april 26 - 28, 1988
amsterdam
the netherlands

traffic safety theory & research methods

Contributions of the invited speakers

SWOV

This book was made possible through the support of 3M Europe

CONTRIBUTIONS OF THE INVITED SPEAKERS

Session 1

Frank A. HAIGHT, University of California, Irvine, USA

Research and theory in traffic safety

Session 2

Ezra HAUER, University of Toronto, Toronto, Ontario, Canada

Research on the effect of road safety measures; A personal view (Paper outline)

Session 3

John A. MICHON, University of Groningen, The Netherlands

Driver models: How they move (Preliminairy version)

Session 4

Mike MAHER, Transport and Road Research Laboratory, Crowthorne, United Kingdom

Statistical models for accident data

Session 5

A.C. HARVEY, University of London, United Kingdom

Intervention analysis and structural time series modelling

RESEARCH AND THEORY IN TRAFFIC SAFETY

Frank A. Haight
Institute of Transportation Studies
University of California, Irvine

Abstract. This paper discusses traffic safety research in various aspects: the need for research, methodology of research, difficulties in research, institutional siting of research and concludes with some areas which could be included in a research agenda.

1. Introduction

The theme of this conference -- theory and research -- makes it, I believe, unique. In contrast with other traffic safety meetings, we are dealing here not with intervention, but with understanding. Many of us have heard demands that we "do something" about traffic accidents, but it is only recently that there have been suggestions that we should "know what we are doing", before we begin to do it. These suggestions have come, not from the general public, or even from those responsible for countermeasure programs, but from the research community (cf. Hauer 1988, Evans 1988). It is a curious aspect of traffic safety that so much action has been based on so little knowledge.

Therefore, I'd like to start off by paying tribute to SWOV and its directors for sponsoring such a radical departure from the usual traffic safety conference.

In making this sharp verbal distinction between research and intervention, I don't intend to gloss over the links which obviously exist between the two. It is a fact that more knowledge of any phenomenon generally leads to better methods for controlling it, and that ignorance is no acceptable basis for action. In the field of traffic safety this wisdom is unfortunately often perverted by the substitution of "passion" for "knowledge" and "objectivity" for "ignorance".

Of course it is also true that a good research agenda should be framed by knowledge of the consequences of past interventions.

But it is important to remember that, although research and intervention are closely connected, the practice of research and the practice of intervention are -- and should be -- separate professions. We don't expect a molecular biologist to practice medicine, or even to make recommendations for public health programs. I believe that there is need for both impartial investigators and partisan advocates, and also, obviously, for suitable means of communication between them. It is through scholarly publications and scientific meetings, that advances in learning and in doing have the opportunity to interact.

My paper is divided into five parts: the need for traffic safety research, the methods of traffic safety research, the pitfalls of traffic safety research, the institutional siting of safety research and, in conclusion, some suggested research areas.

2. Need

There is a surprising resistance to learning more about safety. Frequently we hear questions like these: Isn't it enough to apply existing knowledge, so that the world-wide drop in fatality rate per distance traveled will continue to decline? Isn't the overwhelming evidence of driver error in accidents sufficient to define an agenda of intervention? How is it possible to justify spending money on research that could be better spent on saving lives?

I'd like first of all to provide some answers to these -- and similar -- rhetorical questions. In answer to the first one, I would say that although certain rates are dropping in every country for which sufficient information exists for accurate calculation, the world-wide number of road deaths is now approaching half a million per year -- with corresponding numbers of injuries. Personal injury has been called by the U. S. National Academy of Sciences the leading public health problem in the United States today. (Houk, et al 1987) To me, this suggests a topic worthy of continued objective investigation.

The answer to the second question -- can programs of intervention be inferred from knowledge of driver error? -- is a simple "no". "Fault" gives little if any guidance in designing countermeasure programs. It is a basic legal concept, used by policemen in regulating road traffic, and by lawyers in their search for their own and their clients' compensation, but it has proved to be a will-o-the-wisp for the program designer. The traditional appeal by fault-finders begins "If only we could convince drivers that . . . ." But cost-effective means of persuasion seem to be as elusive today as at the dawn of motorization.

The third question -- querying cost of research vis-a-vis cost of intervention -- is based on an assumption: that by spending money, a fairly sure and proportional benefit could be obtained. It would certainly be convenient if this were true, but the evidence is hardly convincing. For many reasons, some of which are sketched later in my paper, the relationship between safety investment and safety return is obscure at best, and often quite unknown.

There are numerous examples of expensive programs which are not known to have produced significant results. This is especially true in the area of drunk driving countermeasures. Better criteria for program selection would be one payoff from research.

The argument for research can also be expressed in positive terms. I hope I will not offend anyone if I first mention a basic motive in much science: curiosity. We find plenty of unanswered questions about the accident phenomenon, and many of us would like to know the answers. As Hauer (1988) has demonstrated in great detail, there are difficulties in making precise such presumably simple relationships as those between safety performance and safety engineering. If we consider interventions other than those which are based on engineering, the confusion is greater. And if we go even further to investigate the relationships which seem to exist between safety and parameters which may affect safety without being particularly designed to do so, we are in a region where little is really known, but much is suspected. In short, there is no lack of interesting and presumably valuable topics for study.

The value of research relates specifically to competing threats to public health. If research funding is properly proportional to the cost to society of specific conditions, then accidents are, it is generally agreed grossly underfunded. This is especially true in comparison with certain relatively rare diseases which attract public sympathy. Even if the relationship between social cost and research investment is not linear, but needs to be modified to take into account the probability of payoff from the research, it still seems clear that the claim of underfunding must be accepted.

3. Methodology

Now, I would like to turn to the question of research methodology, and specifically to discuss the extent to which the techniques of safety research are shared by other disciplines, and the extent to which they are unique.

It is clear that many studies which contribute to safe road transport fall within traditional disciplines: vehicle design and performance within mechanical engineering, driver behavior within psychology, alcohol effects within human physiology and sociology, and the road environment mainly within civil engineering.

However, even in these established fields, there are, for traffic safety research, peculiar difficulties in formulation of objectives, in the design and analysis of experiments, in arriving at scientifically based conclusions and in presenting policy alternatives. Let me mention briefly three categories of problems confronting the safety analyst.

First, safety is not an isolated goal, which can be easily compared to the eradication of disease. Accidents are but one side-effect of an industrial, social and economic system which is substantially based on road transport, and specifically on individual decisions regarding road transport alternatives. To study safety without regard to the transportation context in which it is embedded can and often does yield results which, although possibly true, are inconsequential. I'll give some examples in the discussion of research pitfalls.

A second important characteristic of safety research relates to experimental design. Many aspects of the road transport system are essentially out of the reach of the research worker. It is virtually impossible to carry out a designed experiment, using procedures that have been so successful in agricultural trials, industrial research, or medical investigations. The fundamental concepts of placebo, randomized treatment, control group, etc. are difficult to apply in the transportation context. The physical facilities, vehicle fleet and operating systems are more or less fixed, or at least have long lifetimes, and road users are nearly immune from experimental studies. Also the managers of the system are generally unsympathetic to requests for tampering with traffic in the interests of science.

The third difficulty arises from the fact that the independent variables affecting safety are so numerous, so complex and so interrelated, as to present nearly insuperable problems in multivariate analysis. In addition to all the factors inherent in the transport system there are economic factors, social and climatic factors, political, legal and institutional variables ... the list is nearly endless. Any one of these seems at some time, in some location, to have influenced the dependent variables which characterize road safety.

For example, if aggregate death rates are chosen as measures of level of safety it is nearly impossible to attribute changes to corresponding changes in specific independent variables. Although this may be partly a problem of experimental design, it is nevertheless also a constraint on research methodology. It means that dependent variables always need to be substantially disaggregated, to have a better chance of correlating the results with exogenous variables.

But disaggregation also entails difficulties, especially those associated with small sample size. Often a newly implemented program is judged on an apparent discontinuity in a carefully chosen safety related time series. But reliance on visual evidence of a curve with a kink in it as conclusive isn't any longer acceptable. It may bring joy to the heart of the program administrator, with a little luck it may even be statistically significant, but it often turns out that the new trend won't stay in place, or that the next-door jurisdiction also has a kink in its curve without the intervention, or, most frequently, that other constituencies press forward to "claim" the kink as the result of their own efforts. It is interesting that the benefits forecast in the 1960's as a result of the Federal Motor Vehicle Safety Standards arrived in the early 1970's just about as predicted, but were then attributed to the energy panic.

All three impediments to scientifically based conclusions are illustrated in the results of a committee commissioned by the U. S. National Research Council (1984) to study the effects of the national 55 mph speed limit. Although the composition of the panel and much of the methodology suggested political aims rather than scientific ones, there was an attempt to quantify both the safety and mobility effects of the speed limit. A controlled experiment was evidently not considered, and conclusions were based on other publications and existing data. Mobility loss was calculated roughly, but not balanced against gains in safety. The problem of confounding variables was explictly recognized, but little was offered as a solution. The following excerpt from the committee's report suggests the problems confronted: "In determining the effects on safety, analysis is confounded by the difficulty of isolating the effects of the speed limit from other causes of the improved safety record. Indirect estimating techniques must be relied on, and assumptions must be made in the process. The committee believes that in spite of the difficulties some rough estimates can be offered, based on the plausibility of the techniques employed and the similarity of findings that emanate from different data and different statistical methods. Nevertheless, an exact determination of a specific number of lives saved by the 55 mph speed limit is not possible. Data on the effects on serious injuries are particularly sketchy, and any estimate of the effect of the speed limit on injuries is essentially an educated guess."

I don't want to suggest that there are no other fields of research similarly handicapped by experimental constraints and multiple causes. Although it separates accident analysis from those traditional fields closely connected with the development of statistical inference, that does not make it unique. It is my impression that many of the constraints on accident research which I have outlined, apply also in the field of economics. But I do believe that economists have more reliable data than we do, in fairly long series, so that their problems relate more to finding appropriate ways to manipulate the data, rather than to attempting inference from haphazard, biased and fragmentary information.

Given then, that relationships which may exist between safety parameters and variables which influence them have not so far been amenable to traditional statistical methodology, the search for " needles in the haystack" seems bound to depend to a considerable extent on quasi-experiments, which will be discussed by other speakers, and by the methods of epidemiology.

Epidemiology, originating in the study of infectious diseases, developed many valuable concepts (Glass 1986) which should be carefully considered for their application to accident analysis: taxonomy of victims, life table analysis, standardized mortality rates, person-years at risk, dose-response relationships and years of potential life lost.

Some of these concepts, in order to be applied in accident epidemiology, may need some modification. For example, in view of the multiplicity of independent variables, a simple dose-response formulation would perhaps be replaced by several such relationships, one for each significant factor. The effect of randomness in accident experience and especially in severity is probably greater than in disease. The "agent of harm" is in our case, also an "agent of good", namely mobility. The adaptation of epidemiology to safety is a conceptual as well as a statistical problem.

The nearest analogue to the concept of dosage is that of "exposure", which has been used principally in the context of a single individual, usually a driver. But it seems clear that exposure to traffic produces a much feebler response than exposure for example to typhoid; with a smaller mean, and much larger variance. Thus, dose-response techniques -- for example, probit analysis -- would need to be applied to populations of considerable size for meaningful conclusions to be likely. These populations could be defined in many different ways, being based on classes of roads, of road users, of vehicles or operating/enforcement systems. It seems probable that the most fruitful analysis would involve populations which cut across these categories.

One goal of epidemiology is to identify meaningful clusters of significant events. We are already familiar with the identification of black spots, or clusters in space. It would be more helpful if the pins on the map which define a black spot also carried information about various attributes of the victims of accidents at these spots. Beyond this, if enough data were available, it might be possible to find clusters in higher dimensions of time, space, personal characteristics, vehicle type and history and so forth.

Another aspect of classical epidemiology, is the identification of so-called "subclinical" manifestations of disease before the disease itself is detectable. In our context, such conditions might be, in addition to traditional traffic conflicts, climatic/geometric environment combinations; for vehicles, a history of minor collisions; for individuals, a variety of social, economic and psychological indicators. Even "accident proneness", now quite discredited because of exaggerated claims, poor experimentation and statistical naivete, might prove to have some small merit as a subclinical indicator.

To summarize: epidemiology does in my opinion, hold out the prospect of many new directions in accident analysis. I would also like to add at this point that the research methods of accident epidemiology should apply equally to all accident types: not only transport accidents, but also those occurring in industry, the home and in recreation. The interventions which follow are of course different and usually specific to the type of accident, but the research methodology is remarkably similar.

4. Pitfalls

After talking about research methods, it seems appropriate to discuss research difficulties: the pitfalls and problems that may be encountered in our attempt to build up a scientifically valid body of knowledge about accidents.

I think the first problem is simply to stay research-oriented, that is, to work with the goal of discovering and publishing objective information. It may be surprising to newcomers to the field to learn how difficult that can be. There is always the temptation -- often supported by considerable offers of funding -- to do "something useful." I've put these last two words in quotation marks, because the usefulness is often only in the eye of some person or organization which is committed to a particular agenda, and the work being proposed is wanted only as evidence to support that agenda.

Sponsors of accident research often demand that, before being funded, the research worker first demonstrate how any knowledge which might be forthcoming project can be transformed into a "life-saving" intervention during the current budget year. Some projects, for example a search for subclinical indicators would be especially vulnerable to this requirement. Most research proposals require a good deal of ingenuity and often some downright prevarication to satisfy assure sponsors of immediate payoff.

Maintaining objectivity seems to be more difficult in safety research than in other branches of science. It would be too daring of me to say that objectivity is a dirty word in the safety profession, but I will tell you that it does not exactly have the luster we find in other disciplines. The traffic safety field is dense with advocates, often with only a frail basis for their advocacy. Believing sincerely in a particular program, the advocate may be willing to insist that 2 + 2 = 7 if he or she sincerely believes that by doing so, lives will be saved.

Whether or not disinformation will in fact save lives -- or, more accurately, postpone deaths -- it seems to me an obligation of the research worker to insist that 2 + 2 is 4, whatever the consequences. For example, although it is clearly desirable from the point of view of society that seat belts be worn at all times, I see no point in denying the fact that fastening a seat belt on a particular journey -- even the most hazardous -- has extremely low probability of producing any benefit to the wearer. It is precisely for this reason that we can justify belt-wearing laws; if the case for voluntary wearing made sense, the laws would be unnecesary. Similarly, the often repeated demand that the drinking driver be persuaded that he or she will surely be caught and severely punished lacks a basis of truth. "Research" to determine how best to convince people of untruths is somewhat outside the customary agenda of science.

In addition to avoiding false research, it is also desirable to avoid trivial research. We don't need any more experiments to show that those under the influence of alcohol are unable to steer around traffic cones (or unwilling to do so -- I wonder what would happen if they were offered a large cash reward). It is in my view equally unnecessary to demonstrate statistically that if blue-eyed people were to be deprived of driving licenses, they would experience significantly fewer traffic violations, car crashes and injuries; that indeed it would be a life-saving countermeasure for blue-eyed drivers. Conclusions of this sort, although supported by tests of significance at the five percent level, are not really worthwhile. They are usually conducted for purposes of persuasion, and thus also fall outside the boundaries of science.

One particular type of trivial research has already been alluded to: that which ignores the the role of the road transport system in maintaining an industrial society. This might be called "suboptimizing on safety". If $m < n$, then other things being equal, a speed limit of m, provides more safety and n more mobility. The real task of the research worker is not to quantify the obvious, but to devise methods for finding an optimal balance between the two socially desirable goals of safety and mobility, with due regard to a third element, cost. The fundamental triad is safety, mobility and cost. By spending more money, it is possible to increase both safety and mobility. By decreasing mobility it is possible to increase safety and decrease cost. Several other permutations come to mind. Little has been done to address optimization of the triad, although a rough framework appears in a paper by Kamerud (1988).

Still another group of difficulties arises from a tradition of using false taxonomies in traffic safety. These include the categories used in police reporting and legal proceedings, which involve "fault" as a basic descriptor as well as classifications oriented around specific intervention strategies. In the latter category, we would find, for example, "alcohol" mentioned if involved in an accident, rather than, for example "poverty" simply because alcohol use is supposed to be more amenable to safety measures than poverty, or that drinking and driving is more the "fault" of the driver. Thus "alcohol" has become an accepted category either for moral reasons, or for reasons of intervention, but not because it has yet been shown to belong to a meaningful or useful taxonomy.

An important consequence of adopting the epidemiological point of view is that it does not admit either morality or intervention as a sufficient basis for taxonomy. I once made the suggestion that a parallel to a countermeasure-oriented taxonomy would be to classify research into anatomy of the cockroach into categories of poisoning, crushing, burning, etc. Actually, the categories should emerge from the data, and should be based on victims and their characteristics.

Bad taxonomy undoubtedly arises from lack of adequate data, but at the same time it also contributes to some important gaps in information. With fault an accepted as the basis for classification, it is overwhelmingly assigned to the road user, whether it be "driver error" in the case of the driver, "dart-out" in the case of the child pedestrian or "failure to have due regard to the circumstances" when all else fails. This system leads to official secrecy about the statistical characteristics of victims, since the victims have mostly already been assumed to be "guilty". The logic seems to be that if research indicated that a particular category of individuals were more likely to be the victims of traffic accidents, someone -- perhaps the press -- would decide that people in this category were "bad drivers."

As I have emphasized in another context, the blame attached to young male drivers comes not only from a valid statistical basis, but also from the scarcity of categories of road user to choose from. With only age and sex given, the research worker is in the difficult situation of spending his time on some new aspect of the young-male-syndrome, or, if he is clever, trying to work out some file-linkage procedure to discover, for example, family income, place of residence, ethnicity, travel behavior or criminal history.

It appears that in medical research, not only are well-designed experiments acceptable -- as in the case of recent aspirin trials reported in the press -- but also it is reasonably convenient for research workers to have access to certain confidential records. In disease research, the enemy -- if there is one -- is nature, and so it is easier to regard the victims' personal characteristics objectively. With the "fault" concept in accident studies, it becomes difficult to concentrate on victims, when there are easily available scapegoats, faulty drivers. It would help our understanding of the accident syndrome if a few slots in large surveys were reserved for research questions.

A final problem relates to the competence and qualifications of research workers. Most of us have to come to the accident field through some other disciplines, and have learned slowly and sometimes painfully the principles which I have discussed in the earlier part of this talk. New workers in the field, lacking any curriculum of professional training, must tread the same path and this usually means writing naive papers, rejected by editors, until they have found their way amongst the pitfalls in accident analysis.

A current example of the lead-in time needed for research sophistication can be found in the work of the newly established National Center for Injury Control of the U. S. Public Health Service Centers for Disease Control. The papers submitted to <u>Accident Analysis and Prevention</u> by members of that group have received mostly bad referees' reports reflecting mainly the authors' naive approach to the subject. Among other things, many authors seem unaware of the existing literature and are painfully trying to start from scratch. Both Evans and Hauer have commented on the tendency in accident research for each new recruit to begin at the beginning rather than to build on earlier results. It is a pity that there does not exist any systematic way to become acquainted with the basic literature of the accident field.

5. Siting

Next, I would like to discuss briefly the question of institutional siting for accident research. There is a clear need for research, like evaluation, to be conducted away from the pressures of project formulation and implementation. For this reason, it is a temptation for me to proclaim that only in universities is it possible to maintain objectivity, resist pressure groups, and carry out a coherant, long term research agenda.

There are, however, a few partial counter-examples. Twenty years ago, the SWOV, in its annual report, expressed the need for and dedication to "fundamental knowledge," and has proved successful in some specific fields. There was a time when the Road Research Laboratory in the United Kingdom also contributed to fundamental knowledge. In Sweden, the VTI has a history of basic research sponsorship, as have BASt, ONSER and ARRB. General Motors Research Laboratories has conducted some useful studies and of course many national transportation agencies have safety research components.

I should also acknowledge that very few universities have thus far provided support for a research institute in accident studies. The comparison with proliferating institutes of transportation is especially noteworthy. The few groups which do exist in universities are mostly living on the fringes of the academic mainstream, funded by soft money, usually in the form of short term contracts.

There are two basic reasons why universities have not been more congenial to accident research. First, the departmental structure does not accommodate interdisciplinary subjects easily, and safety research spans quite a goodly number of traditional disciplines. More importantly, the fundamental teaching mission of academic institutions has not been addressed by the safety community in the form of standard textbooks. Just as a core curriculum is needed to educate research workers, so the same kind of curriculum must be developed if accident research is ever to be independent of casual funding. From an academic point of view, the best arrangement would appear to lie in grafting the teaching mission onto public health schools rather than onto transportation institutes, at least until transportation earns its way to departmental status.

In spite of all these constraints, I do believe that universities are able to provide the unique, most needed ingredient for safety research, namely independence and objectivity.

6. Areas for Research

The conclusion of my paper consists of a short, and admittedly subjective, list of some areas which might form the basis for a research program in an academically based institute. I'll omit epidemiology, which I have talked about enough already, and engineering, which is covered in the paper by Hauer (1988).

One interesting category of problems concerns theoretical models relating fatality rates to one or two time-dependent variables. The first of these was

Smeed's (1949) formula which purported to give death rates in terms of human and vehicle populations. There is a large literature which derives from Smeed's simple idea, much of it claiming far more for the formula than Smeed ever intended. The recent paper of Andreassen (1985) shows how the formula arose, how it has been misused and in doing so provides an important critique of the basic idea. Second, the time dependent formula for deaths per unit of mobility, a one-parameter curve that looks negative exponential for almost any jurisdiction, but which has never been systematically fitted, with the result that the parameter value has not been related to motorization level, or indeed to any other independent variable. Third, the curve proposed by Marchetti (1983) for percapita fatalities as a function of time, which seems to have been independently hypothesized by Jørgensen (1985). This model would appear to require three parameters for specification and gives reasonable eyeball agreement with data from Denmark, the United States and Japan. Fourth, a model proposed by Oppe (1987) which is based on negative exponential fatality rates combined with logistic travel growth.

I believe there is also an opportunity for further basic research on project evaluation. Specifically, we should have more accurate information on value of time versus value of life. These important ingredients for planning and evaluation need to be made more precise, not just by averaging numbers adopted by various agencies, but by seriously analyzing the conceptual questions -- value to who? -- and by some realistic measurements. Evaluation research also requires better measures of effectiveness than have so far been used. An example from the FHWA evaluation handbook on the installation of stop signs assesses costs only to the agency which installs the signs, omitting the cost of bringing a car to a halt and then starting up again. It is not surprising that stop signs, suboptimized on safety alone, omitting effects on mobility and using only a fraction of total cost, turned out to be extremely cost-effective. Evaluation research should address the length of time needed for new operating systems to reach equilibrium, better costing procedures, and especially help to develop better measures of effectiveness.

Another promising research area relates to public policy towards risk: how much individual risk-taking behavior deserves to be considered an area of social responsibility and how much rests with the risk-taker. For example, are the social-cost arguments for compulsory seat belt wearing equally valid for hazardous recreational activities?

I would also endorse the suggestions made by others for more complete information on the driving task, and especially as it relates to category of road user. In the field of transportation research, the topic of driver information systems is receiving a good deal of attention. This area has important relationships with accident experience. It is linked with another important question: the relationship between traffic safety and demographic changes in the population, specifically the increasing size of the cohort of the aged.

There is also more to be done in the area of exposure measurement techniques, particularly with respect to pedestrian exposure. There are some curious discrepancies between the industrialized nations in the percentage of casualties who are not vehicle occupants.

Another area of interest is the relationship between traffic safety and economic indicators. There have been a number of good papers on the subject, (Wagenaar 1984, Partyka 1984, Joksch 1984) but no really systematic pursuit of the causes and consequences of the linkages.

I have already mentioned curriculum development, among other things for bringing novices in the field up to a reasonable level of sophistication. At the moment, this need is particularly felt by those without prior professional involvement in the field of transportation. This group includes public health specialists for whom the disease paradigm is a natural approach to accident research. The blending of transport principles and public health principles seems to me to be an ideal area for academic research.

Still another good research area would be the objective study of compensation by road users not only to deliberate safety measures but more generally to all kinds of variations in the road and vehicle environment. There has been a good deal of hypothesizing but very little experimentation and still less theoretical formulation.

International comparisons of accident data, of institutional arrangments, of operating systems, of legal/judicial sanctions is another area which deserves some serious attention. There is general agreement that transport safety problems are especially troublesome in developing countries; the way in which safety parameters change with increasing motorization deserves to be investigated.

In earlier papers, I have emphasized, perhaps excessively, the desirability of obtaining socio-economic profiles of accident victims. Here, there may even be room for further studies on the relationships between alcohol consumption and road user behavior. I'm not thinking of quantifying impairment any further, but rather experiments designed to separate the effects of physiological impairment from those induced by attitude change, and to relate the two effects to personal and psychological variables.

Thank you for your attention and patience.

References

Andreassen, D. (1985) Linking Deaths with Vehicles and Population. ARRB Internal Report 000-225. Australian Road Research Board, Vermont South, Victoria, Australia.

Evans, Leonard (1988) Commentary on two papers on mandatory safety belt use laws, and reflections on broader issues. In Traffic Safety as Injury Control (John Graham, ed.) Auburn House, Dover MA.

Glass, Roger I. (1986) New prospects for epidemiologic investigations. Science 234:951-956.

Hauer, Ezra. (1988) A case for science-based safety design and management. Presentation at ASCE Highway Safety at the Crossroads. San Antonio, March 28-30

Houk, Vernon N., Brown, Stuart T. and Rosenberg, Mark L. (1987) Injury prevention and control comes of age. Public Health Reports 102:574-576.

Hutchinson, T. P. (1987) Road Accident Statistics Rumsby Scientific Publishing, P. O. Box 76, Rundle Mall, Adelaide SA 5000, Australia

Jørgensen, N. O. (undated) Traffic safety towards 2000. (unpublished mimeo) The Technical University of Denmark, Lyngby.

Kamerud, Dana B. (1988) Evaluating the new 65 mile per hour speed limit. In Traffic Safety as Injury Control (John Graham, ed.) Auburn House, Dover MA.

Marchetti, Cesare (1983) The automobile in a system context: The past 80 years and the next 20 years. Technological Forecasting and Social Change 23:3-23.

Oppe, Siem (1987) Macroscopic models for traffic volumes and traffic safety. Institute for Road Safety Research SWOV, The Netherlands. (unpublished mimeo)

Smeed, R. J. (1949) Some statistical aspects of road safety research. Journal of The Royal Statistical Society, Series A 112:1-23.

RESEARCH ON THE EFFECT OF ROAD SAFETY MEASURES; A PERSONAL VIEW

E.Hauer

Paper Outline

Haight draws to our attention the distinction between the world of "safety knowledge" and the world of "safety action", the difference between the practice of "research" and the practice of "intervention". He notes that "the practice of research and the practice of intervention are -- and should be -- separate professions". My task here is to comment on the "practice of research", specifically, about research on the effect of road safety measures. Accordingly, the point of view which I take is that of the professional who practices such research.

From this vantage point it is tempting to be introspective, to speak of the methods and theories which help us to do our research work. However, our research is but a means to an end. We, researchers and theorists hope, that eventually, our collective effort will lead to improved "safety knowledge" and thereby to better "safety action". Even when we see that, in spite of our endeavours, improved safety knowledge is slow to emerge, even if we note that what safety action takes place is only marginally influenced by what safety knowledge already exists, it is still natural and convenient for us, researchers, to strive to do ever better research. Accordingly I will devote the second part of this paper to questions of method.

However, so it seems to me, progress towards better safety knowledge is obstructed not only or primarily by inadequacies of method and theory. It is obstructed also, and perhaps mainly, by the very world of "safety action" which we, researchers, intend to support by our work. In fact, so I will argue, it is not so much the limitations of "safety knowledge," which cause the sorry state of "safety action", it is more that the real world of "safety action" tends to obstruct progress towards better "safety knowledge". Because this topic is in my view important yet seldom discussed, I will devote to it the first part of this paper.

1. Research and the Delivery of Road Safety.

By "delivery of road safety" I mean the set of road safety related actions which are the responsibility of government. Thus, the delivery of road safety consists of the licensing of drivers; the setting of vehicle standards; the prescription of the rules of the road; the enforcement of these rules; the design, building and maintenance of roads; the management of traffic on these roads; the provision of emergency medical services and the like. I have chosen the "responsibility of government" to be the defining feature of the "delivery of road safety". While alternative definitions are possible, the actions listed above are in fact actions by government and their employees and collectively do give a satisfactory interpretation to the phrase.

Be the "action" of local importance (such as to install a STOP sign at some intersection) or of broad significance (such as to increase the national speed limit). In North America it is usually only mildly affected by a knowledge-based anticipation of its safety consequences. Perhaps with the exception of vehicle standards, this is how it is now and this is how it always was. As Haight observes, "traditionally action has preceded knowledge".

Typically road safety delivery actions (being actions of the government) tend to be associated with legislation, budgets, programmes, standards, codes, administrations, jobs, careers etc. The popularity, success, perpetuation and growth of such action, once taken, become the self-interest of many. Conversely, any intimation that what action has been taken is not cost-effective, is a threat. Therefore, once action has been taken, it is usually convenient not to ascertain its real safety impact or at least do so "in-house". To do otherwise, is to risk not only embarrassment but also to do real harm to a variety of real interests which that action brought into existence.

The net result is predictable. At the time the action was taken, knowledge of fact did not (and often could not) exist. Once the action has been taken there is no compelling reason why factual knowledge should be acquired while there are strong reasons not to do so. This is what brings about the reign of ignorance.

Furthermore, those who control the "action" of road safety delivery also exert strong influence over what research is funded, who does the work and what is reported. As a result, there is a much larger inclination to herald success than to publicize failure. Therefore, since factual knowledge is neither required nor encouraged to grow, and what passes for knowledge is polluted by self-interest, one should not hold scarce or poor reseach to be responsible for the slow progress towards cost-effective delivery of road safety. Rather, the culprit is the world of road safety delivery which has little use for factual knowledge and is inhospitable to research about the effect of measures which have been implemented.

Another repercussion is the tendency for road safety delivery to be symbolic, rather than safety performance oriented. Thus, e.g., the police are not known to measure the effect of their speed enforcement activity on the speed distribution on the road. They count the number of speeding tickets instead. Ascertaining, what relationship the annual harvest of tickets has to the speed at which people drive, and thereby to road safety is not regarded to be in the domain of police responsibility. The symbol (the action of apprehending a violator of the law) becomes the product instead of the intended result (the reduction of accidents). Similarly, highway engineers design crest curves to give drivers a nominal distance to stop if there is on the road an obstacle of given height. What the relationship between this "sight distance" and the occurrence or severity of accidents does not seem to be known. Thus, "sight distance" -- a symbol, is what governs design, not a fact-based anticipation of how safety changes with sight distance.

I have discussed these issues at length in two recent papers (Hauer,1987 and 1988) providing, what I hope is sufficient anecdotal evidence, to support the claim that this diagnosis fits reality (in North America). Therefore, in the full version of this paper I will only provide selected illustrations of the seriousness and the pervasiveness of the problem.

At the root of the problem is the universality of self-interest. We recognize that the private sector is motivated by self-interest and look for the government to provide oversight and regulation when needed. Because we are so used to think of government as a possibly inefficient but certainly benign protector, it is perhaps not easy to recognize that self-interest, albeit of a different kind, is also behind actions by government. As a result, there are no well developed institutions to protect the public from government self-interest.

In the case of road safety delivery the government is the sole "producer". It appears that it has little self-interest in finding out what the safety effect of various actions is and there is often definite interest in not finding out. For this reason it perpetuates a style of road safety delivery which is not supported but fact-based knowledge.

The remedy to this ailment is not simply to insist that more research be done or that it be done better. One has to aim at the core of the problem. If there is a natural tendency not to ascertain the safety effect of actions, the duty to do so much be enshrined in law. If there is a natural tendency to control the results of such research, one must insist that it be done by experts who have no stake in the outcome. This will usually mean a complete separation between the agnecy which initiates and implements and the people who evaluate.

2. Some questions of method.

The point has been made that the prevailing societal arrangements for the delivery of road safety create an inclement environment for the growth of factual knowledge about the safety effect of interventions. This explains much of the prevailing state of ignorance. Another part of the explanation must be ascribed to the real difficulties of finding out what works and how. These real difficulties are two in kind.

First, we would be able to learn a great deal faster if it was possible to conduct large-scale randomized trials. That the conduct of such experiments is deemed "impossible" is in part a result of a certain lack of determination. After all, if it is possible to conduct randomized trials about the effect of by-pass surgery it is not readily apparent why it is impossible when it comes to the examination the safety of, say, vehicle-actuated signals. Nevertheless, one has to admit that in many cases it is genuinly difficult to think of randomized trials and one has to learn from retrospective studies.

The need to extract defensible information from retrospective studies gives rise to the second kind of real difficulty; variables are many, interactions are complex and one can not stop the world. What methods, strategies and approaches promise to deliver results under these conditions?

What follows is neither in the nature of advice nor is it a summary of current consensus. I only intend to discuss five issues which may contain elements of an answer to the above question. But first, let us define what the task is.

All research about the road safety effect of a measure reduces to the following pair of questions:

1. What is the safety of the entity with the measure in place
2. What would have been the safety of the same entity had the measure not been implemented.

We use a variety of ruses (experimental designs) to guess at the answer to question 2. Sometimes we use only a few years of "before" data to quess "what would have been", perhaps refining our method by using a "control system"; at other times we use a longer sequence of data and place our trust in the extrapolation of some regularity over time; a third popular choice is to use similar entities which remained without the measure to make inferences about what would have been the safety had "our" entity remained without the measure. In any case, the second question is about an event which has not occured and is therefore is not observable. We must be content with inductive validity. It is in this context that questions of method arise. What methods and strategies serve to enhance the inductive validity of our inferences.

The following five issues will be discussed:

a. How to estimate

b. Do not test hypotheses

c. How to let knowledge grow

d. When to use a "control system"

e. What is worth knowing

REFERENCES

Hauer, E. (1987), The Reign of Ignorance in Road Safety: A Case for Separating Evaluation from Implementation. Proc. Transportation Deregulation and Safety, The Transportation Center, Northwestern University, pp. 113-140.

Hauer E. (1988), A Case for Knowledge-Based Safety Design and Management. Proc. "Highway Safety at the Crossroads", ASCE Specialty Conference, San Antonio.

DRIVER MODELS: HOW THEY MOVE

John A. Michon
University of Groningen

INTRODUCTION

Driver models are the toys of traffic safety research. But, are they *good* toys? The principal characteristic of a good toy is its propensity to generate surprisingly complex behavior from a few very simple principles. The vehicles described by the neuro-physiologist Braitenberg (1984), for instance, prove that hardly more is required to interpret behavior than the principle of feedback. And since the principle of feedback pervades the universe, it would seem the ideal vehicle for modeling of driver behavior. But is it? Isn't homeostasis too general, and thereby too weak a principle? And, aren't there perhaps other, equally ubiquitous and equally generative principles that qualify as foundational for behavior models?

Consequently, to start a discussion about driver models, a convenient approach would seem to categorize them according to what makes our toys move. Here we have a whole spectrum of possibilities at our disposal.

On the one hand there are models that are moved by *magic* or what amounts to the same, by hand and by chanting "*vvvrrooommm!*" In traffic research we are occupied with many such models although sooner or later we may hope to recognize them for the curve-fitting tricks they really are. At the other extreme of the spectrum we find models that move autonomously and by doing so learn from their experiences. Such models - if they existed - would be able to cope in a *reasonable* way with the environment.

Reasonable, that is indeed the proper term! Reasonable, or *rational*, is what we tend to call the performance of models that react to external inputs while "keeping their goals in mind". But if we choose to talk about simple feedback models - such as Braitenberg's vehicles, for instance - in such intentional[1] terms, we take out a rather formidable loan on the explanatory power of our models. By simple I mean in this case all such forms of feedback that aim at maintaining a specific output variable at a constant level without structural modifications. Models that represent this kind of feedback do not adapt their internal structure on the basis of their experiences.

It requires yet another category of models, models that are driven by concepts and adaptive rules, and that are consequently able to learn. Only models that fit into this last category can ultimately be said to move autonomously and they will tend to move better all the time. They constitute the only class of driver models that ultimately has scientific survival value.

In this paper I shall consider various prominent and less prominent driver models with this criterion in mind. I wish to emphasize, however, that this is a meta-theoretical, not an empirical criterion. I also wish to point out that I am not going to deal with empirical merits that specific models may or may not have. They should account for the facts they address, although I know that this is a very strong, and in some cases untenable assumption. Empirical fit and theoretical plausibility are orthogonal properties of models, but both matter for the purpose of evaluation.

SMEED 'S RULE: THE MAGIC OF CURVE FITTING

In 1949 Smeed formulated an empirical relation between the number of fatalities on the road (D), the number of motorized vehicles (M) in a particular geographical region and its population P. The formula $D = c(MP^2)^{1/3}$ has described this relation for many years and in many countries (Smeed, 1949; 1968;

1 Intentionality or *aboutness* is a fundamental characteristic of human activity. Whether or not one may attribute intentionality to the behavior of animals and machines is a matter of considerable philosophical debate. In this article I follow the argument of Dennett (1978, 1987).

1972; see also Adams, 1985). There is, however, a fatal problem with "Smeed's Law:" no one has ever succeeded in offering a plausible explanation in terms of underlying social or psychological processes. Moreover, in the mid-seventies Smeed's law suddenly and drastically broke down. Thus the formula became what it had been in the first place: a magical toy and a brilliant piece of curve-fiting.

Despite the devastating consequences real world statistics since 1973 have had on the face validity of Smeed's conjecture, there are still authors who remain faithful to it. Adams (1985), for instance, claims that "the law is still holding up remarkably well." But alas, he's wrong! We don't understand Smeed's rule and it doesn't describe what it is supposed to describe anymore. It fails on account of both theoretical and empirical validity. Janssen (1986, p. 13) similarly concludes that "Smeed's formula is not suitable as a model of trafic safety. Its empirical validity is insufficient and the formula does not apear to have a conceptual foundation that makes it comprehensible or open to attempts to influence it." After 40 years of service Smeed's rule should finally be put to rest: *Requiescat in pace*!

THE RATIONALITY OF DRIVER BEHAVIOR

With the failure of Smeed's rule in mind I wish to raise the following question. What connection do we actually assume - explicitly but more often implicitly - between the performance of aggrgate models of road user behavior on the one hand, and models of (individual) driver behavior on the other? To illustrate this issue I will consider the Theory of Risk Homeostasis (TRH) proposed by Wilde, and one of the most persistent modeling concepts in the field (Wilde 1982a, 1982b). The concept is attractively simple: accident occurrence at the aggregate level is taken to be a regulatory process by means of which the level of risk in a society is kept constant. This risk is expressed in some measure of disutility or unsafety, e.g. the number of fatalities. When circumstances change in such a way that the objective risk of driving decreases, for instance when the level road maintenance improves, the behavior of the driving population will shift towards more risky forms of behavior.

Unlike Smeed, Wilde has come up with an explanation that makes a lot of common sense. Wilde assumes (a) that risk homeostasis is, in fact, an individual propensity and (b) that the ensemble of homeostatic behaviors of individuals accounts for the homeostatic behavior of the ensemble. Unfortunately neither of these two assumptions is necessarily true, and actually there is a lot to argue against them. Only on the extremely implausible, and much too strong assumption, that the same homeostat is operating im all individuals (rather than weakly, but plausibly assuming that any human behavior is adaptive in some generic sense) can Wilde's model make theoretical sense. But such can be the case only in a world of windowless monads *sensu* Leibniz, all wound up by the Almighty and released at the same time. While Smeed and his followers failed to define what processes can give rise to Smeed's law, the Theory of Risk Homeostasis has come up with just one highly overtaxed principle. Because at the intra-individual level homeostasis is so pervasive that it accounts for almost every form of activity, it is too weak a principle to impose the right kind of constraints on behavior.

In short, ensembles of homeostats do not necessarily produce homeostatic behavior. On the other hand non-homeostatic processes may easily generate homeostatic behavior at the aggregate level.

In his recent work Janssen (Janssen & Tenkink 1988, and also Janssen's presentation at this symposium) has shown that the latter statement is indeed correct. At this point I shall refrain from reiterating Janssen's argument; Janssen convincingly argues that risk homeostasis can be an outcome, at least under special circumstances, of a process of trip utility maximization.

The example of risk homeostasis as an explanatory principle both at the aggegate and the individual level touches directly upon a rather important issue that involves the confounding of levels of explanation.

RATIONAL VS. FUNCTIONAL EXPLANATION

We need to distinguish between two levels of discourse in driver models. The first is frequently called the intentional (or action, or semantic) level, the second the functional (or design, or syntactic) level.[2] Other, related distinctions are competence vs. performance, normative vs. descriptive, or product vs. process.

In making this distinction with respect to driver models, the first general point to observe is that aggregate behavioral models are not really accounts of collective behavior but, almost invariably, descriptions of central tendencies of the behavior of an idealized (but after all *individual*) driver. Such "prototypical" descriptions, based on average behavior of a whole population, a random sample, or perhaps specific segments of the population, rest heavily on the assumption that the average driver will, as a rule, behave rationally (or reasonably, normally, etc.). In other words, given a person's goal and some information about the environment in which the behavior takes place, I can predict with a great deal of success what this person will do, on the simple assumption that he or she will behave *rationally*. To attribute rationality to a behaving system is only a convention, a convenient shortcut to avoid complicated functional explanations that, at least for everyday purposes, would not give much extra predictive mileage. In othert words, whether or not I know if the person is *really* rational (or intelligent, motivated, sensible, or optimally designed) is immaterial and will not affect the quality of my behavioral predictions (Dennett, 1978).[3]

---

2 The distinction has been made by several authors at roughly the same time. This explains the Babylonic terminology. The reader is referred to Dennett (1978, 1987), Newell (1982), or Pylyshyn (1984), for similar expositions.

3 Only when a behaving system acts in a distinctly non-rational fashion, given particular goals and circumstances, we would need to abandon the intentional level of explanation and to turn to the functional process level. Instead of attributing rationality - that is *optimal* design - to such a system, we would begin explaining its behavior in terms of faulty design and malfunctioning.

Individual driver models usually claim to be formulated at the functional (or design) level of theoretical discourse. At this level behavior is described in terms of (mental) functions and processes, operations performed on internally represented facts about the world. However, instead of assuming that the driver is behaving optimally (or rationally), the focus of attention is on actual behavior. Since actual behavior is usually suboptimal, the model is designed suboptimally too, so as to faithfully mimic the driver's performance.

Theorists are constantly facing the risk of confounding these two levels of discourse, the rational level and the design level. Terrible things may happen when they mix, which they frequently do. One of these catastrophes is the introduction of *pernicious homunculi* in one's theory. As an illustration think of Freud's psychoanalytic theory of the human person as a dynamic relation between three sub-personal components, the *Ego*, the *Super-Ego* and the *Id*. Everything would be fine had these three components not been attributed precisely the kind of property (intelligence, motivation, etc.) they are supposed to explain. The consequence will be clear: nothing is gained in the end. Ultimately Freud explained a conscious agent - the person - by postulating (unconscious) agents - homunculi - that were given the same sort of features the conscious agent possessed in the first place instead of intentionally neutral processing features.

My claim is that the Theory of Risk Homeostasis ultimately falls into the trap of homuncularity. In order to *explain* risk homeostasis it *assumes* risk homeostasis in the first place. Janssen, in contrast, successfully avoids this trap; in his model the rational (homeostatic or adaptive) behavior that is to be explained does not sneak in through the back door.

The second issue I wish to bring up is in some sense the complement of the preceding one. It deals with the fact that a good many individual models do *mimic* the elements that a constitute a normative - and therefore aggregate - task analysis.

As an example I take the model that was proposed recently by Böttícher and Van der Molen (1988). These authors attempt to develop a genuinely non-homuncular process theory of the driver. They use descriptive terms that are consistent with my own analysis of driving in terms of a three-tiered division into strategic (planning), tactical (manoeuvring), and operational (skill) aspects of this task (Michon 1971).

The approach adopted by Bötticher & Van der Molen reflects, in my opinion the second kind of confusion that may result from mixing the rational level and the functional level. It is equivalent to the assumption that every step in a cookbook recipe corresponds to a distinct feature of the completed dish. It reflects, in other words, a confusion between syntax and semantics.

I cannot discuss the intricacies of Bötticher & Van der Molen's model in any depth, but refer the reader to several relevant chapters in Rothengatter & De Bruin (1988). One important aspect to be highlighted, however, is Bötticher & Van der Molen's attempt to specify processing units at each of the three levels of the driving task - planning, manoeuvring, skill. As an elementary form of task analysis this distinction has served as a convenient distinction to partition the field of driver behavior research in manageable subdomains. The point that Bötticher and Van der Molen seem to have missed when they set out to implement these three levels in a very detailed process model is that distinctions that are useful to "carve nature at its joints" when we adopt the intentional point of view, need not at all correspond with relevant distinctions at the functional level. Or, to use once again a culinary metaphor: the relevant units of the kitchen syntax do not necessarily map isomorphically or homeomorphically onto the semantics of the dining hall.

STIMULUS-RESPONSE MODELS

By now it should be clear that the problem how to connect behavior at the aggregate level (as tackled by the attribution of rationality, intelligence, etc.) with the sub-personal functional level centers around the proper separation of concepts and terms that belong to and operate at either level. It is, in other words, necessary to specify very clearly and concisely what are elementary processes and building bricks at the functional level, and what are the complex (aggregate) behaviors that are generated by these elementary processes.[4]

This consideration brings me to the concept of rule-based behavior. The idea is familiar since the associationism of James and John Stuart Mill, and it found a detailed but flawed expression in the stimulus-response relations of behaviorism. More specifically we find it instantiated in Fuller's risk avoidance model (Fuller 1984, 1988 and this symposium). Fuller's model is based on the "syntax" of instrumental conditioning. As such it deals at the functional level, with the strength of association between stimuli and responses, and with the corresponding transition probabilities between successive elements in a chain of external (or internal) actions. In Fuller's model these associative mechanisms lead eventually to actions for avoiding harmful or uncertain situations.

It should be pointed out that the model is also distinctly intentional. The behavior that is generated by the internal workings of the rules of association, shaped by reinforcement and punishment, can be perceived as rational: it is maximizing some form of subjective utility.

I have sympathy for Fuller's model because it does (implicitly) separate its assumptions about functional mechanisms (association, reinforcement, etc.) from its assumptions about the intentional, adaptive aspects of behavior. Approach and avoidance are rational relative to the prevailing circumstances and not relative to the principles of association, and Fuller's model is consistent with this position.

There is, however, one pernicious flaw in the type of model Fuller proposes. As soon as we wish to extend it to situations

4 I should point out that this relation can only be specified in terms of sufficient conditions, not necessary conditions. A syntax will generate sentences but not prescribe them.

that are slightly more complex than the one-shot approach-avoidance reaction to a stimulus in an otherwise static situation, we run into a grave difficulty. Again, I cannot deal with this point in nearly enough detail, but the point was already made by Chomsky (1959). in his devastating critique of a book on on "Verbal Behavior" by the arch-behaviorist Skinner, published in 1958. Chomsky's analysis told us, once and for all, that models of the type proposed by Fuller cannot possibly cope with "imbedded" serial behavior. Imbedding or nesting occurs, for instance, when a driver is looking for the next exit from the highway and meanwhile needs to pay attention to the car in front of him, overtakes that car, and then continues his search for the next exit where he left it of. In short, the word *meanwhile* does not exist in the vocabulary of behaviorism. The syntax of avoidance models is that of the Markov process and Markov processes cannot account for the generativity of human behavior.

COGNITIVE SCENARIOS

Cognitive psychology has introduced a new and frankly mentalistic approach to the modeling of behavior. On this view complex situations, are represented internally as frame, scene, script, mental model, or scenario (or whatever other name is used for what amounts to the same thing) and operated upon by computational procedures or rules.

Thinking, understanding and explaining, begins with *reminding*, that is the setting up of a script X for the prevailing situation A, in which you hope to find useful facts that are pertinent to A. If one expects script X to be an explanation for A, but there is no proper fit, then either a memory search is executed for a better fitting script, or X is *tweaked*, transformed to make it more consistent with A. Reminding and tweaking then, are the basic cognitive operations for matching internal representations with prevailing circumstances: "In order to find an explanation, what is required is to find an applicable old pattern, determine to what extent it differs from the current situation and begin to adapt it to fit that situation (Schank, 1986; p. 24). From this it follows that learning is basically failure-driven. As long as all our expectations are met, we do not learn anything new.

Explanation frequently occurs for predictive purposes. The basic planning algorithm is, in fact, to get reminded of a prior plan that is sufficiently close to one's present goals and then to adapt this prior plan to fit the new situation. Like heuristic procedures in general the reminding and tweaking pardigm proposed by Schank (1986) works quite well to generate goal-directed behavior under a good many circumstances. Looked at by an impartial observer the overwhelming impression is that of rationality. Not surprisingly, if you come to think of it, because what we have indicated here by reminding and tweaking is frequently called rationalization in another (intentional) context: reminding and tweaking is indeed close to justification of behavior in retrospect.

In a number of recent studies the cognitive modeling approach as outlined above, has been used in the context of driving, as well as in other tasks that involve decision making under uncertainty. One particularly interesting example is a project of Vlek & Hendrickx (e.g. 1988). Its interest derives in part from the fact that the authors compare a scenario-based model for risk perception with a model that is based on statistical frequency of occurrence. The results of their studies indicate that subjects indeed tend to rely primarily on scenario- or case-based information. Information about the relative frequency of hazardous events works only if the subject has little or no insight in the nature of the events and the processes involved.

This type of research has important consequences for the teaching of risky skills, such as driving. In particular it may offer concrete suggestions about the mental representations or scenarios that drivers require, the operations they need to reason in the context, and the evaluative procedures that should help them to avoid inappropriate scenarios. Especially the latter is a matter of great importance, simply because once a particular scenario has been adopted, one may easily fail to notice a

discrepancy between the plan and the actual situation. The results may be catastrophic.

RULE-BASED MODELS

My peripathetic circumambulatory digressions finally bring me to heart of my matter: the rule-based model or production-system. Chances are still you don't like this! Somehow there is a considerable opposition against rue-based models, and it has kept me busy for quite a while to find out what the motivation can be of those who most strongly argue against such models. The ultimate reason, I figure, is a mistaken belief that the rule-based model falls into precisely the trap that I discussed above: mixing two levels of explanation, the intentional level and the functional level. On that assumption only am I capable of understanding Janssen's (1986) verdict that production-based models are either trivial, or non-trivial because the designer of the model has put in all relevant knowledge by hand. Now, barring the fact that this latter statement would require considerable qualification, Janssen's argument is clearly based on a misunderstanding about the nature of the rules at the intentional or action level and their internal representation at the functional or process level.

Let me illustrate this point by a simple example. Assume that I attribute to a driver the following rule at the action level:

IF the light is red
THEN brake
and stop before the intersection

From the fact that every time the driver approaches a red light he will brake and stop before the intersection, *and* the assumption that the driver will behave rationally, I conclude that the driver must have a mental representation of this rule and that he or she *intends* to use it whenever the circumstances make that a reasonable thing to do. However, it is absurd - so the critics' argument appears to run - to believe that every time an appropriate situation occurs, an interior monologue of the following kind would unfold: "Well, Ok, now there is a light over there, so let's see what color it has. Ok, red. Now there is that rule about red lights that tells me to brake and stop. Nothing else? Ok then, *brake and stop*!"

This is indeed an absurd presumption, but it does not follow that the driver may not *implicitly* be following some version of this rule, behavior being consistent with (but not necessarily translatable into!) this implicit version of the rule. Here is, in fact, the same point I raised earlier with respect to both the model of Wilde and that of Bötticher and Van der Molen. Restated in terms that are appropriate in our present conext: rules that capture behavior at the intentional level (that can be interpreted as rational) need not correspond with rules at the functional level, but rules that determine the processes at the functional level should suffice to generate the kind of behavior that we are willing to call rational (and that is described by the first kind of rule).

THE PRECOGNITIVE LOOP

I consider the resistance to rule-based models surprising for yet an other reason. In the late sixties the idea of the so-called *precognitive loop* in driver models was quite popular. Psychologists accepted it as a useful, if somewhat primitive, means of modeling the complex adaptive behavior of drivers. But the precognitive loop is, in fact, an engineers' trick to incorporate several (structurally different) control components in one model.

Assume for instance that I am driving on a straight, dry road. The outside temperature is dropping. Suddenly, near the canal, the road gets slippery. Some signal - extrinsic, e.g. the road looks wet, or intrinsic, e.g. "ooopps, there I go!" - causes the control of the system to switch to a different dynamic component. Indeed that is what the precognitive loop is: a switch, a behavioral routine selector. Maybe a complicated one, but certainly not a clever one. The clever reader, however, should have understood that the precognitive loop is a very primitive version of a rule-based system that includes such rules as:

IF the road gets wet
THEN switch to cautious driving.

IF the "feel" of steering control becomes unsteady
THEN switch to cautious driving.

A RULE-BASED MODELING FRAMEWORK: SOAR

How should we impose adequate constraints on a rule-based system, without raising the suspicion of triviality? Not by an apriori choice of formal constraints that are too narrow. Perhaps, however, by choosing a structure that is semantically plausible, a structure that offers, in other words a context in which the behavior is supposed to exhibit can be attributed a meaning.

Such a structure - thus far the only of its kind - is Soar, the intelligent architecture developed by Allen Newell and his colleagues at Carnegie Mellon University (Laird *et al.*, 1986; Laird *et al.*, 1987, Newell, *in press*). Soar[5] is literally the *embodiment* of the theory of human problem solving, put forward by Newell & Simon (1972). Having said this, there are two points I need to raise here and now.

The first concerns the form of Newell & Simon's theory, the second relates to the form of Soar. First, problem solving in the view of Newell and Simon amounts to searching a problem space by applying operators to successive states of the problem. A great many activities can be described in these terms. Presently Newell

5Soar (probably) stands for Symbol Operating ARchitecture even takes the most radical position that *any* intelligent activity can be described as a form of problem solving behavior. I wish to point out that I consider this the ultimate extrapolation of my own position, defended since 1974, that *all* travel behavior (which includes driving) may be described as problem solving (e.g. Michon, 1976). This implies that Soar, because it is an exemplary problem solving system, may be considered as a natural medium for modeling driver behavior.

The second issue concerns the form of Soar (or any other rule-based system). Rules that we can understand and that seem so trivial to the critics of rule-based modeling, are in fact *patterns* of interconnected elements in the working memory of the system. Knowledge represented in this format is indeed indeed trivial in the sense that it is presented to the system from the outside in a ready-made form (very much in the way humans derive facts from textbooks and dictionaries. The relevant question is, however, what happens with this knowledge once it is represented internally in the system.

Roughly the following will happen in Soar. Given the top goal, e.g. *travel to the SWOV-symposium in Amsterdam*, and some knowledge about the facts of life, Soar will search its working memory for a pattern of interrelated knowledge elements that has the structure of a problem space and that is, moreover connected with the top goal. One might choose to call such a pattern a scenario). If anything that meets this requirement is found, the system will select it. It may, of course, find nothing that meets the structural requirements of a problem space (e.g. if I happen to be hospitalized because of a severe traffic accident). On the other hand, it may find more than one pattern because it tests the whole content of its working memory in a parallel fashion (e.g. the problem spaces related to driving and public transport respectively). In both cases Soar will find itself in an impasse, which it will recognize as a new problem. Consequently it will specify a new goal. More specifically, depending on the type of impasse it is in, it will select as its next task to solve that particular impasse. It will then search its memory for a pattern that has the appropriate structure of a problem space that can be searched to solve this kind of impasse.

Given the top goal, and having identified a pattern in its memory that qualifies as an adequate problem space for this goal, Soar will then try to find a pattern in its memory it recognizes as a legitimate problem state in the problem space it is currently in (e.g. given that I select the driving problem space *cum* goal "travel to Amsterdam" I will have to reach the motorway).[6]

Finally, Soar's pattern matcher will attempt to find one or more patterns that take the current problem state into a new problem state. Again impasses may occur, and again Soar will set up subgoals to solve these impasses. The pervasive tendency of Soar to set up new goals if it fails to find a suitable pattern in its memory is called *subgoaling*. A second crucial feature of Soar's internal operation is its unique learning mechanism, called *chunking*. If Soar works its way through a series of one or more impasses and finally reaches a solution that allows it to proceed with its primary activity - achieving it toplevel goal - it will first build a new rule that will tell it what to do when it ever hits upon the same impasse again. It will add this rule to its rule-base and subsequently wipe its slate, that is, it will forget all the irrelevant trial and error it went thorough before it found the proper solution.

This should suffice to give a first impression of the internal workings of Soar, as we are currently using it for the modeling complex traffic behavior. Let me just add that our major efforts are presently devoted to the problem of multi-tasking, that is, teaching Soar to learn each of two independent tasks, when the information it gets about these tasks is presented in an integrated fashion. This, in our opinion, is one of the basic features the driving task, one that is closely connected with the

6 This may, of course create another impasse for which I have to find a goal (select which motorway to take) and a problem space (compute whether to take the A28 or the A50 on the basis of some trip utility maximizing criterion).

three-level division - planning, manoeuvring, and skill - discussed earlier in this paper.

At this point I shall only stipulate a few points that summarize the particular value of the rule-based approach to driver models.

(a) The distinction between external rules and internal rules. In this paper I have stressed the difference between rules that describe observable behavior and rules that determine the functional relations that generate such behavior. I have indicated that in various models these two kinds of rules tend to be combined indiscriminately. Using a rule that describes observable behavior as a shorthand expression requires brings with it the risk of infesting one's processing model with pernicious homunculi: technically speaking one is putting the *explanandum* in the *explanans*. Rule-based systems, such as Soar usually make a very neat distinction between the two types of rules, thus reducing the risk that one inadvertently puts the intelligence of the model in by hand.

(b) Production rules and S-R chains. Unlike early versions of associative rules, such as we find them in behavioristic models, production rules can be given the appropriate kinds of constraint. This allows models to display much more complex kinds of activity patterns, including hierarchically nested and imbedded sequences of behavior.

(c) Pattern matching and scenarios. Schema-, frame-, or scenario-theories of performance and action planning are becoming increasingly popular. Research in artificial intelligence has given psychologists a number of tools for developing models along this line. Rule-based systems, such as Soar, provide a natural context for such models. The IF-part of a production rule is in fact a pattern, a script or a schema. If such a pattern is matched, its THEN-part will suggest a certain action. When only one match is found, that action will be executed. However, if several matches are found, an selection problem will arise, and the problem of selecting one from the set of alternatives will have to be solved first, before an action can be taken. If no adequate matching is found, a solution for the resulting impasse in that case will be attempted by some kind of "tweaking" mechanism.

(d) Precognitive loops and the problem space concept. As I pointed out earlier, the precognitive loop is a primitive kind of production system. Now I can be a bit more specific. In Soar the problem space has the same function as the precognitive loop in control systems. Each instantiation of the problem space in Soar is formally equivalent to a point in a model parameter space, that is, it is equivalent with one component in a precognitive control system. The difference is that the problem space is entirely fluid: its charactristics change with the experience of Soar when it is performing its task, depending on the external circumstances (bottom up processing), but also as a result of its internal operations (subgoaling and chunking; top down processing).

(e) Parallelism. Whether human action can be better described by rule-based formalisms or by means of parallel distributed networks (PDP) is a hotly debated issue. It is very likely to come up in driver modeling as well. I am aware of the potential benefits of the PDP approach, but I do not yet see where it would lead road user research. A disadvantage of the PDP-approach is that it may turn out that it can describe anything (there is a relation with the affliction of risk homeostasis). And in the second place rule-based systems can be made to function in a highly parallel fashion too (but without loss of generativity and decidability).

REFERENCES

Adams, J. (1985). Smeed's law, seat belts, and the emperor's new clothes. In L. Evans and R.C. Schwing (Eds.), *Human behavior and traffic safety.* New York: Plenum Press; pp 193-248.

Bötticher, A.M.T., & Molen, H.H. (1988). Predicting overtaking behaviour on the basis of the hierarchical risk model for traffic participants. In J.A. Rothengatter & R.A. de Bruin (Eds.), *Road user behaviour: Theory and research.* Assen: Van Gorcum; pp. 48-65.

Braitenberg, V. (1984). *Vehicles: Experiments in synthetic psychology.* Cambridge, MA: MIT Press.

Chomsky, N. (1959). Review of B.F. Skinner's "Verbal Behavior." *Language, 35,* 26-58.

Dennett, D.C. (1978). *Brainstorms: Philosophical essays on mind and psychology.* Hassocks, Sussex: Harvester Press.

Fuller, R. (1984). A conceptualization of driver behavior as threat avoidance. *Ergonomics, 27,* 1139-1155.

Fuller, R. (1988). Predicting what a driver will do: Implications of the threat-avoidance model of driver behaviour. In J.A. Rothengatter & R.A. de Bruin (Eds.), *Road user behaviour: Theory and research.* Assen: Van Gorcum; pp. 93-105.

Janssen, W.H. (1986). *Modellen van de verkeerstaak: De "state of the art" in 1986.* Institute for Perception TNO, Soesterberg, The Netherlands, Technical Report nr. IZF 1986/C-7.

Janssen, W.H., & Tenkink, E. (1988). Considerations on speed selection and risk homeostasis in driving. *Accident Analysis and Prevention, 20,* 137-142.

Laird, J.E., Newell, A., & Rosenbloom, P.S. (1987). Soar: An architecture for general intelligence. *Artificial Intelligence, 33,* 1-64.

Laird, J.E., Rosenbloom, P.S., & Newell, A. (1986). *Universal subgoaling and chunking: The automatic generation and learning of goal hierarchies.* Boston: Kluwer Academic Publishers.

Michon, J.A. (1971). *Psychonomie onderweg.* Inaugural lecture, University of Groningen. Groningen: Wolters Noordhoff.

Michon, J.A. (1976). The mutual impacts of transportation and human behavior. In P. Stringer & H. Wenzel (Eds.), *Transportation planning for a better environment*. New York: Plenum Press; pp. 221-236.

Newell, A. (in press). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.

Newell, A. & Simon, H.A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.

Rothengatter, J.A., & de Bruin, R.A. (Eds.), *Road user behaviour: Theory and research*. Assen: Van Gorcum.

Schank, R.C. (1986). *Explanation patterns: Understanding mechanically and creatively*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Smeed, R.J. (1949). Some statistical aspects of road safety research. *Journal of the Royal Statistical Society, 112*, 1-24.

Smeed, R.J. (1968). Variations in the pattern of accident rates in different countries and their causes. *Traffic Engineering and Control*, , 364-371.

Smeed, R.J. (1972). The usefulness of formulae in traffic engineering and road safety. *Accident Analysis and Prevention*, , 303-312.

Vlek, C.(A.J.), & Hendrickx, L. (1988). Statistical risk versu personal control as conceptual bases for evaluating (traffic) safety. In J.A. Rothengatter & R.A. de Bruin (Eds.), *Road user behaviour: Theory and research*. Assen: Van Gorcum; pp. ..-...

Wilde, G.J.S. (1982a). The theory of risk homeostasis: Implications for Safety and health. *Risk Analysis, 2*, 209-225.

Wilde, G.J.S. (1982b). Critical issues in risk homeostasis theory. *Risk Analysis, 2*, 249-258.

STATISTICAL MODELS FOR ACCIDENT DATA

Mike Maher, Transport and Road Research Laboratory, Crowthorne, United Kingdom

The following items will be covered:
- log-linear models and generalized linear models
- relations between variables
- the construction of models
- systematic and random variation

This contribution will start from research problems in traffic safety and then discuss the building of models with regard to these problems.
It will also deal with particular difficulties in building models, such as the choice of the disaggregation level, and the small-numbers-problem.

INTERVENTION ANALYSIS AND STRUCTURAL TIME SERIES MODELLING

A.C. HARVEY, University of London, United Kingdom

ABSTRACT

This paper will discuss the rationale underlying the formulation of structural time series models and their relationship to other time series models. The models can be extended to include explanatory and intervention variables. Applications involving the effectiveness of seat belt legislation will be described.

Multivariate structural time series models can be set up and used as a framework for including control groups in the analysis.