	64.45	139.64	446.53
Variable	**Definition**		**Mean**	**Min**	**Max**
Multilane	Multilane Indicator, 1 if the segment has more than 2 lanes, 0 otherwise		65%	0	1
Median Presence	Median indicator, 1 if the segment has divided with a median, 0 otherwise		68%	0	1
Safety	Safety indicator, 1 if the segment witnessed at least 1 target crash during the month, 0 otherwise		6%	0	1
Urban/Rural	Urban/Rural indicator, 1 if		45%	0	1
Low Speed vs High	Speed limit indicator, 1 if speed limit less than or equal to 35, 0 otherwise		15%	0	1
Flat vs Rolling	Terrain indicator, 1 if flat terrain, 0 otherwise		22%	0	1
Division 1	NC division indicator, 1 if the segment is located in district 1, 0 otherwise		22%	0	1
Division 5	NC division indicator, 1 if the segment is located in district 5, 0 otherwise		56%	0	1
Low PMR	Low PMR level indicator, 1 if PMR 139 < retroreflectivity ≤200, 0 otherwise		35%	0	1
High PMR	High PMR level indicator 1 if PMR 280 < retroreflectivity ≤ 446.53,0 otherwise		18%	0	1
Med PMR	Med PMR level indicator 1 if PMR 200< retroreflectivity ≤ 280, 0 otherwise		47%	0	1
Total Sample Size = 5918					
PMR bins are based on a sample discretization policy					

Table 4 shows the summary statistics of the discretized dataset and Table 5 shows the different discretization bins used for PMR to evaluate the sensitivity of the results of average causal effect (ACE) to the categorization of PMR.

As it can be seen from Table 5, the only free parameter was chosen to be the upper bound of bin 1 (i.e., the low PMR level). This was done for two reasons. First, it is of interest to find if there is any evidence for a minimum level of PMR threshold above or below which the safety reduction is drastic. Second, the bin sizes were divided in a way to allow for a tradeoff between varied bins and the number of data points in each bin. Moving two bin levels at the same time, while trying to maintain a sufficient amount of data in each bin, was difficult to achieve.

Table 4 Discretized Dataset for Bayesian Network

Variable	**Definition**	**Percentage**
Right Shoulder Width	1: Below 7 feet; 0: Otherwise	20.9
Adjusted AADT	1: below 30,000 vpd; 0: Otherwise	51.4
Pct. Trucks	1: Below 18 %; 0: Otherwise	53.1
Multilane	1: No. of lanes > 2; 0: Otherwise	65%
Median Presence	1: If Median Present; 0: Otherwise	68%
Safety	1: If at least 1 target crash ; 0: Otherwise	6%
Urban/Rural	1 if Urban area; 0 Otherwise	45%
Low Speed vs High	1 if speed limit < = 35 mph; 0 Otherwise	15%
Flat vs Rolling	1 if Terrain is flat; 0 Otherwise	22%
Division 1	1 if NC District 1; 0 Otherwise	22%
Division 5	1 if NC District 5; 0 Otherwise	56%
Low PMR	Low PMR level indicator, 1 if PMR 139 < retroreflectivity ≤200, 0 otherwise	35%
High PMR	High PMR level indicator, 1 if PMR 280 < retroreflectivity ≤ 446.53,0 otherwise	18%
Med PMR	Med PMR level indicator, 1 if PMR 200< retroreflectivity ≤ 280, 0 otherwise	47%

Table 5 Discretization Policy for PMR

Policy Number	Upper Bound Bin 1	Lower Bound Bin 3
1	270	350
2	265	350
3	260	350
4	255	350
5	250	350
6	245	350
7	240	350
8	235	350
9	230	350
10	225	350
11	220	350
12	215	350
13	210	350
14	200	350
15	195	350
16	190	350
17	185	350
18	180	350
19	175	350
20	170	350
21	165	350
22	160	350
23	155	350
24	150	350
25	145	350
Average	207.6	350

Chapter 7 Estimating Causal Effects

To estimate the causal effect of PMR on safety, the casual model must be learned from observational data. The following sections provide information on learning the causal model.

7.1 Causal Grammar and Causal Diagram

Causal diagrams provide an intuitive way to represent causal relations between a set of variables. Apart from using data, the qualitative relationship between variables can be extracted from domain experts as well as past research studies, and represented in the form of a DAG. This provides an excellent way to encode and build upon past research. The quantitative relations can also be extracted, to a certain extent. However, some of the causal relations uncovered in past research may not possess generalization abilities. Causal grammar is used to address this problem. The concept of causal grammar was introduced by Tenebaum and Niyogi (2003).

Causal grammar can be defined as a framework to characterize the causal knowledge in a domain with some degree of abstraction. Linguistic grammar provides abstract classes (like different parts of speech) and rules about the relations between these abstract classes. Like grammar for a language, causal grammar specifies abstract classes of entities (variables or nodes in a domain, instead of parts of speech) and rules about the relationships (causal relations, instead of syntactic relations) that may exist between entities of various types. The principal function of causal grammar, much like natural language grammar, is to generate a constrained space of hypotheses that can be considered for further analyses.

One of the goals of this thesis is to introduce this concept and illustrate its use in safety studies. Causal diagrams and causal grammar provide a very powerful and rich framework to encode causal relations. Knowledge representation using DAGs makes the implicit causal assumptions very transparent, thus eliminating ambiguity. The use of causal grammar helps incorporate results from past studies and development of abstract principles in a domain. Causal grammar can also be used to convey generic assumptions made in a causal inference study, which will be illustrated below. To draw a parallel with a field from which safety studies often borrow several methodological developments, consider cause-effect relationships in epidemiology. Simple causal grammar in this field could involve three abstract sets: disease, behavior (a person's daily food habits, etc) and symptoms. Three possible (hypothetical) causal grammar relationships are shown in Figure 7-1. For instance, grammar A shows that behavior causes symptoms, and behavior and symptoms cause disease. Grammar B shows that behavior causes symptoms and disease. Thus, any causal diagram generated by grammar B cannot have symptoms causing disease. Similarly, grammar C asserts that behavior cannot cause symptoms directly; any causal diagram generated by this grammar must also have the same assertion.

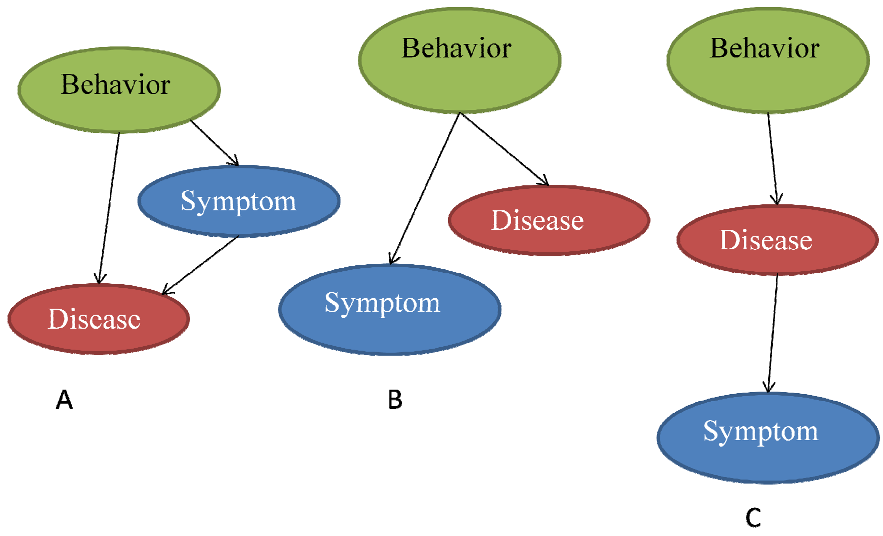

Figure 7-1: Hypothetical Causal Grammar in Epidemiology

Similar grammar relationships in safety studies can be constructed. In the current thesis, the general grammar relationships illustrated in Figure 7-2 are used. This grammar consists of five classes or attributes of any segment. The design class includes the geometric design variables of the segment, such as the horizontal curve radius, posted speed limit, number of lanes, lane and shoulder widths, etc. The operation class consists of macroscopic traffic flow parameters like ADT, percentage of trucks in the traffic stream, traffic density, PMR, etc. The operations class could also consist of microscopic traffic flow parameters such as the speed of a vehicle, lateral position of the vehicle in the travel lane, etc. The operation class variables can be influenced by the design variables, but not the other way around. The environment class consists of weather-related variables such as rainfall, snow, fog, etc., and the roadside setting along a roadway segment. The human factor class consists of variables related to driver and vehicle attributes, such as the age of the driver, reaction time of the driver, experience of the driver, vehicle type, etc. The last class is the safety class, which can include the number and type of crashes, and their severity. The nodes in red denote classes that are completely unobserved. The important thing to note is that in this causal grammar, human factors are independent of all other variables. Also, the environment class does not influence the operation class. This may not be completely true, since severe weather conditions can reduce traffic volumes and vehicle operating speeds. Nonetheless, this assumption is made since weather-related variables are not observed. Similarly, the human factors class variable influences the microscopic operational characteristics of traffic flow, such as speed and lane position. Collection of human factors data for non-crash vehicles is nearly impossible, and crash event data can be subjective in police accident reports. To eliminate these problems, it is assumed that human factors class variables do not influence operational class variables. Since the analysis is performed at the segment level, and macroscopic variables such as ADT, percent heavy vehicles in the traffic stream etc. are not affected by individual human factors, this assumption possesses some validity. This assumption becomes clear when the causal grammar is modified to divide the operation class into two subclasses – marcoscopic and microscopic, as shown in Figure 7-3. The definition of safety is modified accordingly, as noted previously in Section 6.2

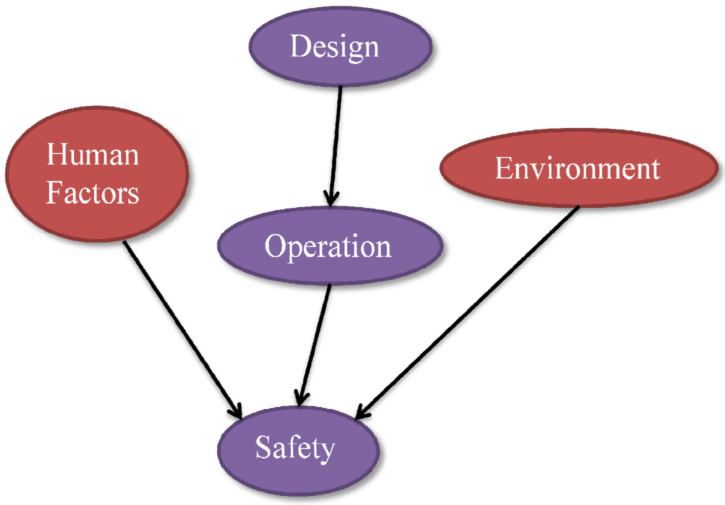

Figure 7-2: Causal Grammar used in the current study

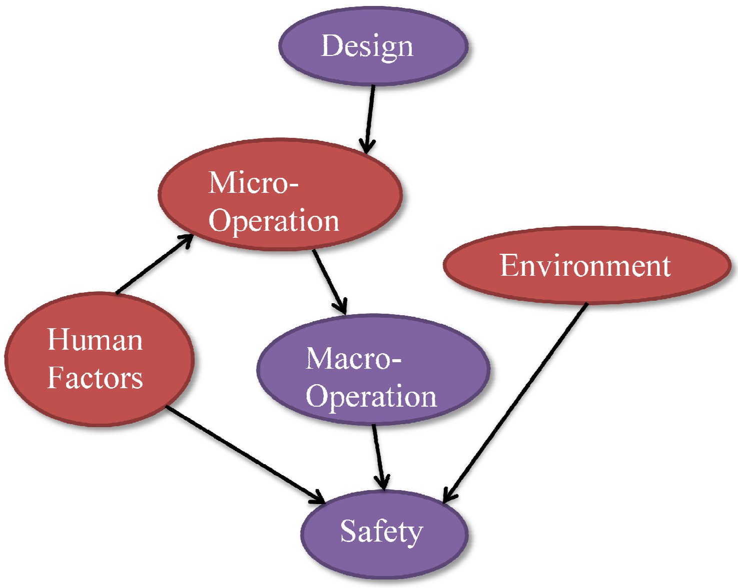

Figure 7-3: Modified Causal Grammar of Figure 7-2

Such causal grammar relationships can be used to specify constraints on causal diagrams when learning causal relations from data. Thus, the search algorithm would be constrained to recover causal graphs that are compatible with one or more selected causal grammar relationships. This issue is explained in the next section on causal structure learning.

7.2 Learning the Causal Structure

The problem of learning a Bayesian Network can be stated as follows. Given a database $D = \{x_1, x_2, \dots x_n\}$ containing independent instances of variables V, find a network $M = \langle G, \Theta \rangle$ that best matches V where G is a directed acyclic graph and Θ is a probability distribution of the variables in D. It has been noted that observational data limits our ability to learn network structures only up to the D-separation equivalence. Hence the search is carried out for an equivalence class of networks represented by a pattern (Pearl et al., 1991; Spirtes et al., 1991) that best matches D.

The issue of ability to learn causal networks from observational data has received a thorough treatment in the literature (Pearl et al., 1991; Spirtes et al., 1991; Heckerman, 1999). The most common strategies fall into two different classes called constraint-based learning and score-based learning. A simple combination of both approaches is used to learn the structure of the Bayesian Network. Irrespective of the strategy used to learn a causal model from observational data, the following three assumptions must be satisfied.

<u>Causal Markov Assumption</u>: It is assumed that the (unknown) casual structure of the domain of variables satisfies the Markov condition (i.e., a variable X_i in the DAG is independent on all its nondescendants, given its parents Pa_i [Pearl et al., 1991]). This assumption has two implications:

1) A commitment to include in the model every variable that is a cause of two or more other variables, and

2) Reichenbach's (1956) common cause assumption, stating that, if any two variables are dependent, then one is a cause of the other or there is a third variable causing both.

This assumption also implies that if the child-parent relationship in a causal model is specified as a deterministic function instead of the conditional probability $P(x_i|pa_i)$, it imposes the same independence constraints on the resulting distribution and leads to the same recursive distribution specified by equation (1). This holds regardless of the choice of the functions and the error distributions. In other words, once $P(x_i|pa_i)$ is measured or estimated, it is not necessary to specify the functional form or the distribution of the error terms in advance. This feature provides a great advantage in Causal Inference for phenomena where the mechanism between a cause and an effect is unknown or non-linear, as in the present case.

<u>Faithfulness:</u> The assumption of faithfulness requires that all conditional independence relations observed in the dataset are due to probabilistic independence and not due to chance (i.e., the independence assertion remains even if there is a slight change in the probability distribution). This assumption can be stated in a different form known as stability (Pearl, 2000) which conveys that all the independencies embedded in the graph G are stable (i.e., they are entailed by the structure of the model S and hence remain invariant to any change in the parameters Θ).

<u>No Latency</u>: It is assumed that there are no hidden variables in the dataset. This assumption may not be true if the entire database of crashes is taken into account for modeling. As noted in section 6.3, only target crashes are used to develop the model. This is

done to ensure that the assumption of absence of hidden variables is satisfied. This issue is discussed in detail in the section on graphical identifiably of casual effects.

With the assumption of stability and absence of latent variables, the joint distribution of variables has a unique causal pattern up to d-separation equivalence.

Constraint-based strategies work by constraining the DAG structure to represent the conditional independence relations observed in the dataset. The constraint-based strategies begin with a fully connected graph and use queries of statistical tests of conditional independence tests in the data to eliminate edges that represent such conditional independences. Conditional independent relations in a DAG can be read using the d-separation criterion provided by Pearl (2000). Thus, these methods are limited by the statistical power of the conditional independence tests and are affected adversely by sparseness in the database.

On the other hand, score-based learning strategies work by minimizing a score based on the distance between the recursive joint probability distribution (specified by eq 1) generated by the DAG and the actual observed joint probability distribution of data and the complexity of the network. These methods work by using a random (or user specified) initial value of the DAG structure. Random moves are proposed in the structure of the graph (such as deletion or insertion of an edge) and the search proceeds until a graph with the minimum score is selected. The stopping criterion is specific to the method used to implement the random search and the scoring function. The score-based methods are not limited by the problems created by statistical independence tests in sparse data, but they may end up in recovering a graph may not be the global optimum.

A simplified strategy is used to combine and gain the power of both methods. The causal graph recovered is assumed to be generated by the grammar provided in Figure 7-2. The constraints specified by the grammar are supplied as initial constraints to the learning algorithm. A constraint-based method is used to recover a DAG structure from the data. This DAG is supplied as an initial structure to the score-based learning strategy, which then attempts to find an optimum DAG structure. Bayesian inference is used to learn the structure of the DAG. The PC algorithm (Sprites et al,. 1991) is used as a constraint-based algorithm and the simulated anneal strategy of (Hartemink, 2005) is used to search for the optimum scored Bayesian Network. The scoring function proposed by (Heckerman et. al., 1995) is used. The algorithm is implemented in Java using the library of banjo (Hartemink, 2005). The details of the PC algorithm are provided in Appendix A. Since the scoring method need not produce the globally optimum structure of the Bayesian network, we use the 10 best networks recovered by the algorithm and perform Bayesian model averaging to estimate the average causal effect. This issue is further discussed in section 7.7.

It is very important to realize that the structure learning algorithm does not make causal discoveries by magic. It simply searches for a Bayesian Network that best represents the data, constrained by the assumptions made. Causal interpretation of such a network is possible only because of these assumptions, which come from subject matter experts.

After the causal structure is recovered from the database, the next step is to learn the conditional probability distributions of the DAG.

7.3 Learning the parameters of the Causal Model

A Causal Bayesian Network is fully specified by its structure (S) and the conditional probability distributions of the variables (Θ). In the present case, where the data are discrete, the parameters are specified by using conditional multinomial distributions in the form of conditional probability tables. The parameters of the model are learned from the data set using Bayesian Estimation and the configuration of the parameters (Θ_s) that maximizes the likelihood function is computed. The observed data set $D = \{x_1, x_2, \dots x_n\}$ is treated as a random sample from the true joint probability distribution of the variables where x_i is a vector representing a single case of the random sample. A prior distribution of the parameters of the DAG $P(\Theta|S)$ is specified where S is the DAG. The posterior distribution of the parameters $P(\Theta|D, S)$ that maximizes the log likelihood function (2) is then estimated.

$$g(\Theta_s) = log(p(D|\Theta_s, S) * p(\Theta_s|S)) \tag{2}$$

The Dirichlet distribution that has the following form is used to specify the prior probability distribution of parameters in the conditional probability table.

$$p(\theta_{x_i|\pi_i}|D) = \frac{\Gamma(\sum_{x_i} \alpha_{x_i,\pi_i})}{\prod_{x_i} \Gamma(\alpha_{x_i,\pi_i})} \prod_{x_i} \theta_{x_i|\pi_i}^{\alpha_{x_i,\pi_i}-1} \tag{3}$$

where $\theta_{x_i|\pi_i}$pertains to variable X_i in a state x_i given that its parents are in joint state π_i $(i = 1, \dots, n; n$ is the total number of variables in the domain). The normalization terms in equation 3 involve the Gamma function $\Gamma(.)$ which satisfies $\Gamma(x+1) = x\Gamma(x)$ and $\Gamma(1) = 1$.

The posterior density of the distribution given the observed dataset D is given by the following closed form equation:

$$p\left(\theta_{x_i|\pi_i} \middle| D, S\right) = c * \prod_{x_i} \theta_{x_i|\pi_i}^{n_{x_i,\pi_i} + \alpha_{x_i,\pi_i} - 1} \tag{4}$$

where c is a normalizing constant and n_{x_i,π_i} are the cell counts from the dataset D.

The expected value of the parameters has the simple expression:

$$E_{p\left(\theta_{x_i|\pi_i} \middle| D,S\right)}\left[\theta_{x_i|\pi_i}\right] = \frac{n_{x_i,\pi_i} + \alpha_{x_i,\pi_i}}{n_{\pi_i} + \alpha_{\pi_i}} \tag{5}$$

There are several advantages of using Dirichlet priors:

1) The Dirichlet distribution is a conjugate family of distributions for the multinomial family which is evident by equation (4),

2) The Dirichlet prior is tied to the desirable likelihood-equivalence property of graph structures that are Markov equivalent which states that data cannot help to discriminate network structures that represent the same assertions of conditional independence, and

3) The positive hyper parameters can be interpreted as pseudo-counts implied by the prior belief (Heckerman et al, 1994).

The hyper parameters α_{x_i,π_i} are specified by following (Heckerman et al, 1994):

$$\alpha_{x_i,\pi_i} = \alpha * p(x_i, \pi_i), \tag{6}$$

where p is a (marginal) prior distribution of *pseudo-counts*; this ensures the likelihood-equivalence of Markov equivalent structures (Heckerman et al, 1994). The distribution p is chosen to be uniform for all the variables (representing non-informative prior).

7.4 Estimating Causal Effects

After learning the structure and the conditional probability table from the data, the next step is to estimate the causal effect by using the manipulative theory of Causation. As noted previously, Casual Bayesian Networks can be regarded as models of interventions when several casual assumptions are made. Causal effects can be estimated from such models if it is assumed that the Casual Bayesian Networks model the Causal Mechanism which generated the data. This approach is based on nonparametric structural equations modeling suggested by Pearl (1994) based on early works in econometrics (Frisch, 1938; Haavelmo, 1943; Simon, 1953). In this approach, the edges in a DAG stand for physical mechanisms amongst the variables of interest and these mechanisms are represented by functional relationships with random disturbances. The relationship between a child and a parent node in a DAG can be represented by the following function:

$$X_i = f_i(pa_i, e_i), (i = 1,2, \dots n) \tag{7}$$

where pa_i denote the parents of variable x_i in G and e_i $(1 \leq i \leq n)$ are mutually independent, arbitrarily distributed error terms (Pearl et al, 1991). The errors are due to exogenous factors that are not included in the analysis. It must be noted that if any of these factors influence two or more variables in the model, violating the independence assumption (or the Casual Markov assumption), the factor must be taken into account in the analysis (if the factor is not measured, it must enter the analysis as a latent variable). As noted previously, one of the implications of the Causal Markov assumption is that the functional form of the relation and the distribution of errors in equation (7) need not be specified as long as the Causal Markov assumption is satisfied.

The representation of edges in a directed graph in the form of a deterministic function provides a convenient way to specify the changes in the joint distribution due to external intervention. In the case of the simplest intervention where a node X_i is forced to take a particular value x amounts to lifting the existing mechanism on X and putting it under the influence of a new mechanism whose action is to force X_i to the value x, keeping everything else the same. This action is mathematically represented by $set(X_i = x)$. As explained by Pearl (2000), such interventions, called atomic interventions can be modeled in a DAG G by creating a new mutilated DAG G_x from G. In G_x, the links between Pa_i and X_i are removed, keeping the rest of the graph the same. The distribution imposed by the new graph G_x under the condition $X_i = x$ represents the effect of intervention and is called the post-intervention distribution. An example is shown in Figure 7-4. The causal effect of a variable X on Y can then be obtained by the following definition:

<u>Causal Effect:</u> Given two disjoint sets of variables X and Y, the casual effect of X on Y, denoted by $P(y|do(x))$ is a function from X to the space of probability distribution on Y. For each realization of x of X, $P(y|do(x))$ gives the probability of $Y = y$ induced from the mutilated graph G_x and substituting the value of X as x in this graph.

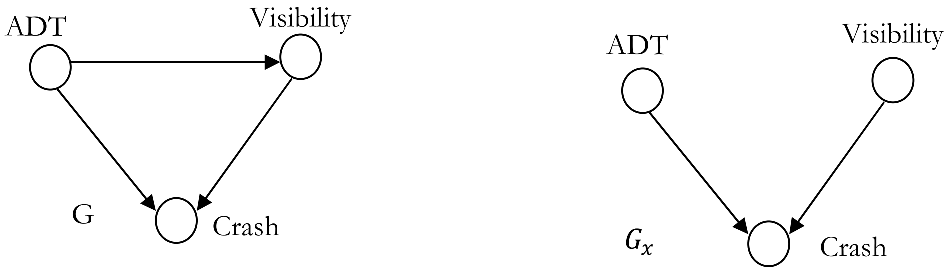

Figure 7-4: An example of the interventional distribution. The graph G represents the original DAG. The mutilated graph G_x under the intervention of forcing visibility to take a particular value is obtained by deleting the arcs between visibility and its parents (i.e ADT).

Mathematically, in the case of atomic intervention $set\ (X_i = x_i^{'})$, the joint probability distribution of variables in V can be obtained by the following equation:

$$P(x_1, \dots, x_n|do(x_i^{'})) = \begin{cases} \{P(x_1, \dots, x_n)\}/\{P(x_i|pa_1) & if\ x_i = x_i^{'} \\ 0, & if\ x_i \neq x_i^{'} \end{cases} \tag{8}$$

The immediate implication of this equation is that, given a causal diagram in which all the parents of manipulated variables is observed, the post-intervention distribution can be inferred from the pre-intervention distributions; the casual effect can be estimated from passive or non-interventional data. However, when some parents are not observed, $P(c|pa_i)$ may not be estimable in all cases. A graphical test has been provided by Pearl

(2000) to find out when $P(y|do(x))$ is estimable from the observed data. This issue is dealt with in the section on latent variables (Section 7.6).

The Average Casual Effect (ACE) of a variable X on Y, when X has two possible states (x_1 and x_2) is given by the following equation

$$ACE = \frac{P(Y = y|do(X = x_1)}{P(Y = y|do(X = x_2)} \quad (9)$$

where $P(Y = y|do(X = x_i)$ is the marginal probability of $Y = y$ under the atomic intervention $X = x_i$. Computing the marginal probability of a variable from the conditional probability distributions specified by a DAG under an evidence or intervention is not a trivial task when the number of variables is large. This is because the joint probability table increases exponentially with an increase in the number of variables (Jensen, 2001). There are several methods in the literature (Cowell, 1998) to efficiently perform inference in a Bayesian Network. The junction tree algorithm that performs exact inference in a Bayesian Network is used in the present research.

7.5 Causal Markov Assumption and Faithfulness

In this section, the assumption of Markovian Parents and Faithfulness is discussed in detail.

The Causal Markov assumption is an assertion of the fact that each variable is independent of its nondescendents, given its parents or immediate causes. The natural question that arises is what are the immediate causes of a crash? For instance, is driving at a high speed considered an immediate cause? To understand the Causal Markov assumption, it is important to first consider the factors that cause a crash.

Factors causing crashes can be divided into the following three broad categories: road user (driver), roadway characteristics (geometric elements, environment conditions, roadway volume, etc.), and the vehicle. Keeping in mind the hierarchy introduced in section 6.1, if one were to include driver level data, the immediate causes of crashes (75 percent of which are due to human error (Neville et al, 2009)) may become very specific to a particular crash and are governed by complex human behavior which is difficult to model. This rules-out the possibility of inferring type causes such as "speeding causes crashes". Safety research is more inclined at finding out general causes rather than token causes such as "Playing with the radio caused Jane's crash" because such statements are more appropriate when considering tort liability cases for specific crashes. Hence the notion of immediate cause is modified to include general factors like roadway and weather conditions; the specific human factors are included in the disturbance or error terms considered to be stochastic in nature.

This idea becomes clear when the Markovian assumption is treated as a convention rather than an assumption, where it merely defines the granularity of the model which is being considered. We can start from the deterministic extreme where all the variables that affect a system and its outcome (in this case, the road segment and the occurrence of a crash) are included in a microscopic model of the system. Such a hypothetical model would include variables representing individual driver behavior and its complex interaction with the roadway environment and other drivers. There would not be any level of randomness involved. We can move up to the next level of macroscopic abstractions by aggregating variables and introducing probabilities to summarize omitted variables and error terms to take into account the variability unexplained by the model. The other end of the extreme would be to aggregate all the variables and model crashes at the highest level of the hierarchy. Clearly, we need to decide at what stage the abstraction has gone too far and where the useful properties of causation are lost. The abstraction also depends on the specific problem at hand and the kind of data at our disposal. For instance, when dealing with a particular crash where our goal is to identify the party at fault (this is very common in tort liability cases), it would not be advisable to ignore driver-level data. Such problems demand a different definition of causation, closely related to the current definition, called the counterfactual models, whose discussion is out of the scope of this thesis. On the other hand, consider the current problem at hand where we are interested in knowing the behavior at the population level. In such a case, including the driver-level information would be not only impossible, but create several problems with statistical inference. The Markov condition guides us in selecting this level of abstraction and deciding when a set of parents Pa_i is considered complete in the sense that it includes all the relevant causes of a variable X_i. If the set Pa_i is very narrow, there will be disturbance terms that will simultaneously influence several variables and the Markovian assumption will no longer be valid.

When considering the factors that cause a segment to be accident prone (rather than considering the causes of an individual driver to be involved in a crash), certain causes dealing with human factors which are too specific can be ignored and included in the disturbance term. It must be noted that only those human causes that do not influence other factors like weather and roadway conditions and ADT can be eliminated from the model. The same is not true with environmental factors like weather conditions and a certain class of human factors. Removing these from the model would violate the Markovian assumption. For instance, the presence of a work zone influences both the ADT and occurrence of a crash; weather conditions influence both ADT and visibility of pavement markings. The problem is that data related to these factors is not available for non-crash locations. Hence these factors must enter the model in the form of latent variables. Inclusion of these factors in the form of latent variables destroys our ability to perform casual inference due to the problem of confounding bias and the absence of a set of variables to control for this bias. This issue is dealt with in the next section on graphical identification and latent variables (Section 7.6).

The second assumption of the Causal Model to learn structure from data is that of faithfulness. By assuming that a causal graph is Causally Markov, it is assumed that any population produced by this causal graph has the independence relations obtained by applying d-separation to it. It does not follow, however, that the population has exactly these and no additional independencies. If there are any independence relations in the population that are not a consequence of the Causal Markov condition (or d-separation), then the population is unfaithful. By assuming Faithfulness we eliminate all such cases from consideration. Assuming that a population is Faithful is to assume that whatever independencies occur in it arise not from coincidence but rather from structure.

7.6 Latent Variables and Graphical Identifiability

In this section the rationale behind the assumption of no latent variables in the model is explained. Latent variables impose two problems when estimating casual effects from observational data. The first problem deals with the difficulty involved in recovering DAGs with latent variables and the second deals with the estimation of causal effect from observational data and DAGs with latent variables. The former has several solutions but the latter problem, called Graphical Identifiability, can restrict the extent to which causal inference can be made from observational data (or in some cases, eliminate the ability to perform any kind of casual inference).

When latent variables are present in the model, the observed joint probability distribution need not be stable relative to the observed set of variables V. The constraints imposed by a latent structure upon a distribution cannot be completely characterized by any set of conditional independence statements. Certain sets of those independence constraints can be identified (Verma and Pearl, 1990) and there are algorithms which can recover latent structures from the observed data. The Fast search Causal Inference algorithm by Sprites et al. (1999) was used in recovering the DAG assuming that there are latent variables. A part of the DAG recovered by the algorithm is shown in Figure 7-5.

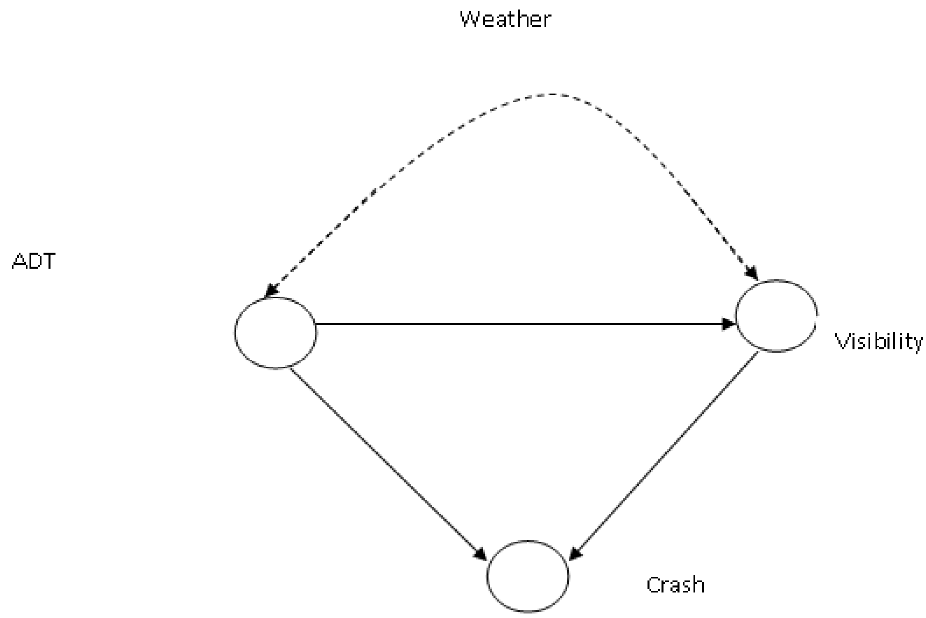

Figure 7-5: Part of DAG with latent variables

The graph reflects the fact that occurrence of nighttime crashes on a segment is effected by the amount of traffic flowing through the segment (ADT) and the visibility of pavement marking retroreflectivity on that segment (visibility). Both ADT and visibility are affected by the weather conditions on the segment. As noted previously, estimation of the causal effects of X on Y (crashes) may not be possible in the case where some parents (weather) of X are not observed. As it will be seen, the graphical identification criterion of Pearl (2000) is not satisfied by the graph of Figure 7-5, thus preventing the estimation of casual effects of visibility on crashes.

For details on the Graphical Identifiability, refer to (Galles and Pearl, 1995). The necessary and sufficient condition for a model to be a *nonidentifying* model, a model in which the total causal effect of X on Y is not identifiable, are noted here.

Necessary Condition: The graph contains unblockable back-door paths between X and Y, that is paths ending with arrows pointing to X which cannot be blocked by observed nondescendents of X.

Sufficient Condition: The graph contains a confounding path between X and any of its children on a path from X to Y.

The casual graph of Figure 7-5 satisfies both of these conditions; hence, the casual effect of visibility on crashes is not estimable.

To overcome the problem of identifying Causal Effects, to define the safety of a segment, only a subset of crashes is considered in the analysis. Crashes that satisfy the criteria defined by the definition of target crashes in Section 6.3 are considered for modeling. The target crashes are defined in such a way, so as to eliminate all those crashes whose conditions influence more than one variable in the model and these conditions are not observed (they are hidden). Such filtration ensures that only those crashes are considered where the common causes are measured (and hence can be included in the model) and the removal of specific factors of these crashes from the model does not violate the assumption of Markovian Parents.

7.7 Statistical Confidence in the Average Causal Effect

The definition of Average Causal Effect (ACE) given in section 7.4 assumes that the parameters of the Bayesian Network and its structure are fixed, and hence provides only a point estimate of the effect of Y on X. Both the parameters of the Bayesian Network and its structure are estimated from a finite sample of data and hence do not represent true probabilities and/or structure. Therefore, a point estimate of the ACE does not take into account the uncertainty associated with the parameters and the structure of the Bayesian Network. This uncertainty is taken into account by using the concept of Bayesian Model Averaging. The extent to which the data supports the estimate of ACE, given the uncertainty in the model structure and its parameters, is measured by the confidence in the value of the ACE estimated from the data. The Bayesian Credible set is estimated within which the ACE is expected to be present with a posterior probability of 0.95. Bayesian Model Averaging and a brief implementation to handle both uncertainties are described next.

7.7.1 Bayesian Model Averaging

Bayesian Model Averaging (BMA) (Hoeting et al., 1999) takes into account the uncertainty in model selection. The underlying idea is that when there are several competing models that fit the data reasonably well, instead of selecting a single model to be the final one, an average of the several competing models is used, thus taking into account the uncertainty involved. In the present context, if ACE is the average causal effect and $D$is the observed data, then

$$P(ACE|D) = \sum_{k=1}^{K} P(ACE|M_k, D) * P(M_k, D) \tag{10}$$

Equation 10 represents the average of the posterior distributions under each of the $K = \{1, \dots . k\}$ models considered, weighted by their posterior probability. The posterior probability of each model is given by:

$$P(M_k|D) = \frac{P(D|M_k) * P(M_k)}{\sum_{l=1}^{K} P(D|M_l) * P(M_l)} \tag{11}$$

where,

$$P(D|M_k) = \int P(D|\theta_k, M_k) * P(\theta_k|M_k) d\,\theta_k \tag{12}$$

Equation 12 is the integrated likelihood of the model M_k and θ_k is the parameter vector of the model. The posterior mean and variance of the ACE are given by:

$$\begin{aligned} E[ACE|D] &= E\big[E[ACE|D, M_k]\big] \\ &= \sum_{k=1}^{K} E[ACE|D, M_k]\, p(M_k|D) \end{aligned} \tag{13}$$

where the outer expectation is over the distribution of models M_k

$$\begin{aligned} Var[ACE|D] &= E[(ACE|D)^2] - E[ACE|D]^2 \\ &= E\big[E[(ACE|D, M_k)^2]\big] - E[ACE|D]^2 \\ &= \sum_{k=1}^{K} (Var[ACE|D, M_k] + E[ACE|D, M_k]^2) * p(M_k|D) \\ &\qquad - E[ACE|D]^2 \end{aligned} \tag{14}$$

To compute the mean and variance of ACE, we need to compute the following:

1) A set K of DAGs over which the averaging is to be performed,
2) The posterior probability distribution $p(M_k|D)$ of the models in K, given the data, and

3) For each DAG $\in K$, an estimate of the mean of the ACE ($E[ACE|D]$) and the variance of ACE ($Var[ACE|D, M_k]$), given the data.

There are some difficulties associated with performing BMA and computing the above three parameters for a dataset D. Firstly, the number of models in the summation of equation 10 (set K) can be very large. For a DAG structure with N variables, the number of possible structures is super-exponential in N. Thus, summing over all possible DAGs is clearly impossible; approximations must be selected. Secondly, even if the cardinality of set K is reduced to a manageable size, computation of the integrals in equation 12 is not trivial and generally, computationally expensive. Thirdly, although computation of the expected value of the ACE is easy (comparatively, as will be mentioned in the coming sections), computation of variance is not straightforward. The next sections describe solutions to these problems.

7.7.2 Reducing the Cardinality of K

There are two strategies to reduce the cardinality of the set K. They can be roughly classified as pre- and post-data strategies. The first strategy, which can be performed even before the data are obtained, is to eliminate from the complete set of models, those models which contradict existing science. Thus, we eliminate infeasible models based on prior scientific evidence. For instance, consider a DAG with three nodes, PMR, Safety, and ADT. There are 25 possible Markov-equivalent graphs. Based on prior scientific evidence and logical reasoning, one can eliminate all those graphs in which Safety precedes ADT. One may argue that over time, Safety may influence ADT, but such feedbacks are rare and may represent the transient behavior of a segment (however, such loops can be incorporated in some graphs, but they must be in the form of dynamic Bayesian Networks or cyclic graphs; these graphs are out of the scope of this work). Such graphs which contradict scientific results can be eliminated by providing a temporal ordering of the variables. Other methods to systematically exclude such graphs from the set K would be to incorporate their absence in the form of prior knowledge. Thus, in score-based methods, the prior probabilities of such graphs, or the edges, that are deemed to be impossible, would be 0. In constraint-based methods, such graphs can be excluded by specifying constraints on the edges between variables. In the current study, we impose such restrictions by specifying a temporal ordering of the variables and exclude the edges that contradict the temporal order.

The second strategy of reducing the number of DAGs to be averaged in the set K, is carried out after the data have been observed. This strategy is based on the principle of Occam's razor. Models that predict the data far less than the best model available are excluded from the set K (i.e. models that do not belong to the set) as follows:

$$A' = \left\{ M_k \colon \frac{max_l\{p(M_l|D)\}}{p(M_k|D)} \leq C \right\} \tag{15}$$

Secondly, complex models that receive less support from the data than their simpler counterparts are also excluded. Thus, models belonging to the set:

$$B = \left\{ M_k \colon \exists\, M_l \in A, M_l \subset M_k, \frac{p(M_l|D)}{p(M_k|D)} > 1 \right\} \tag{16}$$

are excluded from the analysis. The models in set $K = A'/B$ are used to compute equation 10.

7.7.3 Computing Integrals

Once the set of models are identified, the next step is to compute the posterior probabilities of the models in set K. For discrete Bayesian Networks with Bdeu priors, and Dirichlet parameters, the posterior distribution of the parameters is a closed expression as noted in Section 7.3. The integrated likelihood and the posterior probability of the DAGs in set K can be obtained using Markov chain Monte Carlo (MCMC) techniques which are described in several places in the literature (Heckerman et. al., 1995; Hartemink, 2005).

The BANJO search procedure uses the log scale of posterior probability of a DAG as a score to search the model space of all possible DAGs. The scores are normalized with respect to the best DAG in the model and the posterior probability of each DAG is computed as per equation 17.

$$P(M_k|D) = \frac{\exp\,(S_k - \overline{S_k})}{\sum_{i=1}^{K} \exp\,(S_i - \overline{S_i})} \tag{17}$$

where $\bar{S} = \frac{1}{K} * \sum_{i=1}^{K} S_i$ and S_i is the log score of the DAG mode from the search procedure. The equation for the log score is given in Heckerman et. al. (1995).

7.7.4 Computing the Mean and Variance of Causal Effect

After identification of the models in set K and computing their posterior probabilities, the next step is to compute the posterior distribution (mean and variance) of the ACE using equations 13 and 14. Computing the posterior distribution of causal effects is not straightforward. Mean estimates of Causal Effects can be computed by averaging over the set of DAGs in K by replacing the parameters of the DAGs by mean values (Cooper and Herskovits, 1992). For a given DAG k in set $K,$ if ACE_k is the causal effect of interest (the dependence on data D is excluded for simplicity) and k has no latent variables, then ACE_k can be estimated by performing two queries in a mutilated graph k_x. Thus ACE_k is the ratio of two query variables in the mutilated graph k_x (section 7.4). Every query in a Bayesian Network is a function of the parameters Θ of the Bayesian Network and can be written as $q_k^i(\Theta)$ where Θ denotes the set of parameters of the CPT in a Bayesian Network, $i = 1,2$ for each query.

For example, if we are to estimate the ACE of changing PMR (X) from med to low, on safety (Y), $i = 1$, corresponds to manipulating PMR = low in k which is equivalent to introducing the evidence PMR = low in the mutilated graph k_x. Similarly, $i = 2$ corresponds to manipulating PMR = med. The ACE is estimated as the ratio of two queries in the mutilated graph k_{x}., given by the following equation:

$$ACE_{med\ to\ low} = \frac{P(Y = y|do(X = low)}{P(Y = y|do(X = med)} = \frac{q_{k_x}^{low}(\Theta)}{q_{k_x}^{med}(\Theta)} \qquad (18)$$

If the distribution of the query variables is known, then in principle, the distribution of ACE can be computed as the distribution of the ratio of the two query variables. For all query variables in a Bayesian Network it has been shown that (Cooper and Herskovits, 1992):

$$E[q_{k_x}(\Theta)] = q_{k_x}(E[\Theta]) \tag{19}$$

Equation (19) states that the expected value of the probability distribution of the q's is equal to the q's in the Bayesian Network parameterized by the expected values of the parameters. However, such a straightforward expression does not exist to estimate the variance of a query in a Bayesian Network. It was conjectured that the query variable in a Bayesian Network with Dirichlet parameters, Bdeu priors and fully observed data would be Beta distributed (Kleiter, 1996). But recently, this claim was proved to be valid only for special cases and incorrect, in general (Hooper, 2007). Allen et al. (2008) provide an asymptotic approximation to the distribution of any query in a Bayesian Network. They use the delta method and propagate the sensitivity of the query distribution to each parameter in the network, approximating the variance of the query by a Beta distribution. The algorithm by Allen et al., with a few modifications as per Castillo (1997) is used in the present work to compute the variance of the query variables for each network in K. The details are provided in Appendix B.

As mentioned, the distribution of the ACE's from each DAG can be computed using the ratio of the distribution of the query variables. Since the query variables are approximated by Beta distributions, the ACE for each DAG is distributed according to the ratio of two beta distributions. Although closed-form expressions for the ratio of two generalized beta distributed variables are available (Pham-Gia, 2000) we use simulation to estimate the variance to ease further computation. (We need to estimate a mixture of ratios of beta variables, as per equation 13 and 14. The analytical form of the beta ratios requires numerical integration and mixing such distributions further complicates the analysis; hence, we use simulation).

The posterior distribution of the query variables for each DAG in K are approximated using beta distributions as explained in Appendix B. To evaluate the distribution of the ACE for each DAG, a sample of 100,000 is simulated from the approximate beta distributions of corresponding queries and the ACE is computed using the following equations:

$$ACE^k_{high\ to\ low} = \frac{P(crash\ =\ 1|do(visbility = low)}{P(crash\ =\ 0|do(visibility = high)}$$

$$ACE^k_{high\ to\ med} = \frac{P(crash\ =\ 1|do(visbility = med)}{P(crash\ =\ 0|do(visibility = high)} \tag{20}$$

$$ACE^k_{med\ to\ low} = \frac{P(crash\ =\ 1|do(visbility = low)}{P(crash\ =\ 0|do(visibility = med)}$$

The sample from each DAG is then weighted according to the posterior probability of the DAG $p(M_k|D)$. The mean and variance are computed from this weighted sample per equations 13 and 14. The posterior credible set in which the ACE is expected to lie with probability 0.95 is estimated. The next section provides the results of the analysis.

Chapter 8 Results

The results of the analysis are presented in this section. The structure of the best graph recovered by the structure learning algorithm is shown in Figure 8-1. A total of 10 graphs were used to compute the ACE's using Bayesian Model Averaging (BMA). Table 6 shows the structure of all 10 graphs. In all of the graphs, safety has a probabilistic dependence on PMR. This shows that PMR levels indeed have an influence on the occurrence of a crash. The strength of this relation is specified by the estimate of the ACE parameters.

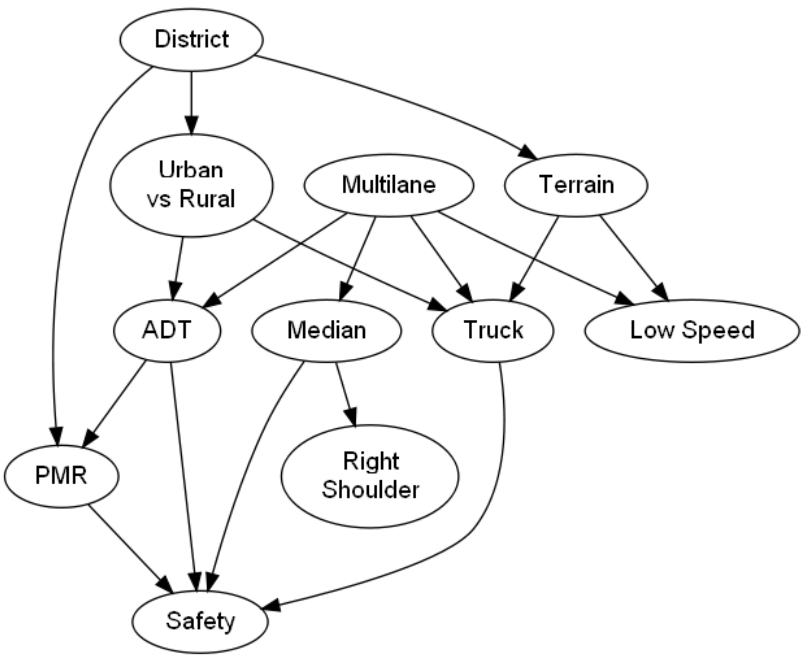

Figure 8-1: Best Network recovered from the Data

Table 6 Best 10 Graphs used for BMA

	Parents of V										
V	D	T	M	Ur	Speed	Med	Shld	ADT	Trk	PMR	Safety
1	-	D	-	D	T, M	M	Med	M,Ur	T , M, Ur	D , ADT	Med,ADT,Trk,PMR
2	-	D	Ur	D	T, M	M	Med	M,Ur	T , M, Ur	D , ADT	Med,ADT,Trk,PMR
3	-	D	-	D	T, M	M	Med	M,Ur	T , M, Ur	D , ADT	M,ADT,Trk,PMR
4	-	D	Ur	D	T, M	M	Med	M,Ur	T , M, Ur	D , ADT	M,ADT,Trk,PMR
5	-	D	-	D	T, M	M	Med	M,Ur	T , M, Ur S	D , ADT	Med,ADT,Trk,PMR
6	-	D	Ur	D	T, M	M	Med	M,Ur	T , M, Ur S	D , ADT	Med,ADT,Trk,PMR
7	-	D	-	D	T, M	M	Med	M,Ur	T , M, Ur S	D , ADT	M,ADT,Trk,PMR
8	-	D	Ur	D	T, M	M	Med	M,Ur	T , M, Ur S	D , ADT	M,ADT,Trk,PMR
9	-	D	-	D	T, M	M	Med	M,Ur	T , M, Ur	D , ADT	M, Med,ADT,Trk,PMR
10	-	D	Ur	D	T, M	M	Med	M,Ur	T , M, Ur	D , ADT	M, Med,ADT,Trk,PMR

D = District; T = Terrain; M = Multilane
Ur = Urban vs Rural; Speed = Low Speed; Med = Median Presence
Shld = Right Shoulder; Trk = Percentage of Trucks

The other relationships discovered by the graph are also intuitive. The amount of vehicular traffic flowing on the road depends on the number of lanes and the location of the segment (urban vs rural). District location does not have a direct influence on the traffic volume on a segment. This indicates that given the urban/rural nature of the segment, the ADT level is independent of the district location of the segment. Similarly, the percentage of trucks is influenced by the number of lanes and the urban/rural nature of the segment; given these variables, it is independent of the district. The amount of truck traffic also depends on the terrain of the segment, whether the segment is on rolling terrain or not. But the terrain variable has no influence on ADT, indicating that the traffic volume on a segment does not depend on the grade. The speed limit on a segment depends on the number of lanes and the terrain. The PMR variable depends on district as well as ADT. The dependence on district is due to the fact that there are weather differences in the three North Carolina districts. The weather variable lies on the causal pathway between district and PMR but is not observed. Other variables captured by district could be differences in the PMR application process. Presence of a median is influenced by the number of lanes; median presence influences the shoulder width as well as the safety of the segment. Safety is also affected by the traffic volume, the percentage of trucks and the PMR levels. It is surprising to see that the number of lanes, shoulder width, and the speed limit has no direct influence on the safety of a segment. This could be due to the fact that, in the current sample, segments with low speed limits generally do not have a median and hence the effect of speed limit is captured by the presence or absence of medians.

In some of the other networks generated by the algorithm, safety of a segment is affected by the number of lanes. But, in some of these networks (2,3,6,7) the dependence of safety on number of lanes renders it independent with the presence of median. Again, this could be due to the fact that multi-lane roads are generally associated with medians, thus the effect of one variable is being captured by the other.

As noted earlier, the conditional probability table (CPT) of a Bayesian Network is exponential in the number of parents and the classes of each parent. For example, in the top scoring graph, the variable safety has 4 parents, each of which has 2 states. Thus the complete CPT of the node safety has 2^4 free parameters. Hence, quantitative information for all the nodes is difficult to provide. However, we provide insights into quantitative relations between some of the variables for the top scoring graph. The full CPT of this graph is provided in Appendix C. The graph can be reconstructed in any Bayesian Network software such as Genie (Druzdzel, 1999), SamIam (SamIam, 2005), and others (Murphy, 2005), and inference can be performed by using the CPT provided. It must be noted that the results may be sensitive to the discretization policy used, and this sensitivity was not tested for all the variables except for PMR. Also, these values represent point estimates and do not provide complete information about the effect of a variable on the other. Unfortunately, computing the variances is intensive, hence the variance of these queries has not been provided.

The probability of a crash, when the ADT is high is 0.1682 and the probability of a crash, when the ADT is low is 0.0370. Thus the chances of a target crash increase by almost 4.5 times when the level of ADT increases. Similarly, the probability of a target crash is 0.0627 when a median is present and 0.2937 when a median is absent. Thus, in general, medians reduce the probability of target crashes by 78 percent. It must be noted that this a marginal value of the effect of median on safety. For a segment with low ADT and a median, the probability of a target crash is 0.0349. For a segment with low ADT and no median, the probability of a target crash is 0.0516. Thus, even for a fixed value of ADT, medians reduce the chances of target crashes. On the other hand, for a segment with a median, the probability of a target crash when ADT is low is 0.0349 and the probability of a target crash when ADT is high is 0.2937, indicating that higher ADT increases the probability of target crashes. The interpretation of such results is limited due to the discretization and very few classes (less than or equal to 3) for each variable.

This is a result of sparse safety data. We now present the main results: the effect of PMR on safety.

As discussed in section 7.4, the causal effect can be estimated by performing inference in the mutilated graph. The mutilated graph for the manipulation of PMR obtained from the top scoring graph of Figure 8-1 is shown in Figure 8-2.

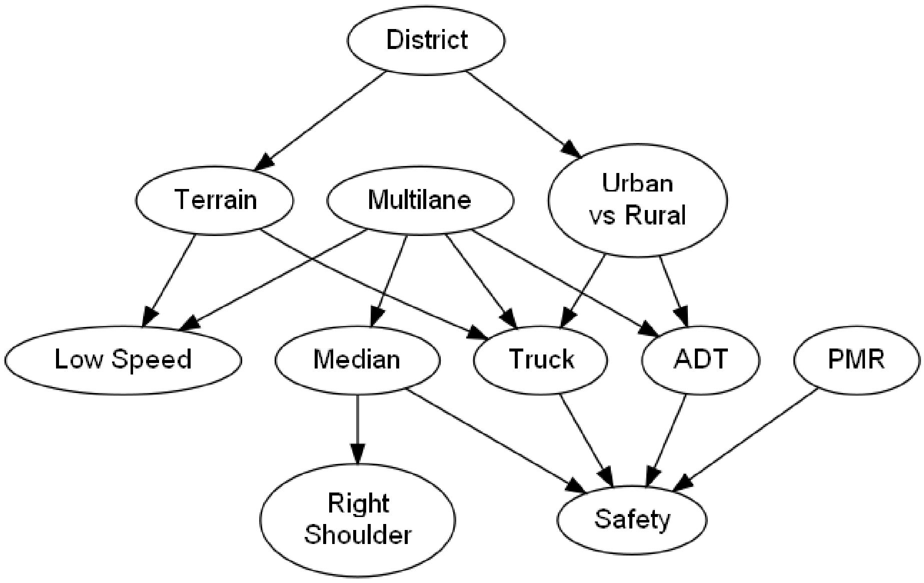

Figure 8-2: Manipulated Graph

All of the edges entering into the PMR node are deleted and the value of PMR is set to a fixed value, representing the manipulation of PMR. Thus, in Figure 8-1, the edges between PMR and nodes of ADT and district are deleted giving rise to the mutilated graph of Figure 8-2. Nature assigns a value of PMR level to each segment based on the characteristics of ADT and location of the segment in a district. Nature has a finer assignment model than what is indicated here; for example, the weather conditions are also evaluated before assigning the PMR level. As complete information about this model is not available, uncertainties are represented in the form of a probabilistic network. The probabilistic network itself is also uncertain; hence, we employ the BMA method to average over several plausible network structures. Nevertheless, the following assumption is made. For the finer model that nature has (of which we lack information), there are no attributes of a segment that were used to assign the PMR levels that affect the safety of that single segment. This has a direct correspondence to the hypothetical experimental design introduced in section 6.2. Thus, by deleting the nodes that enter the PMR node, the influence on PMR levels due to the other attributes of a segment are eliminated, forcing all of the segments to take a fixed value, irrespective of their other attributes. The corresponding concept in experimental design is that of a randomized setting. The treatment variable is assigned randomly to a population ensuring that there is no influence on the assignment by any other covariates, thus eliminating potential confounders. In the current context, all segments are forced to have a fixed PMR level (say low) and estimating the average number of segments that contain target crashes. Next, all of the segments are forced to take a different (but fixed) value of PMR (say high) and estimating the average number of segments that experience target crashes. The ratio of these two values gives the average causal effect (ACE). The average value of the causal effects and their confidence intervals, averaged over all 10 graph structures and all the discretization policies is shown in Table 5. The values in Table 7 represent the posterior values of the ACE of PMR on safety, taking into account the uncertainty of the parameters, the DAG structure, as well as the discretization of the PMR variable into three levels. The discretization bins for these ACE's is taken to be equal to the average of the bin boundaries of all the discretization policies considered. Thus from Table 5, PMR = low $if\ 0 < Retro \leq 207.6\, mcd/m^2/lux$

PMR = med $if\ 207.6 < Retro \leq 350\, mcd/m^2/lux\ and$

PMR = high $if\ Retro > 350\, mcd/m^2/lux$

From Table 7, it can be seen that a change in the PMR level from high to low increases the chances of a target crash, on average, by 87 percent. Similarly, a change in the PMR levels from medium to low increases the chances of a target crash by 66 percent. On the other hand, the effect of changing PMR levels from high to medium on average increases the target crash probability by 11 percent. But the posterior interval contains the ACE value of 1 as well. This suggests that sufficient evidence does not exist in the present sample to conclude any influence of PMR change from high to medium on safety.

.

Table 7 Estimate of ACE - Mean and Credible set

Change in Visibility Level	Point Estimate	95 % limits
High to Low	1.87	[1.65,2.14]
Medium to Low	1.66	[1.48,1.88]
High to Medium	1.11	[0.95, 1.30]

Figure 8-3 shows the sensitivity of the ACE to the PMR bins. The figure illustrates that the sensitivity of the ACE in general increases when the PMR level (upper bin) falls below 160. This is due to the fact that at such low levels, the number of segments reduces and hence the causal effects are not reliable. This was one of the reasons why the lower bin of the upper level was not tuned. If the sensitive bins are excluded, the average effect of change in PMR from high to low levels on target crashes ranges from 1.2 to 2. Similarly, the average effect of a change in PMR from medium to low levels ranges from 1.1 to 1.7. The average effect of a change in PMR from high to medium levels remains fairly constant at 1.1 and most of the confidence intervals cover the value 1, thus indicating that there may not be a change in the probability of target crashes by changing PMR levels from high to medium.

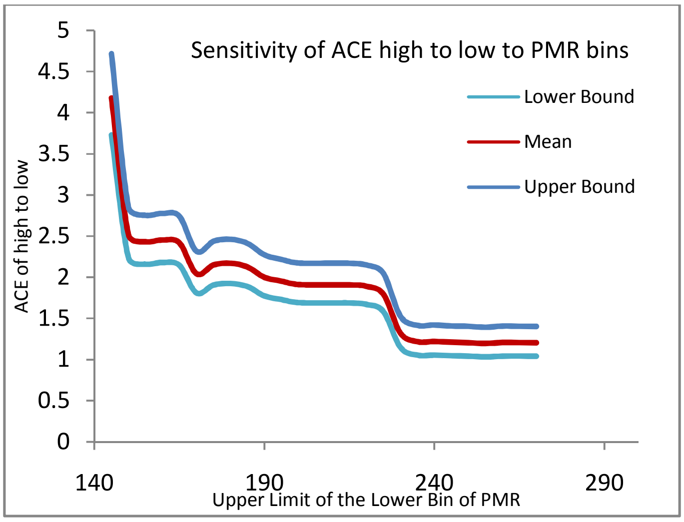

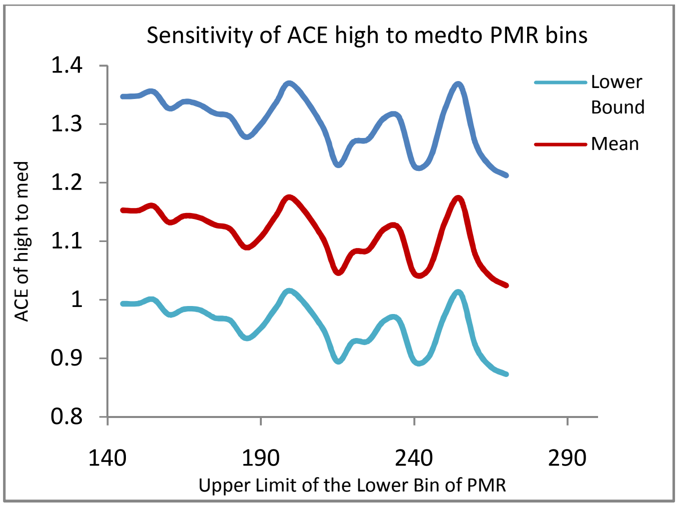

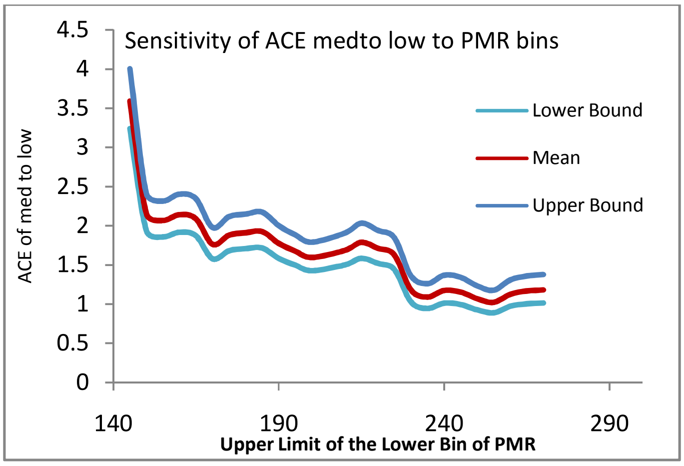

Figure 8-3: Sensitivity of ACE to PMR bins

Chapter 9 Conclusions and Future Work

This work is an exploratory analysis that examines the applicability of causal inference methods in transportation safety. The relationship between PMR levels and crash frequency was considered and modeled using a causal Bayesian network framework to illustrate an application of this proposed analytical method. The results are plausible and consistent with engineering intuition, suggesting that high and medium levels of PMR both produce lower crash frequency probabilities than low levels of PMR. A practical limitation of the present study is that PMR levels were all relatively high (even in the "low" category) when compared to the proposed minimum levels for inclusion in future versions of the MUTCD (DeBallion et al., 2007). It is possible that including lower levels of PMR in future studies may produce different results, and therefore, additional research is required to investigate this possibility using data from a variety of states and for a variety of pavement marking material types (only durable thermoplastics were included in the present sample).

Though the mechanisms that generate crashes are complex, this work shows that there is considerable promise in applying Causal Models to transportation safety studies. Particularly important is to take into account the assignment mechanism of the treatment variable of interest. Bayesian networks provide an efficient and intuitive way of incorporating expert knowledge and combing it with knowledge from data.

In this work, the assignment mechanism of the treatment variable was fairly straightforward. The PMR levels depend on the age of the markings, the ADT on the roadway segment, and the weather conditions (based on the district indicators). This knowledge was incorporated into the Bayesian Network and the structure of the network was recovered from observational data. Uncertainty in parameters of the network and the structure of the network itself was incorporated in the Average Causal Effect by using Bayesian Model Averaging.

Future studies should examine the feasibility of causal inference methods to other road safety treatments or countermeasures. For instance, treatment assignments related to geometric design elements of a segment could be quite complex and are an area of research offering vast potential to further explore the use of causal inference in road safety. As an example, the selection criteria of a particular type of median barrier on a segment could depend on several factors such as past accident history, availability of right-of-way, cost of installation, median cross-section design, median barrier warrant criteria, engineering judgment, etc. In the case of subjective judgments (e.g., accident history), or when roadway data are difficult to collect or quantify (e.g., median cross-slopes), it may be difficult to model the correct assignment mechanism. Advancement can be made by introducing latent variables and examining the appropriate set of covariates using Pearl's (2000) graphical criteria.

Causal diagrams also provide ways to incorporate causal knowledge from past research, both generic and specific, using causal grammar. An example of inferring qualitative causal knowledge from past studies is provided in Appendix D using human factors as an example. Human factors data are an important consideration in safety research that could not be explored in the present study because such data are difficult to collect in retrospective safety studies for all drivers traversing a roadway segment during the analysis period. However, care must be taken before including any causal relations in a Bayesian network, and their validity must be examined by reviewing the relationships with domain experts.

The Bayesian Network framework is quite powerful in representing causal knowledge but computationally expensive. It also imposes several restrictions on the nature of the data. For instance, in the current study, all of the variables were discretized to enable efficient inference and learning in the Bayesian Network. Bayesian Networks can handle mixed distributions, but the field of research is comparatively new and few algorithms exist to handle Bayesian Networks with arbitrary distributions. Alternative methods of Causal Inference such as the potential outcomes model do not have such limitations. However, representation of causal knowledge in an intuitive way, and elicitation of causal relations from experts, is not easy in the potential outcomes framework. An amalgam of both methods is promising and is the direction in which future work of Causal Inference in road safety studies can be directed.

Appendix A: PC Algorithm

In absence of hidden variables, using the PC algorithm, the structure of the DAG can be reconstructed from queries of conditional independencies, assuming that those conditional independencies reflect d-separation conditions in some unknown graph structure S. If all of the common causes are included in the dataset and if a node is independent of its nondescendents, given its parents, then any conditional dependency arising in the dataset can be assumed to be due to a cause and effect relationship between the observed variables. The faithfulness assumption ensures that the conditional independencies present in the dataset are only due to the Causal Markov condition and that no additional independencies arise due to chance.

A detailed description of the algorithm is provided in "Causation, prediction and search" (Spirtes et. al., 1993). A brief introduction to the algorithm is provided below along with the pseudo code of the procedure. The algorithm takes as an input a dependency model M (list of dependences between variables in a domain V) and gives as an output a directed acyclic graph G representing the conditional independencies found in M (or a pattern in case of observational data, which is the case in the present study). The following properties are required for this algorithm.

A dependency model M is said to be isomorphic to a DAG G, iff (Spirtes et. al., 1993):

a) For each pair of nodes X and Y in G, x and y are adjacent iff x and y are conditionally independent given every subset of vertices in G (excluding X and Y)

b) For each triplet of vertices x, y, z such that x and y are adjacent and y and z are adjacent but x and z are not, $x \rightarrow y \leftarrow z$ is a subgraph of G iff x, y and z are conditionally independent given the set of all variables in G excluding y but not x and z.

The PC algorithm works by asking a conditional independence oracle to make judgments about the independence relation between pairs of variables (e.g. X and Y) conditioned on sets of variables (e.g. Z). Conditional independence tests are available for datasets that consist either entirely of continuous variables or entirely of discrete variables. The algorithm can be divided into two stages. In the first stage, the algorithm recovers a complete undirected graph from the dataset based on conditional independence judgments. This phase is called the fast adjacency phase. In the adjacency phase, a complete undirected graph over the variables is initially constructed and then edges $X - Y$ are removed if some set S among either the adjacents of X or the adjacents of Y can be found (of a certain size, or "depth") such that X and Y are independent of each other, given S.

Once the adjacency structure over V has been well estimated by this procedure, an orientation phase is begun. The first step of the orientation phase is to examine unshielded triples and consider whether to orient them as colliders. An unshielded triple is a triple $< X, Y, Z >$ where X is adjacent to Y, Y is adjacent to Z, but X is not adjacent to Z. Since X is not adjacent to Z, the edge $X - Z$ must have been removed during the adjacency search by conditioning on some set $S_{x\,z}$; $< X, Y, Z >$ is oriented as a collider $X \rightarrow Y \leftarrow Z$ just in case Y is not in $S_{x\,z}$. Once all such unshielded triples have been oriented as colliders by this rule, a series of orientation rules is applied (in this case, the complete orientation rule set from Meek [1995]) to orient any edges whose orientations are implied by previous orientations. Below is the pseudo code of the algorithm. In the following, $ad(x)$ is the set of adjacent nodes for node x and the *"/"* operator represents the difference operation on sets.

1. Create a complete graph G on the variables in V
2. $N = 0$
3. Repeat until $Ad(x)\backslash\{y\} < n$ for each set of ordered pairs (x, y)

   3.1.Repeat until all ordered pairs of adjacent variables (x, y) such that $Ad(x)/\{y\} >= n$ and every subset S in $Ad(x)\backslash\{y\}$ have been tested for independence.

   Select an ordered pair of variables x, y adjacent in G such that $|ad(x)\backslash\{y\}| < n$

   Select a subset S of $Ad(x)\backslash\{y\}$ with cardinality n

   If $I(x|S|y)$, then erase $x - y$ from G. Store S in the sets separating (x, y) and separating (y, x)

   3.2.$N = n + 1$
4. For each triplet of nodes x, y, z where x and y are adjacent and y and z are adjacent but x and z are not adjacent, orient $x \rightarrow y \leftarrow z$ if and only if y does not belong to separating (x, z)
5. Repeat until no more arcs could be oriented

   5.1.If the structure $x \rightarrow y - z$ belongs to G, where x and z are not adjacent and no head to head arcs point y, orient $y - z$ as $y \rightarrow z$

   5.2.If there exists a directed path from x to y and the arc $x - y$ exists, turn it into $x \rightarrow y$

Appendix B: Variance in a Bayesian Network

This Appendix describes the algorithm used to estimate the variance of any query in a discrete Bayesian Network. As noted previously, every query in a Bayesian Network is a function of the parameters Θ of the Bayesian Network and can be written as $q_{h|e}^{x|pa}(\Theta)$ where x is any variable in the Bayesian Network and Pa is the set of parents of X, Θ denotes the set of parameters of the CPT in Bayesian Network and $h|e$ is the query with evidence e. We drop the notation $h|e$ from here on, for simplicity. Also, for each network parameter $\Theta_{x|Pa} \in \Theta$, we can consider the partial derivative of the query with respect to that parameter:

$$q^{x|pa}(\theta) = \left.\frac{\partial q(\Theta)}{\partial \Theta_{x|pa}}\right|_{\theta} \tag{21}$$

The following is a restatement of a Theorem from Allen et al. (2008), required to compute the asymptotic approximation to the variance of a query in a Bayesian Network.

Given the "independent Dirichlet property with respect to a given DAG" assumption, the posterior mean of $q(\Theta)$ *is* $E[q(\Theta)] = q(E[\Theta])$. Asymptotically, the variance of the query is given by:

$$\sigma^2 = \sum_{\mathrm{x,Pa}=pa} \frac{1}{1 + m_{x|pa}} * v(x|pa) \tag{22}$$

$$v(x|pa) = \left(\sum_{x \in X} q^{(x|Pa=pa)}(\Theta)^2 * \Theta_{x|Pa=pa} \right) - \left(\sum_{x \in X} q^{(x|Pa=pa)}(\Theta) * \Theta_{x|Pa=pa} \right)^2 \tag{23}$$

Under this framework, the standardized random variable

$$\frac{q(\Theta) - q(\acute{\Theta})}{\sigma}$$

converges to the standard Normal distribution $N(0,1)$.

Here, $m_{x|pa}$ is the pseudo count of the variable-parent combination $x - pa$. The summation in equation (22) is over each level of variable X. The summation in equation (23) is over all rows in a CPT. The mean μ and variance σ^2 of the Normal distribution can be approximated by a beta distribution with parameters $\alpha\ and\ \beta$ using the following equations:

$$\alpha = \mu \frac{\mu(1 - \mu) - \sigma^2}{\sigma^2} \tag{24}$$

$$\beta = (1 - \mu) \frac{\mu(1 - \mu) - \sigma^2}{\sigma^2} \tag{25}$$

Equations (22) and (23) show that σ^2 adds up the influence of each $\Theta_{x|pa}$ on the query $q(h|\ e)$, which is based on the derivative $q^{(x|pa)}\ (\Theta)$, weighted by $(1 + m_{x|pa})^{-1}$. However, the number of terms to be added can be reduced. Not all parameters of the CPT of Bayesian Network contribute to the variance term. The algorithm by Castillo (1997) is used to identify the parameters sufficient to compute the derivative of the query.

The algorithm takes as input a Bayesian Network $(G;\ CPT)$ and two sets of nodes: the target node, H, and the set of evidence nodes E and gives as output the set of relevant nodes V needed to compute the derivative terms of the variance.

Step 1: Construct a DAG G_0 by augmenting G with a dummy node V_i and adding a link $V_i \rightarrow X_i$ for every node X_i in D.

Step 2: Identify the set V of dummy nodes in G_0 not d-separated from H by E. The node V_i represents the parameters, Θ_i, of node X_i. Step 2 of Algorithm 1 is carried out in linear time using an algorithm provided by Geiger et al. (1990).

Now, we need to compute the derivative of the query with respect to the parameters identified by the algorithm above. The following theorem is required.

Let B be a Bayesian network, θ be a probability parameter, h be a query, and e be the evidence entered into B. The posterior probability $p(h|e)(\theta)$ is a fraction of two linear functions of θ:

$$p(h|e)(\theta) = \frac{\alpha_1\theta\ + \beta_1}{\gamma_1\theta\ +\ \delta} \tag{26}$$

For simplicity, the function can be normalized as follows:

$$p(h|e)(\theta) = \frac{\alpha\theta\ + \beta}{\gamma\theta\ +\ 1} \tag{27}$$

Then the partial derivative of $p(h|e)(\theta)$ on θ can be expressed as:

$$\frac{\partial p(h|e)}{\partial\theta} = \frac{\alpha - \beta\gamma}{(\gamma\theta + 1)^2} \tag{28}$$

To determine the value of α, β, and γ, three message propagations are used, for each given evidence. In the present case, junction tree inference with $\theta's$ set as 0, 0.5, and 1, are used. The values of the coefficients in equation (28) are determined as the following:

$$\beta = p^0$$
$$\gamma = \frac{\beta - p^{0.5}}{p^{0.5} - p^1} - 1 \tag{29}$$
$$\alpha = p^1(\gamma + 1) - \beta$$

where $p^0, p^{0.5}$, and p^1 denote the posterior probabilities of $p(h|e)(\theta)$, when $\theta = 0, 0.5, 1$, respectively. Proportional scaling is used to change the related parameters in such a way that they keep the original proportion. The derivatives are computed using equations (28) and (29) and the variance of the query is computed using equations (22) and (23). The mean and variance are then approximated to a beta distribution using equations (24) and (25)

Appendix C: Parameters of the Best Network

CPT Tables for the Best Graph

CPT for Safety

Parents				Counts		Probability	
Median	Truck	ADT	PMR	Safety = 0	Safety = 1	Safety = 0	Safety = 1
Yes	Low	Low	Low	474.021	10.021	0.979	0.021
No	Low	Low	Low	0.021	0.021	0.500	0.500
Yes	High	Low	Low	279.021	50.021	0.848	0.152
No	High	Low	Low	0.021	0.021	0.500	0.500
Yes	Low	High	Low	171.021	46.021	0.788	0.212
No	Low	High	Low	0.021	0.021	0.500	0.500
Yes	High	High	Low	1613.021	144.021	0.918	0.082
No	High	High	Low	0.021	0.021	0.500	0.500
Yes	Low	Low	Med	1094.021	12.021	0.989	0.011
No	Low	Low	Med	813.021	20.021	0.976	0.024
Yes	High	Low	Med	162.021	5.021	0.970	0.030
No	High	Low	Med	16.021	0.021	0.999	0.001
Yes	Low	High	Med	149.021	33.021	0.819	0.181
No	Low	High	Med	0.021	0.021	0.500	0.500
Yes	High	High	Med	459.021	42.021	0.916	0.084
No	High	High	Med	0.021	0.021	0.500	0.500
Yes	Low	Low	High	151.021	2.021	0.987	0.013
No	Low	Low	High	29.021	0.021	0.999	0.001
Yes	High	Low	High	27.021	1.021	0.964	0.036
No	High	Low	High	0.021	0.021	0.500	0.500
Yes	Low	High	High	31.021	6.021	0.837	0.163
No	Low	High	High	0.021	0.021	0.500	0.500
Yes	High	High	High	72.021	5.021	0.935	0.065
No	High	High	High	0.021	0.021	0.500	0.500

CPT for PMR

Counts			Probability		
Retro = Low	Retro = Medium	Retro = High	Retro = Low	Retro = Medium	Retro = High
2787.333	2805.3333	324.333	0.47107	0.47411	0.0548

CPT for ADT

Parents		Counts		Probability	
Urban	Multilane	ADT = Low	ADT = High	ADT = Low	ADT = High
Yes	Yes	1568.125	0.125	0.999	0.0001
No	Yes	527.125	2501.125	0.174	0.8259
Yes	No	240.125	240.125	0.500	0.5000
No	No	810.125	30.125	0.964	0.0359

CPT for Truck

Parents			Counts		Probability	
Urban	Multilane	Terrain	Truck = Low	Truck = High	Truck = Low	Truck = High
Yes	Yes	Flat	690.06250	0.06250	0.99991	0.00009
No	Yes	Flat	177.06250	1049.06250	0.14441	0.85559
Yes	No	Flat	480.06250	0.06250	0.99987	0.00013
No	No	Flat	810.06250	30.06250	0.96422	0.03578
Yes	Yes	Rolling	862.06250	16.06250	0.98171	0.01829
No	Yes	Rolling	22.06250	1780.06250	0.01224	0.98776
Yes	No	Rolling	0.06250	0.06250	0.50000	0.50000
No	No	Rolling	0.06250	0.06250	0.50000	0.50000

CPT for Rt. Shoulder

Parents	Counts		Probability	
Median	Rt. Shoulder = Low	Rt. Shoulder = High	Rt. Shoulder = Low	Rt. Shoulder = High
Yes	450.250	4588.250	0.089	0.911
No	788.250	90.250	0.897	0.103

CPT for Low Speed

Parents		Counts		Probability	
MultiLane	Terrain	LowSpeed = Yes	LowSpeed = No	LowSpeed = Yes	LowSpeed = No
Yes	Flat	690.125	1226.125	0.360	0.640
No	Flat	510.125	810.125	0.386	0.614
Yes	Rolling	698.125	1982.125	0.260	0.740
No	Rolling	0.125	0.125	0.500	0.500

CPT for Median

Parents	Counts		Probability	
Multilane	Med = Yes	Med = No	Med = Yes	Med = No
Yes	3718.2500	878.2500	0.8089	0.1911
No	1320.2500	0.2500	0.9998	0.0002

CPT for Terrain

Parents	Counts		Probability	
District	Terrain = Flat	Terrain = Rolling	Terrain = Flat	Terrain = Rolling
1	1290.1667	0.1667	0.9999	0.0001
5	1214.1667	2072.1667	0.3695	0.6305
12	732.1667	608.1667	0.5463	0.4537

CPT for Urban

Parents	Counts		Probability	
District	Urban = Yes	Urban = No	Urban = Yes	Urban = No
1	480.1667	810.1667	0.3721	0.6279
5	1568.1667	1718.1667	0.4772	0.5228
12	0.1667	1340.1667	0.0001	0.9999

CPT for Multilane

Counts		**Probability**	
Multilane = Yes	Multilane = No	Multilane = Yes	Multilane = No
4596.5	1320.5	0.7768	0.2232

CPT for District

Counts			Probability		
Dist = 1	Dist = 5	Dist = 12	Dist = 1	Dist = 5	Dist = 12
1290.333	3286.333	1340.333	0.2181	0.5554	0.2265

Appendix D: An Example of Incorporating Causal Knowledge from Past Research Studies

This appendix shows how causal knowledge from past studies can be summarized in a causal diagram. Such exercises are useful not only to amalgamate existing knowledge in one place, but also in designing new studies. Indeed, the causal diagram can serve as a guidance tool to list the variables for which data need to be collected and the variables which need to be analyzed as latent. The results of the human factors literature review from section 3.3 are summarized in the form of a graph, with several additions from intuition.

The literature review provides the following evidence. Firstly, the behavior of drivers is different on curves and tangent sections. Driver behavior is captured using the vehicle speed, for simplicity. (In actual studies, the lane position could also be used). Hence, the horizontal curve radius is an important variable to be considered. The speed selected by a driver could also depend on the lane width and the shoulder width along a roadway segment, among other geometric features. Also, there is evidence that the PMR levels on a segment may influence the speed of a driver. This influence can be included by including the preview distance (or end detection distance, the distance that a driver can see) on the causal pathway between PMR and speed. The preview distance depends on the amount of light retroreflected back to the driver (PMR) and also on the contrast sensitivity of the driver (which may be a function of age, as well as driving experience). Also, based on the literature review, preview time is an important variable to capture the effect of speed on detection distance. The preview time is simply the time a driver has to preview the road ahead and make a decision accordingly. The decision can be to steer the vehicle, apply brakes, or change lanes, etc. Thus, preview time is determined by the preview distance available and the speed of the vehicle. The preview time denotes the driver performance demanded by the environment. This is the amount of attention a driver must allocate to the driving task. On the supply side, driver performance can be measured by the reaction time of a driver. The reaction time depends on the age, experience and secondary driving tasks of a driver. The safety or the risk of an accident depends on the driver performance and the environmental demand. These relationships are shown in the form of a DAG in Figure D-1. This is a very simplified version to illustrate the application of causal diagrams.

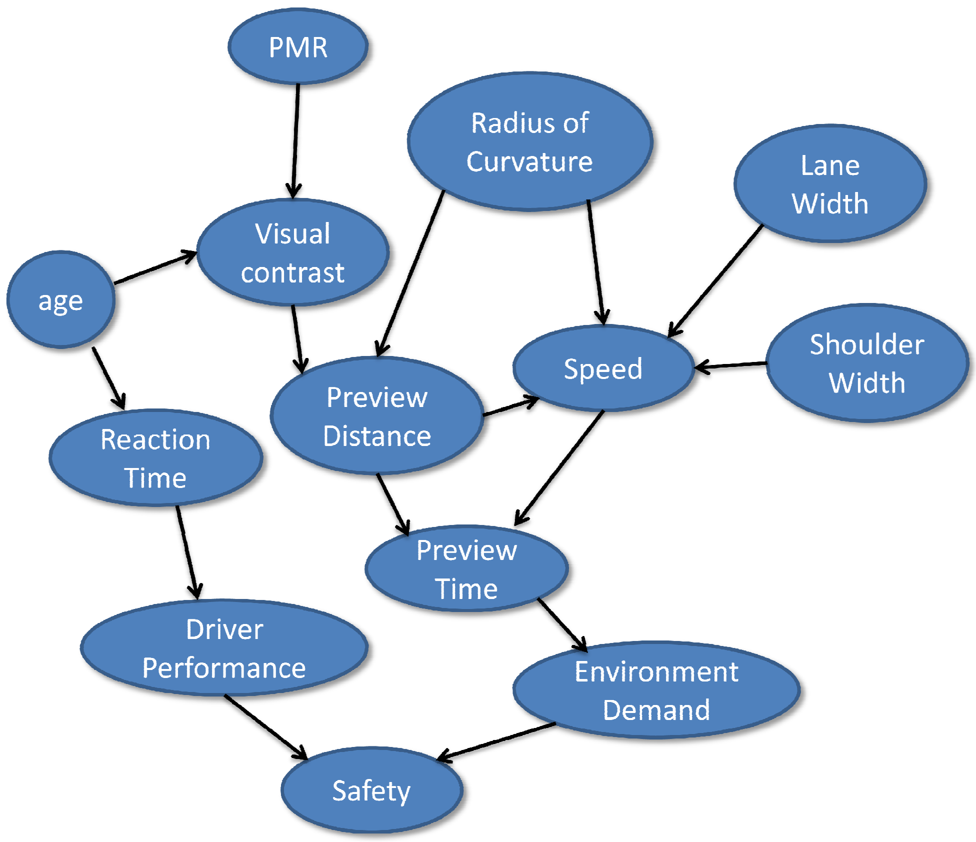

Figure D-1: DAG of Combined Human Factors and PMR Effects on Safety

References

Abboud, N., and Bowman B. L. (2002) "Establishing a Crash-based Retroreflectivity Threshold," Proceedings of the 81st Annual Meeting of the Transportation Research Board, Washington, DC.

Allen, V. T., Singh, A., Greiner, R. and Hooper, P. M. (2008) "Quantifying the uncertainty of a belief net response: Bayesian error-bars for belief net inference." Artificial Intelligence, Vol 172, pp. 483-513.

Al-Masaeid, H. R., and Sinha, K. C. (1994) "Analysis of Accident Reduction Potentials of Pavement Markings," In ASCE Journal of Transportation Engineering, Vol. 120, No. 5.

American Association of State Highway and Transportation Officials, National Transportation Product Evaluation Program. (2003) "Project Work Plan for Field and Laboratory Evaluation of Pavement Marking Materials," Found at: [http://www.ntpep.org].

Austin, R.L., and Schultz, R.J. (2002) "Guide to Retroreflection Safety Principles and Retroreflective Measurements," Gamma Scientific, San Diego, California.

Bahar, G., Masliah, M., Erwin, T., Tan, E., and Hauer, E. (2006) “Pavement marking materials and markers: real-world relationship between retroreflectivity and safety over time,” NCHRP Web-only Document 92, TRB, Washington, DC.

Carlson, P.J., Miles, J.D., Pratt, M.P., and Pike, A.M. (2005) “Evaluation of Wet-Weather Pavement markings: First Year Report,” Research Report No. 0-5008-1, Texas Transportation Institute, Texas A&M University, College Station, Texas.

Castillo, E., Guti´errez, J. M., and Hadi, A. S. (1997) “Sensitivity analysis in discrete Bayesian networks,” IEEE Transactions on Systems, Man, and Cybernetics, Vol. 27, pp. 412–423.

Cooper, G., and Herskovits, E. (1992) "A Bayesian method for the induction of probabilistic networks from data," Machine Learning, Vol 9, pp 309–347.

Cottrell, B. H. and Hanson, R. A. (2001) “Determining the Effectiveness of Pavement Marking Materials,” Report No. VTRC 01-R9 Virginia Transportation Research Council, Charlottesville, VA.

Cowell, R. (1998) “Introduction to inference for Bayesian Networks,” In Proceedings of the NATO Advanced Study institute on Learning in Graphical Models (Erice, Italy). M. I. Jordan, Ed. Kluwer Academic Publishers, Norwell, MA, 9-26.

David, M., Andersson, S., Perlman, M., and Volinsky, C. (1996) "Bayesian Model Averaging And Model Selection For Markov Equivalence Classes Of Acyclic Digraphs," Communications in Statistics: Theory and Methods, pp 2493-2519.

DeBallion, C., P. Carlson, Y. He, T. Schnell, and F. Aktan. (2007) "Updates to Research on Recommended Minimum Levels for Pavement Marking Retroreflectivity to Meet Driver Night Visibility Needs," Report No. FHWA-HRT-07-059, McLean, VA.

Dewar, R. E. (2002) "Roadway Design," Chapter 12 in Human Factors in Traffic Safety (Tuscon, AZ, Lawyers & Judges Publishing Company, Inc., pp. 381-420.

Dickinson, A., and Shanks, D. (1995) "Instrumental action and casual representation," In D. Sperber, D. Premack and A. Premack, Causal cognition. Oxford, England, Oxford University Press.

Druzdzel, M.J. (1999) "Genie: a development environment for graphical decision-analyticmodels," In Proceedings of the 1999 Annual Symposium of the American Medical Informatics Association (AMIA-1999).

Elvik, R. and Vaa, T. (2004) "The Handbook of Road Safety Measures," Elsevier.

Frisch. R. (1938). “Statistical versus theoretical relations in economic macro dynamics. League of Nations Memorandum,” (Reproduced in Autonomy of Economic Relations, Universitetets Socialokonomiske Institutt, Olso, 1948).

Galles, D., and Pearl, J. (1995) “Testing identifiability of causal effects,” Uncertainty in Artificial Intelligence 11, pp. 185-195. San Francisco: Morgan Kaufmann Publishers.

Greenland, S., and Brumback, B.A. (2002) “An overview of relations among causal modeling methods.” Int J Epidemiol, Vol. 31, pp. 1030–1037.

Geiger, D., Verma, T., and Pearl, J. (1990) “Identifying Independence in Bayesian Networks,” Networks, Vol. 20, pp 507 - 534.

Haavelmo, T. (1943) “The statistical implications of a system of simultaneous equations,” Econometrica 11, pg 1-12.

Hartemink, A. J. (2005). Banjo: Bayesian network inference with java objects, Found at [http://www.cs.duke.edu/~amink/software/banjo/].

Hauer, E. “Cause and effect in observational cross-section studies in road safety,” Found at [http://ca.geocities.com/hauer@rogers.com/Pubs/full.doc].

Hawkins, G., Schertz, G., Carlson, J. and Beck, R. (2000) "Minimum Levels of In-Service Retroreflectivity for Pavement Markings: Summary of Workshop Findings," FHWA Operations Core Business Unit Washington, D.C. Found at: [http://safety.fhwa.dot.gov/roadway_dept/retro/sum2000wkshps/].

Heckerman D., Geiger D., and Chickering D. (1994) "Learning Bayesian networks: The Combination of Knowledge and Statistical Data," Machine Learning, Vol. 20, pp. 197-243, 1995. Also appears as Technical Report MSR-TR-94-09, Microsoft Research, March, 1994 (revised December, 1994).

Heckerman, D. (1999) "A Tutorial on Learning with Bayesian Networks," Learning in Graphical Models, M. Jordan, ed.. MIT Press, Cambridge, MA.

Heckerman, D., Geiger, D., and Chickering, D. M. (1995) "Learning bayesian networks: The combination of knowledge and statistical data," Machine Learning, Vol. 20, pp. 24–3.

Heckman, J. J., Ichimura, H., Smith, J., and Todd, P. (1998) "Selection Bias Using Experimental Data," Econometrica, Vol. 66 (5), pp. 1017–1098.

Heckman, J. J., and Vytlacil, E. (2003) "Econometric evaluations of social programs," Leamer E, Heckman JJ (eds). Handbook of Econometrics, Vol. 6, New York: Elsevier, 2003.

Hoeting, J., Madigan, D., Raftery, A., and Volinsky, C. (1999) "Bayesian model averaging: A tutorial", Statistical Science, Vol 4, Issue 4.

Holland, P. W. (1986) “Statistics and Causal Inference,” Journal of the American Statistical Association, Vol. 81 (396): 945–960.

Hooper, P. (2007) "Exact distribution theory for belief net responses," Technical report, University of Alberta. Found at: http://www.stat.ualberta.ca/~hooper/research/papers+talks/exactbeta.pdf.

Jensen, F. V. (2001) "Bayesian Networks and Decision Graphs," Springer.

Johnston, I. R. (1983) "The Effects of Roadway Delineation on Curve Negotiation by Both Sober and Drinking Drivers," Australian Road Research Board Research Report No. 128.

Karwa, V. and Donnell, E.T. (2009) “Modeling Pavement Marking Retroreflectvity using Artificial Neural Networks,” In Review. Submitted to the ASCE Journal of Transportation Engineering.

Kaufman, J. S., and Kaufman, S. (2001) “Assessment of structured socioeconomic effects on health,” Epidemiology, Vol. 12, pp. 157–67.

Klebelsberg, D. (1982) "Verkehrspsychologie," Springer-Verlag, Berlin-Heidelberg-New York, pp. 367.

Kleiter, G. D. (1996). “Propagating imprecise probabilities in Bayesian networks,” Artificial Intelligence, Vol 88, pp. 143-161.

Kopf. J. (2004) “Reflectivity of pavement markings: analysis of retroreflectivity degradation curves,” Report No. WA-RD 592.1, University of Washington, Seattle, WA.

Lee, J. T., Maleck, T. L., and Taylor W. C. (1998) “Analysis of the Correlation Between Pavement Marking Visibility and Nighttime Accidents,” Proceedings of the 77th Annual Meeting of the Transportation Research Board, Washington, DC, January (CD-ROM).

Lindly, J. K. and Wijesundera, R.K. (2003) “Evaluation of profiled pavement markings,” Report No. UTCA 01465, University Transportation Center for Alabama, Tuscaloosa, AL.

Lobanov, E.M. (1980) "Highway Design Taking into Consideration Drivers Psycho-Physiology," Transport, Russia, pp.311.

Loetterle, F. E., Beck, R. A., and Carlson, J. (2004) "Public Perception of Pavement-Marking Brightness," Transportation Research Record 1715 - Paper No. 00-1473, pp. 51-59.

Mackworth, N. H., and Morandi, A. J. (1967) "The gaze selects informative details within pictures," Percept. Psychophysics, Vol. 2, pp. 547-552.

Madigan, D. and Raftery, A. E. (1991) "Model selection and accounting for model uncertainty in graphical models using Occam's window," Journal of the American Statistical Association, Volume 89, pp. 1535–1546.

Maldonado, G., and Greenland, S. (2002) "Estimating causal effects," International Journal of Epidemiology, Vol. 31, pp. 422–438.

Manual on Uniform Traffic Control Devices for Streets and Highways. Federal Highway Administration United States Department of Transportation, Washington, D.C., 2003.

Masliah, M., Bahar, G., Erwin, T. and Tan, E. (2005) "Towards Improved Management of Performance of Pavement Marking and Markers," Annual Conference of the Transporation Assciation of Canada, Alberta . Found at: [www.tac-atc.ca/english/pdf/conf2005/s16/masliah.pdf].

Meek C. (1995) "Causal inference and causal explanation with background knowledge," Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence, Philippe Besnard and Steve Hanks (Eds.), Morgan Kaufmann Publishers, Inc., San Mateo, CA, 403-410.

Migletz J., Graham L, J., Bauer M.K., and Harwood W.D. (2004) "Field Surveys of Pavement-Marking Retroreflectivity," Transportation Research Board, and National Research Council, Paper No. 99-0825, pp. 71-78.

Migletz, J. and Graham, J. (2002) "Long-Term Pavement Marking Practices," National Cooperative Highway Research Program Synthesis 306, Washington, D.C., Transportation Research Board.

Migletz, J., Fish, J. K., and Graham, J. L. (1994) "Roadway Delineation Practices Handbook," FHWA-SA-93-001, Washington, DC, Federal Highway Administration, US Department of Transportation.

Migletz, J., Graham, J. L., Harwood, D. W., and Bauer, K. M. (2001) “Service life of durable pavement markings,” Trans. Res. Rec. 1749, TRB, Washington, DC, 13-21.

Migletz, J., Graham, J. L., Harwood, D. W., Bauer, K. M., and Sterner. P. L. (2000) “Evaluation of All-Weather Pavement Markings,” Contract No. FHWA C1038/3524-04. Federal Highway Administration, Washington, DC.

Molino, J. A., Katz, B. J., Duke, D. A., Opiela, K. S., Anderson, C. K., and Moyer, M. J. (2004) "Field Validation for the Relative Effectiveness of Combinations of Pavement Markings and Retroreflectivity Raised Pavement Markers in Recognizing Curves at Night," Washington, D.C., 83rd Transportation Research Board Annual Meeting.

Murphy, K. (2005) Software Packages for Graphical Models / Bayesian Networks. Found at http://www.cs.ubc.ca/~murphyk/Bayes/bnsoft.html.

Neyman, J. (1923) [1990], "On the Application of Probability Theory to Agricultural Experiments. Essay on Principles. Section 9," Statistical Science, Vol. 5 (4), pp. 465–472. Trans. Dorota M.Dabrowska and Terence P. Speed.

Neville A. S., and Paul M. S (2009) "Human error taxonomies applied to driving: A generic driver error taxonomy and its implications for intelligent transport s systems," Safety Science, Vol. 47, Issue 2, pp. 227-237.

Parker, A. N. and Meja, J. S. M. (2003) "Evaluation of the Performance of Permanent Pavement Markings," Washington, D.C., 82nd Transportation Research Board Annual Meeting.

Pearl, J. (1994) "Causal Diagrams for Empirical Research," Department of Statistics, UCLA, Department of Statistics Papers, Paper No. 1994010101.

Pearl, J. (2000) "Causality: models, reasoning and inference," Cambridge: Cambridge University Press.

Pearl, J. (1988) "Probabilistic Reasoning in Intelligent Systems," Morgan and Kaufman, San Mateo.

Pearl, J., Verma, T. S., Allen, J. F., Fikes, R., and Sandewall, E. (1991) "A theory of inferred causation," Principles of Knowledge Representation and Reasoning.

Pham-Gia, T. (2000) "Distributions of the ratios of independent beta variables and applications," Communications in Statistics - Theory and Methods, 1532-415X, Volume 29, Issue 12, Pages 2693 – 2715.

Reichenbach, H. (1956) "The direction of time," Univ. of California Press, Berkeley.

Rosenbaum, P. R. (2002) "Observational Studies," New York: Springer Verlag.

Rubin, D. B. (1974) "Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies," Journal of Educational Psychology,Vol. 66, pp. 688–701.

Rubin, D. B. (2006) "Matched Sampling for Causal Effects," New York: Cambridge University Press.

Rubin, D. B. (2008) "For Objective Causal Inference, Design Trumps Analysis," Annals of Applied Statistics, Vol. 2 (3), pp. 808–840.

SamIam (2005) “A comprehensive tool for modeling and reasoning with Bayesian networks,” Automated Reasoning Group, UCLA. Found at http://reasoning.cs.ucla.edu/samiam/.

Sarasua, W. A., Clarke, D.B., and Davis, W. J. (2003) “Evaluation of Interstate pavement marking retroreflectivity,” Report No. FHWA-SC-03-01, Clemson University, Clemson, SC.

Sathyanarayanan, S., Shankar, V. N., and Donnell, E. T. (2007) “Weibull analysis of pavement marking inspection data,” Proc., 87th Annual Mtg. of Trans. Res. Board, Paper No. 08-2892, Washington, DC.

Shinar, D. (2007) “Traffic Safety and Human Behaviour,” Elsevier.

Simon, H.A. (1953) “Causal ordering and identifiability,” in: W.C. Hood, T.C. Koopmans (Eds.), Studies in Econometric Method. Cowles Commission for Research in Economics, Monograph No. 14, John Wiley & Sons, Inc., New York, 1953, Chapter III, pp. 49-74.

Sobel, M. (2000) “Causal inference in the social sciences,” J Am Statist Assoc, Vol. 95, pp. 647–51.

Spirtes P., Glymour C. and Scheines, R. (1991) "An Algorithm for Fast Recovery of Sparse Causal Graphs," Social Science Computer Review, Vol. 9, pp. 62-72.

Spirtes, P., Glymour, C. and Scheines, R. (1993) "Causation, Prediction, and Search," New York, Springer-Verlag, 2nd Edition, MIT Press (2001).

Spirtes, P., Meek, C., Richardson, T. (1999) "An algorithm for causal inference in the presence of latent variables and selection bias," Computation, Causation, and Discovery, pp. 211–252 AAAI Press, Menlo Park.

Sprites, P. (2001) "An anytime algorithm for causal inference," In the Presence of Latent Variables and Selection Bias in Computation, Causation and Discovery.

Standard Test Method for Measurement of Retroreflective Pavement Marking Materials with CEN-Prescribed Geometry Using a Portable Retroreflectometer, E 1710-05, American Society for Testing and Materials, West Conshohocken, Pa., 2005.

Tenenbaum JB, Niyogi S. (2003) "Learning causal laws," In Proceedings of the Twenty-Fifth Annual Conference of the Cognitive Science Society, pp. 1152-1157

Thamizharasan, A., Sarasua, W. A., Clarke, D. B., and Davis, W. J. (2003) "A Methodology for Estimating the Lifecycle of Interstate Highway Pavement Marking Retroreflectivity," No. TRB Paper Number: 03-3867, Washington, D.C., 83rd Transportation Research Board Annual Meeting.

Van Driel, C. J. G., Davidse, R. J., and Van Maarseveen, M. F. A. M. (2004) "The effects of an edgeline on speed and lateral position: a meta-analysis." Accident Analysis and Prevention, Vol. 36, pp. 671-682.

Velichkovsky, B. M., Doronhoefer, S. M., Kopf, M., Helmert, J., Joos, M. (2002) "Change detection and occlusion models in road-traffic detection scenarios," Transportation Research Part F, Vol. 5, pp. 99-109.

Verma, T. and Pearl, J. (1990) "Equivalence and synthesis of causal models" , Proceedings of Sixth Conference on Uncertainty in Artificial Intelligence.

Victor, T. (2000) "A technical platform for driver inattention research," Volvo Technical Report for project NUTEK Dnr 1P21-99-4131, Volvo, Goteborg, Sweden.

Willis, P. A., Scott, P. P., and Barnes, J. W. (1984) "Road edgelining and accidents: an experiment in South-West England," TRRL 1117, Transport and Road Research laboratory.

Zhang, Y. and Wu, D. (2006) "Development of methodologies to predict service lives of pavement marking materials," J. of the Trans. Res. Forum, Vol. 45, Issue 3, pp. 5-18.

Zwahlen, H. T. and Schnell T. (2004) "Minimum In-Service Retroreflectivity of Pavement Markings," Paper No. 00-1479 - Transportation Research Record, pp. 60-70.

Zwahlen, H. T. and Schnell T. (1995) "Minimum In-Service Retroreflectivity of Pavement Markings," Transportation Research Record 1715, Vol. Paper No. 00-1479 - Transportation Research Record, pp. 60-70.

Zwahlen, H. T. and Schnell, T. (1997) "Driver Eye-Scanning Behavior as Function of Pavement Marking Configuration," Transportation Research Record 1605, Transportation Research Board, National Research Council, pp. 62-72.