safety directions

Predicting and costing road safety outcomes

Working Paper 6

LAND transport safety AUTHORITY

Acknowledgements
This document was prepared by Tony Bliss, Stuart Badger, Bill Frith, Paul Graham, Jagadish Guria, Wayne Jones, Mike Keall, Craig Lauridsen, Joanne Leung, Tui Patterson, Rachel Petrus, Paul Phipps, Lynley Povey, Nigel Rockliffe and Rachel Stokes.

Independent advice and guidance was provided by Professor Peter Vulcan, Research and Safety Services Pty Ltd, Melbourne and Professor Richard Allsop, Centre for Transport Studies, University College, London.

PO Box 2840, Wellington, New Zealand.
Phone 0–4–494 8600, Fax 0–4–494 8601.

October 2000

ISBN: 0-478-20695-X

SAFETY DIRECTIONS

Predicting and costing road safety outcomes

Working Paper 6

October 2000

Foreword

This Working Paper is the sixth in a series of technical documents which form part of the LTSA's *Safety Directions* Development Programme. This Programme is creating the tools needed to underpin New Zealand's road safety funding cycle. In this way it contributes to the achievement of the LTSA's statutory objective of undertaking activities that promote safety in land transport at reasonable cost.

Working Paper 6 describes procedures for predicting and costing road safety outcomes, in particular those being used to set and cost the targets for New Zealand's proposed Road Safety Strategy 2010. It builds on procedures outlined in our previous Working Paper 4 and provides a useful reference source to support Road Safety Strategy 2010 consultations.

This and the other Working Papers in the series are being published to stimulate debate within the wider road safety community. We would welcome your input and comments to assist this process and to help us to further improve the quality of our safety findings.

Reg Barrett
Director of Land Transport Safety

CONTENTS

Introduction

This working paper describes procedures for predicting and costing road safety outcomes.[1] In particular, it describes the procedures we are using to set and cost the targets in New Zealand's *Road Safety Strategy to 2010*. It will also be of interest to road safety agencies in other jurisdictions which need to formulate road safety targets.

This working paper is linked to others of this series. Working paper 2 (LTSA 1998a) discusses the safety funding cycle; the procedures described in this paper assist in the target-setting stage of the cycle. Working paper 4 (LTSA 1998b) describes the theory underlying the setting of road safety targets; this paper shows how the theory was developed and put into practice. Lastly, working paper 7 (LTSA 2000) describes how specific interventions affect road safety outcomes; these relationships are incorporated into the road safety model described in this paper.

Prediction of outcomes

This section describes a computerised simulation model that predicts road safety outcomes. The model was developed by New Zealand's Land Transport Safety Authority to assist in formulating targets for the *Road Safety Strategy 2010*. It takes information about the factors that affect road safety, and predicts how they jointly affect the number of casualties of various kinds and thereby the social cost of crashes.

There are two versions of the model. The generalised model offers precision at the expense of onerous data demands. But if these data demands cannot be met, there is a simplified model which is a special case of the generalised model. It was this that we used in New Zealand. We describe both next.

Generalised model

The generalised model (*figure 1*) is potentially more accurate and versatile than the simplified version. Only a lack of data inhibits its use in New Zealand (and, potentially, elsewhere). This may not always be so. Traffic data are increasingly being collected automatically by electronic means. If this becomes widespread, the generalised model will be a practicality. It is for this reason that we discuss it here.

Current traffic volume and risk

First, we classify the traffic on the network into unique categories distinguished by the attributes of the traffic that are relevant to road safety. An attribute is deemed relevant if it mediates the link between an intervention and its effect on road safety outcomes (see discussion on dose-response below). Relevant attributes relate to:

- the vehicle (type, age, condition, etc)
- the driver (age, blood-alcohol content, restraint-wearing, speed of travel, etc)
- the road environment (design, region, weather, surface condition, etc).

---

[1] An outcome is any measure of road safety performance, such as the number of casualties of crashes of various kinds, and their social cost. It is our aim to *reduce* this outcome.

Figure 1

Logic of the generalised model

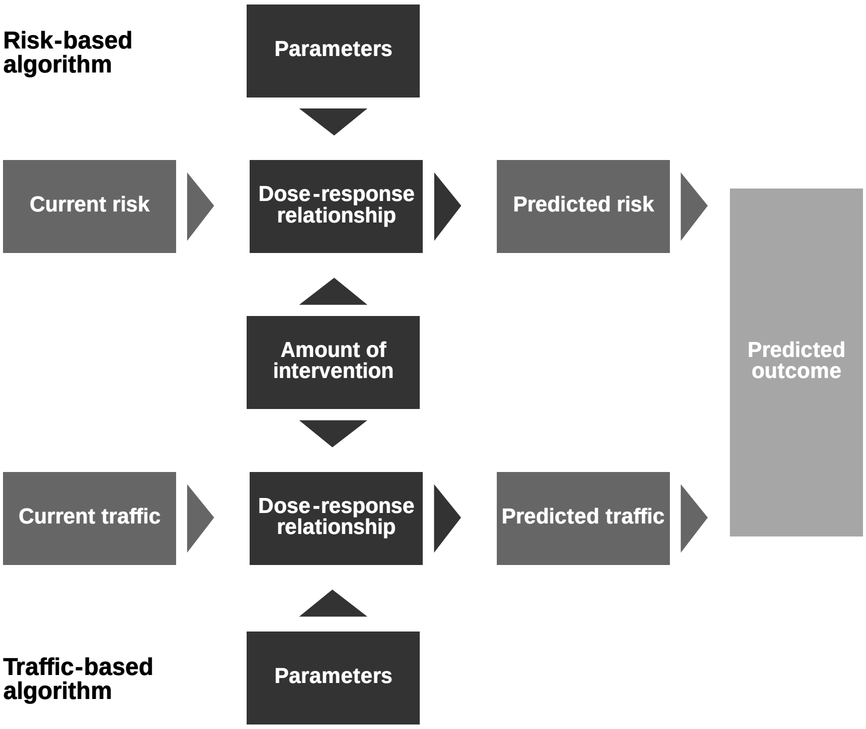

In practice not all of these variables will be used, either because the required information is not available, or because their effect on outcomes is not well enough understood to permit their incorporation into the model. This naturally entails some loss of accuracy.

Let every unit of traffic on the network (vehicle-km, say) be characterised by n attributes, each of which can have a number of discrete values. We then define the set I of possible traffic categories as the product set formed from the n sets of attribute values I_1, I_2,..., I_n:

$$I = I_1 \times I_2 \ldots \times I_n \quad \text{(1)}$$

We then define the set V of traffic volumes, of which each element V_i is characterised by the i^{th} element of product set I:

$$V = \{V_i \mid i \in I\} \quad \text{(2)}$$

Likewise we define the set R of road safety risks,[2] of which each element R_i is characterised by the i^{th} element of product set I:

$$R = \{R_i \mid i \in I\} \quad \text{(3)}$$

Dose-response relationships

Next, we characterise the links between interventions (the 'dose') and road safety outcomes (the 'response')[3]. Dose-response relationships can be represented in

---

[2] Risk is defined as road safety outcome (however measured—see later) per unit of traffic.

[3] The term 'dose-response' is drawn from epidemiology, and describes how large an effect (or 'response') is produced by each level (or 'dose') of an intervention. In previous working papers of this series, for instance LTSA (1996), dose-response relationships are referred to as 'benefit functions'. The former term sees road trauma as an epidemic to be mitigated; the latter as an economic cost to be reduced. Both are apt.

either of two ways. In the risk-based algorithm an intervention is considered to reduce the risk pertaining to certain traffic categories; in the traffic-based algorithm it is considered to reduce the volume of traffic in certain risk categories. Many interventions work in both ways simultaneously: they both alter risk for certain categories of traffic and alter the volume of such traffic.[4] But some lend themselves to specification solely in terms of altered risk or altered traffic volume. The choice of specification is a matter of modelling convenience.

Risk-based algorithm

Let F_{ij} be the risk adjustment factor by which intervention j affects the risk R_i attributable to traffic category i. Since interventions normally combine multiplicatively,[5] the predicted risk R_i^* is given by

$$R_i^* = R_i \prod_j F_{ij} \quad \text{.......} (4)$$

We establish the risk adjustment factor F_{ij} as follows:

$$F_{ij} = f_{ij}\left(R_i, Q_j\right) \quad \text{.......} (5)$$

where function f_{ij} is a computable rule that determines how intervention j affects risk adjustment factor F_{ij}; R_i is the prior risk attributable to traffic category i; and Q_j is the amount of intervention j applied to the network.[6]

Traffic-based algorithm

Let G_{ij} be the volume adjustment factor by which intervention j affects the traffic volume V_i of type i. The predicted traffic volume V_i^* is given by

$$V_i^* = V_i \prod_j G_{ij} \quad \text{.......} (6)$$

We establish the volume adjustment factor G_{ij} as follows:

$$G_{ij} = g_{ij}\left(V_i, Q_j\right) \quad \text{.......} (7)$$

where function g_{ij} is a computable rule that determines how intervention j affects volume adjustment factor G_{ij}; V_i is the prior traffic volume of type i; and Q_j is the amount of intervention j applied to the network. For simplicity, the total volume of traffic is exogenously determined in this formulation of the model, and is not subject to influence by particular interventions.[7] Hence traffic that is removed from one category must be placed in another. This can be accomplished as part of the modelling procedure.

---

[4] Consider drink-driving, for instance. Police patrols both discourage such trips and make drink-drivers drive with more care. This reduces the number of such trips, and makes those that remain safer.

[5] LTSA 1998b explains why this is so, and discusses exceptions to the rule. Sometimes interventions combine additively, which can be readily accommodated with minor changes to the logic of the model. Occasionally interventions combine synergistically, which is usually best handled by combining interventions into indivisible packages, for instance enforcement and its associated publicity.

[6] One may omit prior risk as an independent variable if the risk adjustment factor does not significantly depend on the level of risk to which it is applied—that is, if an intervention reduces risk by the same proportion whatever its prior level. An analogous comment applies to the traffic-based algorithm (see next). But see also the discussion of the simplified model.

[7] Some interventions may indeed alter the volume of traffic. For instance, drinkers could be encouraged to use taxis or public transport, or to drink at home. Although not shown here, this effect could be readily incorporated into the model provided we could quantify it.

Reducing the number of dose-response relationships

There are typically thousands of traffic categories, and the discussion so far implies that there must be a dose-response relationship for each. This is clearly impractical. Fortunately, it is possible to reduce greatly the number of dose-response relationships required.

In practice, dose-response relationships never depend on all n traffic attributes that are used to form the set I of traffic categories, and most use fewer than three. Suppose a given intervention depends on a proper subset of attributes consisting of n' attributes (see *table 2*). We then define the set I' of traffic categories required for this particular intervention as the product set formed from the n' sets of attribute values I_1, I_2,..., I_n:

$$I' = I_1 \times I_2 \ldots \times I_{n'} \qquad (8)$$

It can be demonstrated that (1) I' must have fewer elements than set I, and (2) there exists a computable function that maps set I into set I'. In other words, there are fewer dose-response relationships than traffic categories, but for every traffic category there still exists a single dose-response relationship. Furthermore, since individual dose-response relationships typically make use of so few traffic attributes, the number of relationships needed to characterise each intervention is rarely more than a handful.

Predicted outcome

Having predicted both the risk and volume of traffic in each category, we calculate predicted road safety outcome O_i^* for traffic type i by multiplying the predicted traffic volume V_i^* by the predicted risk R_i^*:

$$O_i^* = R_i^* . V_i^* \qquad (9)$$

This formula should not be taken to imply that every intervention alters both the risk and the volume of traffic in each affected traffic category. In cases where an intervention leaves either risk or traffic volume materially unchanged, we may regard the current risk or traffic volume as equal to the predicted; and it is the current value that is used in the above formula.

Note that the predicted outcome is measured in the same units as those used to express the numerator of risk. For instance, if risk is expressed as social cost per unit of traffic (cents per vehicle-km travelled, say), then the outcome is expressed in terms of social cost. Alternatively, if risk is expressed as the probability of a specified type of casualty per unit of traffic (deaths per billion vehicle-km, say), then the outcome is expressed in terms of persons killed. We shall typically use both kinds of outcome measure.

Simplified model

The simplified version of the model (*figure 2*) is less data-hungry than the generalised. This offers advantages because it does not require disaggregated estimates of traffic volume or risk, both of which are notoriously hard to measure across a large and complex network. Its data requirements are in fact of a kind that is satisfied by the electronic crash databases that exist in nearly all developed countries and in some developing ones.

A special case

The simplified model is a special case of the generalised model. In the generalised model we simulate an intervention by adjusting the risk and/or volume of each traffic category according to rules embodied in the model. Then, to get the predicted outcome, we multiply predicted risk by predicted traffic volume. Since risk is defined as outcome divided by traffic volume, we are in effect first dividing by traffic volume, then multiplying by the same thing.

It is instead mathematically simpler to factor the outcome directly, and it avoids our having to estimate the traffic volume in each traffic category, which is difficult. There are two objections to this, neither very compelling, which is why in the end we can disregard them. First, by introducing traffic volume into the calculation, we allow the dose-response relationship to depend on prior risk wherever risk is a function of traffic volume. This is desirable where a given intervention has one effect where risk is high and another where it is low. In omitting traffic volume from the model we limit our freedom to describe dose-response relationships of this kind; we are confined to saying that a given intervention causes the same (proportional) change in the number of casualties however frequently those casualties occur. But at this stage our dose-response relationships are insufficiently sophisticated to take advantage of this freedom, so it is not missed. The other reason to use traffic volume in the model is that it allows us to employ a traffic-based algorithm. But again, this is no more than a modelling convenience.

Figure 2

Logic of the simplified model

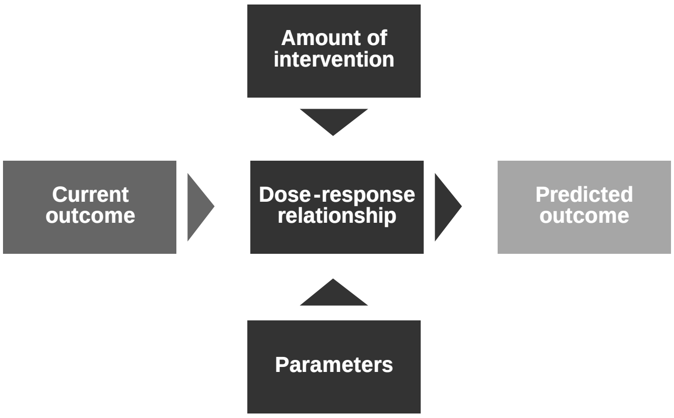

On the other hand, the advantages of *not* using traffic volume in the model are substantial. The volume of traffic on each road segment in the network can often only be guessed at. Moreover, it must be estimated not just globally but disaggregated by the attributes used in the model. Right now we have no easy way of doing this.

The simplified model was devised to overcome these problems. It has all the advantages of excluding traffic volume, which are large; while tolerating the disadvantages, which are currently small. But things will not always be this way. Traffic data will become easier and cheaper to obtain as electronic collection methods become widespread. We are a long way from this yet, but when it happens the generalised model will become a practicality. Until then, however, the simplified model offers a practical alternative.

Current outcome

First, we classify the current outcome[8] into unique categories distinguished by attributes relevant to road safety. This is analogous to the classification of traffic in the generalised model. We define the set O of outcomes, of which each element O_i is characterised by the i^{th} element of product set I (as defined for the generalised model):

$$O = \{O_i \mid i \in I\} \qquad (10)$$

Dose-response relationships

In the simplified model an intervention is considered to reduce the outcome pertaining to certain, typically high-risk, traffic categories. Let F_{ij} be the outcome adjustment factor by which intervention j affects the outcome O_i attributable to traffic category i. Since interventions normally combine multiplicatively, the predicted outcome O_i^* is given by

$$O_i^* = O_i \prod_j F_{ij} \qquad (11)$$

We establish risk adjustment factor F_{ij} as follows:

$$F_{ij} = f_{ij}(Q_j) \qquad (12)$$

where function f_{ij} is a computable rule that determines how intervention j affects outcome adjustment factor F_{ij}; and Q_j is the amount of intervention j applied to the network.

Outcomes for 2010

This section describes how we employed the simplified version of the predictive road safety model assist in the development of New Zealand's proposed *Road Safety Strategy 2010* (*figure 3*).

Data requirements

The model uses a historical dataset to characterise the current safety performance of the road network. The dataset consists of individual records for all reported casualties in the period 1996 through 1998, some 40 000 in total. To lessen the computational load, casualties with identical attributes were grouped into some 12 000 homogeneous categories.

Accuracy

The historical dataset is regarded as a sample representing the current underlying structure of New Zealand's road safety performance.[9] Thus the number of casualties in each category is only an estimate of the typical—but unknown—number for a given period under current circumstances. But with so many categories in comparison to the absolute number of casualties, each category contains on average only about three casualties. Hence we must naturally be concerned about the accuracy of the estimates.

---

[8] In principle it does not matter how the outcome is measured. It is normally convenient to measure outcome in terms of casualties of a particular type (fatalities, serious and minor injuries, say) rather than as social cost. The data are generated in this form anyway; moreover it allows one to apply different dose-response algorithms to each type of casualty.

[9] Strictly speaking it represents our current road crash problem, not our safety performance. But since the magnitude of the problem depends in part on the interventions we have undertaken to remedy it, it can be said to reflect our safety performance indirectly.

Figure 3

Logic of the model as used for proposed Road Safety Strategy 2010

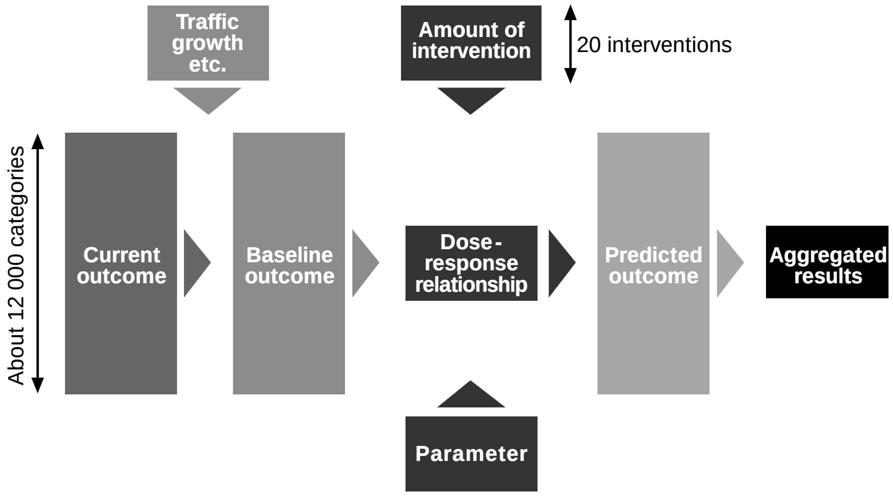

We tackled this problem in two ways. First, the analytical method itself provided a large measure of protection. We have already noted that each intervention is characterised by far fewer dose-response relationships than categories of casualty. Hence each dose-response relationship typically affects not several casualties but hundreds or thousands. To put it another way, the historical dataset contains over 12 000 categories, but in practice it can be regarded as containing far fewer. Since different dose-response relationships require the data to be disaggregated in different ways, it is convenient to categorise the data into a large number of small categories and to aggregate them as needed.

Second, we based the sample on three years' data, not one. New Zealand's road safety outcome changed relatively little over the period, so we estimated our base year from the average. This approach is acceptable provided the data are not contaminated by uncontrolled variables that changed over the period. But are they? In fact traffic grew faster than strategic police hours[10] during the period, although safety outcomes remained within a narrow band. This can be explained by a small secular efficiency trend, probably due to improvements in the road network, the vehicle fleet (*see later discussion*), and the deployment of police resources. This trend was allowed for in our predictions.

Sourcing

The dataset of current casualties was extracted from New Zealand's comprehensive electronic database of road crashes known as the Crash Analysis System (CAS). The contents of this database are similar to those of its analogues in other developed countries and will not be elaborated here. It is sufficient to say that it contains many of the data attributes that are required by the dose-response relationships in the model. Naturally this entailed some compromise in the specification of dose-response relationships.

Casualty attributes

Ten casualty attributes were extracted from CAS (*table 1*). Most were selected because they were required in at least one dose-response relationship. For

---

[10] The number of strategic police hours increased slightly over the last three years while the total number of road safety police hours remained constant.

instance, we want to know if alcohol was involved in the crash that caused the casualty, as these are the very casualties that are most affected by interventions to combat drink-driving. A few attributes were selected because we wished to use them to disaggregate road safety outcomes, not because they were required by the model; for instance the sex of the casualty is included so we can express outcomes in terms of male and female casualties, if desired. The casualty's age was included so that we could take account of the changing age profile of the population. Lastly, some attributes would be very valuable in dose-response relationships but are not available. Speeding and restraint use are examples, but are too unreliable to be of use: it is a rare driver that will admit to speeding or failing to wear a restraint, and when the driver dies, the information is hard to establish independently.

Note that we distinguish between attributes of the casualty and those of the driver.[11] For instance, we record the age of the casualty so that we can accommodate demographic changes in the population. We record the age of the driver for quite a different reason: we want to know if young or elderly drivers were involved as there are specific interventions aimed at these age groups.

Table 1

Contents of casualty dataset

Data item	Values		
Attributes of casualties			
Region	14 admin. regions		
Road type	Motorway	State highway	Other open road
	Major urban road	Minor urban road	
Road user	Driver (light veh.)	Driver (heavy veh.)	Pax (light veh.)
	Pax (heavy veh.)	Pedestrian	Cyclist
	Motorcyclist		
Age of casualty (yrs)	0–14	15–19	20–24
	25–64	65–74	75+
Sex of casualty	Male	Female	
Alcohol involvement	Yes	No	
Heavy vehicle involv't	Yes	No	
Young driver involv't	Yes	No	
Elderly driver involv't	Yes	No	
Severity	Fatal	Severe	Minor
Quantification of casualties			
Casualties	No. of casualties		
Social cost	Amount in dollars		

Dose-response relationships

The model embodies dose-response relationships for 18 interventions, and can readily accommodate more. Since interventions are evaluated sequentially, there is no theoretical or computational limit to the number that can be incorporated into the model. Each dose-response relationship uses as its independent variables a subset of attributes from the casualty dataset (*table 2*).

---

[11] We recognise of course that in a purely technical sense driver attributes are also by definition casualty attributes as they describe an aspect of the crash that gave rise to the casualty—namely the driver or drivers involved.

Table 2

Specification of dose-response relationships, by intervention

	Attribute							
	All casualties	Severity	Region	Road type	Road user	Alcohol involvement	Heavy veh. involvement	Young driver involvement
The road environment								
Standards and rules								
Blackspot treatments		■						
Existing construction	■							
Expanded construction	■							
Trauma management		■						
Enforcement								
Speed management (urban)		■		■				
Speed management (open road)		■	■	■				
Compulsory breath-testing						■		
Restraint-wearing		■			■			
The vehicle								
Light vehicles	■				■			
Heavy vehicles		■					■	
The road user								
Standards and rules								
Driving age								■
Stricter licensing conditions						■		■
Reduced BAC						■		
Zero BAC young drivers						■		
Use of headlights		■						
Enforcement								
Vehicle impoundment		■				■		
Licence suspension		■				■		
Alcohol interlocks		■				■		

Specification

The ideal dose-response relationship expresses the outcome of an intervention as a continuous function of the amount applied. So far, dose-response relationships are assumed to be of this kind. But it is often difficult or impossible to express dose-response relationships in this way. One reason is that we rarely understand them well enough. The other is that many interventions are by nature discrete; they do not lend themselves to continuous analysis: for instance, either we impose a certain vehicle standard or we do not. Our dose-response relationships are therefore commonly expressed as sets of point observations of the form: amount X of intervention *Y* will produce outcome *Z*.

The reasoning underlying the point observations is not currently incorporated in the model. Since interventions differ greatly in their underlying mechanisms, each was best analysed *ad hoc* and independently. A working paper in this series (LTSA 2000) shows how they were derived.

Derivation

The model's dose-response relationships come from various sources (see *LTSA 2000*). Many are derived from the experience of jurisdictions that have tried the intervention in question. Some incorporate well established research results, for instance the link between casualties and traffic speed.

Outcomes

The aim is to predict the road safety outcome in 2010 resulting from the application of selected interventions. But between now and 2010 other factors will also cause the outcome to change. Therefore our interventions will not be acting on the road safety outcome that exists now, but on the one that would exist in 2010 in their absence. Below we discuss the factors that take us from the current outcome to the target outcome (*figure 4*).

Current outcome

The *Current outcome* outcome attributable to traffic category *i*, O^{act}_i, is the estimated situation at the beginning of the strategy period, and the starting point for the analysis. It is based on records for all reported casualties in the period 1996 through 1998, these being the latest available data for complete years.

Baseline 2010

Two factors act to increase social cost: traffic growth[12] and speed creep (the tendency for traffic to travel at ever higher speeds). Unless countered by appropriate interventions they would raise social cost to the *Baseline 2010* level. We therefore allowed for them before simulating the impact of interventions.

The *Baseline 2010* outcome attributable to traffic category *i*, O^{base}_i, is calculated from the *Current* outcome:

$$O_i^{\text{base}} = O_i^{\text{act}}(1+v_i)(1+s_i) \quad \text{..} \quad (13)$$

where v_i and s_i are the proportional changes in social cost for traffic category *i* due to change in traffic volume and speed creep respectively.

Traffic growth

During the planning period we expect traffic to grow substantially (LTSA 2000). This will raise social cost pro rata. The volume of traffic depends on the level of population and its mobility.[13] Both are expected to change between now and 2010, and the changes are expected to vary across age-groups and regions. *Baseline 2010* outcomes were predicted by pro-rating *Current* outcomes by the forecast proportional changes in population and mobility. Population forecasts by region and age cohort came from official sources, and future mobility was extrapolated from the trend over recent years.

---

[12] In LTSA (1996: 14) we argue that road safety outcomes are directly proportional to traffic volume.

[13] Mobility is the average road distance travelled per person per year. In New Zealand it is about 10 000 km/person-yr, roughly the same as in many other developed countries.

Figure 4

Composition and derivation of social cost outcomes

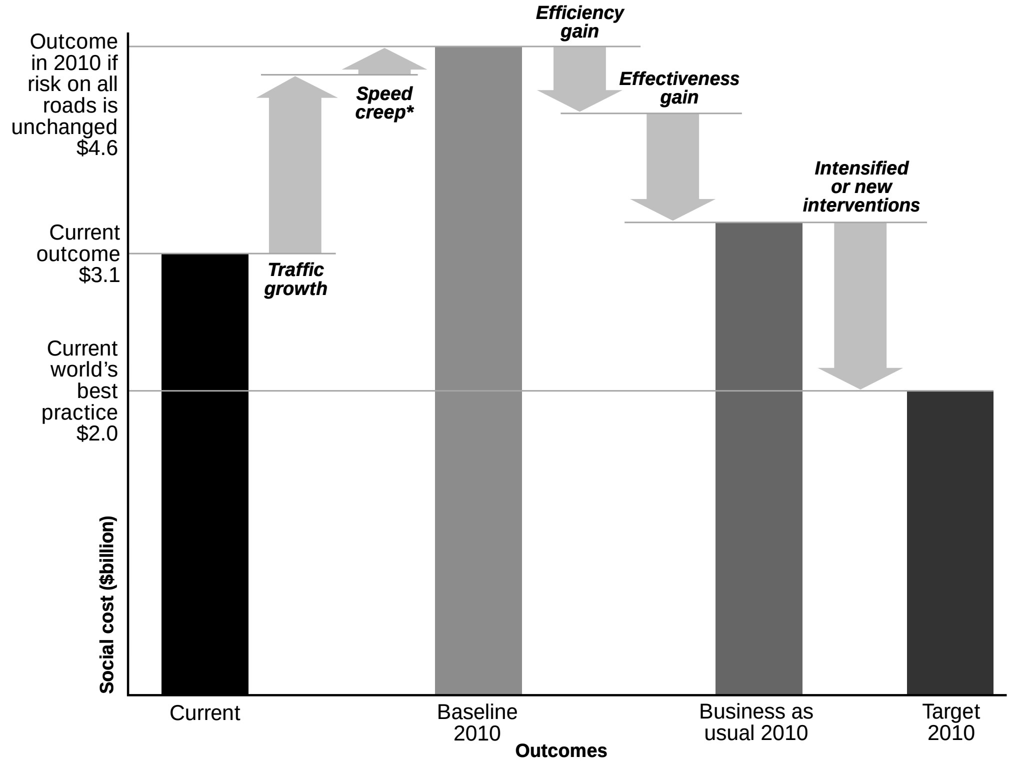

* Unless checked, average traffic speeds will creep up, and will erode safety and raise social cost.

Speed creep

Because modern cars are faster and more comfortable to drive, in the absence of countermeasures average traffic speeds would not just stay where they are—they would creep upwards, and if unchecked would erode safety and raise social cost (LTSA 2000). They have not done so in New Zealand in recent years only because we have stepped up our speed management effort.

Business as usual 2010

Although the *Baseline 2010* outcome is a useful theoretical construct, it is unlikely to arise in practice as some of our existing road safety interventions are becoming more efficient and some more effective without any further action on our part.[14] Among these we include the rising quality of the road network (*see later discussion*). So even if we continued to employ the same mix and level of resources as now, social cost would decline from the *Baseline 2010* level to *Business as usual 2010.*

The *Business as usual 2010* outcome attributable to traffic category i, O^{bau}_i, is calculated from the *Baseline 2010* outcome:

$$O_i^{\text{bau}} = O_i^{\text{base}}(1+e)\prod_{j\in J}(1+q_{ij}) \quad \text{......................................} \quad (14)$$

---

[14] An intervention is said to be more *efficient* if it produces the same road safety outcome for less cost; and more *effective* if it produces a better road safety outcome for the same cost.

where e is the proportional change in social cost due to a change in efficiency, q_{ij} is the proportional change due to intervention j for traffic category i, and J is the set of interventions that undergo a gain in effectiveness.[15]

Efficiency gain

Road safety interventions in general are becoming more efficient (LTSA 2000). Learning explains some of this: with the passage of time we are learning to get more from the resources at our disposal. Better management practices and analysis tools allow us to place resources where they have the more impact, and to use them more productively once there. Another reason is improved technology: there can be no doubt for instance that enforcement devices and safe road construction methods are advancing steadily. Also, road users may be increasingly internalising safe behaviours, that is, they behave safely not because they are forced to but because they choose to.

Effectiveness gain

Certain existing interventions are set to become more effective (LTSA 2000). One reason is that they have already been put in place but their benefits are still to be felt. Vehicle standards for instance only take effect when new vehicles find their way into the national fleet. Other examples are interventions that are committed or have already been imposed but whose impacts had not been reflected in the historical dataset we used for forecasting purposes.

Interventions of a capital-building nature also give rise to growth in effectiveness. Suppose we upgrade a segment of road . The benefit recurs not once but in *every* subsequent year. Suppose we then upgrade another segment of road. In subsequent years we experience the benefit of both road segments. The longer we go on investing, the greater the annual stream of benefits, since it consists not only of the stream from the latest upgrading but also of all the upgradings that preceded it. The effect of this appears as a growing benefit arising from a constant annual expenditure.

Target 2010

Continuing 'business as usual' will reduce social cost substantially, but not enough to achieve what we regard as an acceptable road safety outcome. We therefore first used the model to explore the realm of the possible—that is, the range of road safety outcomes that might reasonably be achieved with the means at our disposal. Some existing interventions could be intensified and others could be introduced for the first time. We found that with sufficient effort and at reasonable cost we could match the road safety performance of the current world leaders in road safety, which have about 1.2 fatalities per 10 000 vehicles. After some discussion we selected this 'current world's best practice' as our target outcome—*Target 2010*. Finally, we used the model to develop several alternative options, differing in their mix of interventions, for reaching our chosen target.

The *Target 2010* outcome attributable to traffic category i, O^{tgt}_i, is calculated from the *Business as usual 2010* outcome:

$$O_i^{\mathrm{tgt}} = O_i^{\mathrm{bau}} \prod_{j \in K} (1 + q_{ij}) \quad \text{..} \ (15)$$

where q_{ij} is the proportional change in social cost due to intervention j for traffic category i, and K is the set of new or intensified interventions.

---

[15] Since the change in efficiency reduces the road safety outcome, e is negative; and in this formulation is assumed to be proportionally the same for all traffic categories. Likewise, since interventions reduce the road safety outcome, q_{ij} is negative.

Options

Within limits we can choose what interventions to conduct and to what degree. We evaluated three combinations of interventions—or 'options'. All emphasise interventions relating to the road environment, since this is where most safety gains are to be made, but differ in the stress they place on enforcement as against road engineering interventions.

- The *Enforcement emphasis* option is weighted towards enforcement. Because enforcement is relatively less costly than road construction, this option demands little extra funding but could be unpopular.
- The *Mixed* option combines elements of both the *Enforcement emphasis* and *Engineering emphasis* options. Any mix of elements is possible; we chose to evaluate one that occupies the middle ground.
- The *Engineering emphasis* option is weighted towards engineering: the improvement of existing roads and building of new ones. Because roads are costly to build, this option demands much more funding but it is likely to be popular with road users.

Aggregation of results

We are dealing with approximately 12 000 traffic categories, each of which is potentially affected by 18 interventions and various other factors, namely traffic growth, speed creep, and an efficiency gain. This is capable of producing about a third of a million individual results. How are they to be combined?

Combining the effects of interventions and other factors

The effects of interventions and other factors are combined multiplicatively within outcome categories because they affect the same crashes. For instance, if intervention A reduces the number of casualties by 20%, and intervention B by 20%, the combined reduction is 36%,[16] not 40%. This is because no crash can be prevented (or mitigated) twice. Thus, for each traffic category:

$$\begin{bmatrix}\textit{Target 2010}\\ \text{outcome}\end{bmatrix} = \begin{bmatrix}\textit{Current}\\ \text{outcome}\end{bmatrix} \times \left[\text{Factor 1}\right] \times \left[\text{Factor 2}\right] \times \ldots \times \left[\text{Factor } n\right]$$

Some interventions are an exception to this general rule. When interventions apply to different types of traffic their effects must be combined additively, not multiplicatively. For instance, the effects of bridge treatments must be added to those of intersection treatments, since crashes on bridges do not overlap with those on intersections.[17] Likewise the effects of helmets must be added to those of seatbelts, since road users do not use helmets and seatbelts at the same time.

---

[16] This is given by: 1 - [(1 - 0.2) x (1 - 0.2)]

[17] Note that if the position of the crash—at a bridge or at an intersection—were included as an attribute in the dataset of current outcomes these interventions would no longer be an exception. Crashes on bridges and those at intersections would be placed in separate outcome categories, and would therefore automatically be combined additively by the model.

Combining outcomes for different traffic categories

The model combines the effects of different interventions multiplicatively *within* outcome categories but additively *between* them. This is because by definition their effects on different categories are distinct and separable. Thus, for each type of outcome:

$$\begin{bmatrix}\text{Outcome for}\\ \text{all traffic}\end{bmatrix} = \begin{bmatrix}\text{Outcome for}\\ \text{traffic category 1}\end{bmatrix} + \begin{bmatrix}\text{Outcome for}\\ \text{traffic category 2}\end{bmatrix} + \ldots + \begin{bmatrix}\text{Outcome for}\\ \text{traffic category } n\end{bmatrix}$$

Presentation of partly disaggregated results

The results of the model are capable of disaggregation in precisely the same way as the dataset of current outcomes. However, the more one disaggregates, the greater the coefficients of variation of the disaggregate estimates; normally this limits disaggregation to no more than two attribute dimensions.

Resource cost

This section shows how we estimated the resource cost of achieving the *Target 2010* outcome. In that exercise we defined resource cost as costs borne by government and its agencies; compliance costs borne by road users were excluded.[18]

Method

The resource cost of achieving a stipulated outcome is given by the reduction in social cost divided by the benefit–cost ratio of the interventions employed to achieve it:

$$\begin{bmatrix}\text{Resource}\\ \text{cost}\end{bmatrix} = \frac{[\text{Reduction in social cost}]}{[\text{Benefit-cost ratio}]}$$

Estimates of social cost are provided by the predictive road safety model, discussed above, and estimates of benefit–cost ratios come from LTSA research. But though simple in principle, the method accommodates the following complexities in its application.

Only some things need funding

The method recognises that road safety outcomes are brought about by various factors, only some of which need government funding. We accommodate this by considering only the reduction in social cost between the *Business as usual 2010* outcome to the *Target 2010* outcome (*figure 4*), as this is the only one that needs funding (other levels of outcome are achieved through factors that are effectively costless to the government).

Interventions differ in their benefit–cost ratios

The method recognises that interventions differ in their benefit–cost ratios. To achieve the *Target 2010* outcome we must impose two broad types of intervention—enforcement[19] and engineering—both of which have different benefit–cost ratios. We accommodate this by considering each type separately.

---

[18] Compliance costs, which may be substantial, include travel time losses as a consequence of slower traffic speeds, and the costs of vehicle safety features. In principle, compliance costs could be included in the method for estimating resource costs described here.

[19] For brevity we refer to all non-engineering interventions as enforcement, even though a few are not.

Interventions differ in their timing

The method recognises the timing of interventions. For instance, we plan to emphasise enforcement early in the strategy period as it offers quick results; later, road engineering, which has longer lead times, will begin to take effect. The method also allows us to reflect changes in benefit–cost ratios, which are expected to decline during the strategy period, since road safety often exhibits decreasing returns to scale.

Capital versus recurrent interventions

The method recognises that some interventions, mainly road engineering, generate a stream of benefits, whereas most others, including most enforcement, produce a benefit while the intervention is actually being conducted or for a short time thereafter.[20] So a constant rate of spending on road engineering continuously reduces the level of road trauma as the benefits from previous years accumulate; but a constant rate of spending on enforcement produces a one-off reduction in road trauma (which, however, is maintained as long as the enforcement is maintained at the same level).

Modelling resource costs

We used a computer model to calculate resource costs. The model allows us to vary the mix and timing of road safety expenditures, and accommodates declining marginal benefit-cost ratios. Since costs and benefits are spread over a number of years, the model discounts the combined cost stream to a stipulated base year.

Enforcement

Enforcement interventions reduce social cost from the *Business as usual 2010* outcome in year t, S^{bau}_t, to the *Target 2010 (enforcement component)* outcome, S^{enf}_t; and have an average benefit-cost ratio of b^{enf}_t. Hence their resource cost in year t is given by

$$C_t^{\text{enf}} = \frac{S_t^{\text{bau}} - S_t^{\text{enf}}}{b_t^{\text{enf}}} \quad \text{...} \ (16)$$

Engineering

When calculating the cost of engineering we must add two refinements to the logic we used for enforcement interventions.

- *The benefit-cost ratio depends on social cost.* Specifically, for every additional dollar of engineering expenditure, the safety benefit is a constant proportion ρ of the social cost that would otherwise occur.[21]

---

[20] Not always. Some enforcement interventions produce a 'ratchet' effect whereby a behaviour (such as seatbelt wearing), once acquired, is maintained after the enforcement that caused it is removed.

[21] The constant, ρ, is derived as follows. Let b^{eng} be the incremental benefit-cost ratio of engineering improvements in the base year. Then for each dollar of engineering expenditure the total benefit is b^{eng} dollars, being the present value of benefits over the economic life of the road (25 years in this analysis). Transfund NZ estimate that in the base year 35% of the benefits from engineering expenditure are safety benefits. Hence the total value of safety benefits, SB, from every dollar expended is $0.35 * b^{\text{eng}}$ dollars. Let φ be the present value of a \$1 annuity over 25 years at a discount rate of 10%. Then the annual value of safety benefits from each dollar of engineering expenditure is SB/φ dollars. Suppose the social cost in the base year is S_0. Then the proportion of social cost reduced by a dollar of engineering investment, ρ, in one year is $SB/(\varphi S_0)$.

- *Road engineering interventions generate a stream of benefits that endures for many years.* Since each dollar of engineering expenditure produces a benefit stream, the safety benefit in year *t* arises not just from expenditure in that year alone but from expenditure throughout the entire preceding period. Therefore the benefit due to expenditure in year *t only* is given by the totality of benefits in year *t*, $(S_t^{enf}-S_t^{tgt})$, less the sum of all benefits, $\sum_{i=0}^{t-1} B_i^{\mathrm{eng}}$, due to expenditures in previous years.

Road engineering interventions reduce social cost from the *Target 2010 (enforcement component)* outcome in year *t*, S_t^{enf}, to the *Target 2010* outcome, S_t^{tgt}. Hence the engineering resource cost in year *t* is given by

$$C_t^{eng} = \frac{\left(S_t^{\mathrm{enf}} - S_t^{\mathrm{tgt}}\right) - \sum_{i=0}^{t-1} B_i^{\mathrm{eng}}}{S_{t-1}^{\mathrm{tgt}} * \rho} \quad \text{...} \ (17)$$

Since resource cost depends on the social cost in the preceding years, the model must be run recursively, starting with year 1 and finishing in year *t*.

Conclusions

This working paper describes procedures for predicting and costing road safety outcomes. Its main aim is to serve as a record of how we set and costed the targets in New Zealand's proposed *Road Safety Strategy 2010*.

The procedures we describe will also be of interest to road safety agencies in other jurisdictions which need to formulate road safety targets, and where road crash data are available in electronic form. These agencies will find the model useful both for target-setting and for allocating resources for greater efficiency.

The model we describe is capable of further refinement. As more data and knowledge become available, they can be readily incorporated into the model so as to increase its accuracy and versatility.

References

LTSA, 1996. *A road safety resource allocation model.* Working paper 1. Land Transport Safety Authority, Wellington, New Zealand.

LTSA, 1998a. *The* Safety Directions *Development Programme.* Working paper 2. Land Transport Safety Authority, Wellington, New Zealand.

LTSA, 1998b. *Setting road safety targets.* Working paper 4. Land Transport Safety Authority, Wellington, New Zealand.

LTSA, 2000. *Estimated effects of interventions on road safety outcomes to 2010.* Working paper 7. Land Transport Safety Authority, Wellington, New Zealand.

Nilssen, G., 1982. *The effect of speed limits on traffic accidents in Sweden.* VTI Report No 68. National Road and Traffic Research Institute, Linköping, Sweden.