# A management tool for allocating road safety resources

Nigel Rockliffe[1], Dimitris Tsolakis[2]
[1] Economic Outcomes Pty Ltd, Melbourne, VIC, Australia
[2] ARRB Group Ltd, Melbourne, VIC, Australia

## 1 Introduction

### 1.1 What it does

This paper describes a management tool for identifying the optimal mix of road safety interventions to achieve a given target at minimum cost, or alternatively to maximise safety for a given budget. The tool embodies a computable simulation model to predict the road safety outcomes of a user-specified set of interventions. Though not an optimising model (it does not formally identify the optimal mix of interventions), it can be used in an exploratory mode to seek the optimum.

It differs from most other predictive models (see *Section 2*) in two main ways. First, it explicitly represents road safety interventions in considerable detail. This means that agencies can use it to explore whether particular interventions are justified, where and to what intensity. Second, it uses standard datasets and software, which makes it readily adoptable by most road safety agencies in developed countries.

### 1.2 Why it is needed

Road safety agencies need to predict road safety outcomes for several reasons. One is efficiency: unless the agency can predict what effect its policies will have, it cannot allocate resources sensibly. Another is target-setting: if a road safety target is to be not just achievable but demonstrably so, it must be transparently linked to the measures intended to bring it about. Lastly, forecasting: road agencies need to predict safety outcomes so that resources can be mobilised and consequences prepared for.

### 1.3 How it came about

This paper and the model it describes arose out of work done for Queensland Transport (QT) to assist them with the preparation of their forthcoming road safety plans. In the words of the Brief, QT sought our advice to help them spot future 'issues' and to prioritise 'needs' (Tsolakis, Rockliffe and Cairney 2005).

There were, we considered, two parts to our task. One was build a management tool that gave advance warning of what the future held. The other was both to flag problems (the 'issues' of the Brief) and to rank them and their possible remedies (the 'needs'). In other words, QT needed more than a tool to tell the future; *they needed a tool that would tell them what to do about it*.

The predictive model was the easy part. But to identify issues and needs we also needed a measure of value. We took the view that QT's aim was to manage road safety so as to maximise the benefit to the people of Queensland. That so, our measure of value was social cost and benefit: the model had simultaneously to rate the social benefit of a road safety outcome and the social cost of achieving it.

To do so, a model must satisfy two additional requirements. First, it must formally link resource inputs to safety outcomes; that is, it must show what happens to the number and severity of crashes whenever a change is made to any given road safety intervention. Second, it must express resource inputs and safety outcomes in monetary terms; that is, it must tell us what we are paying for safety and what that safety is worth. The model described in this paper aims to do both.

## 2 The research context

### 2.1 Macroscopic, microscopic or mesoscopic?

The literature documents dozens of predictive models of road safety. With so many to choose from, do we need another one? We argue below that existing models are either too macroscopic to be useful for resource allocation, or too microscopic to be useful for target-setting. Our proposed model combines the advantages of both.

Smeed (1949) was perhaps the first to analyse road safety at the macroscopic level. He exemplified all that was good and bad about the approach: good, because his prediction was very accurate—for a time; bad, because after that it was wildly wrong. And for a good reason: it took no account of road safety interventions. Later macroscopic models, notably Oppe (1989, 1991), Gaudry (1984), Vulcan and Corben (1998), Broughton et al. (2000) Cameron et al (1993), Cameron (2003), Newstead et al (1998), and Bijleveld et al (2002) have all moved beyond Smeed's early work by including safety interventions, or their proxies, as explanatory variables. Such approaches are valuable for forecasting, target-setting, and justifying individual interventions. But they do not assist in resource allocation as they are spatially aggregated and do not embody the full suite of interventions that any road safety agency would normally command.

Microscopic road safety models, many proprietary, are equally numerous and widely used by consultants and others in the design of individual road safety treatments and the cost-benefit analysis of individual interventions. Some are econometrically estimated, such as Nilssen (1982); others empirically calibrated or derived from engineering theory. Again, these approaches are valuable, but are unsuited for making system-wide predictions.

The proposed model incorporates microscopic algorithms within a macroscopic framework. We know of no other 'mesoscopic' model that fulfills this role.

### 2.2 Strategic plans

How then do road agencies currently tackle their strategy plans? For without a predictive model to underpin it, a strategy plan is hampered by the difficulty of demonstrating that targets are achievable and the resource mix optimised.

Most rely on macroscopic models, for instance, the Australian plan (Australian Transport Council (2000)) and the UK plan (documented in Allsop (1998)). Only one to our knowledge, New Zealand's plan, uses something like the approach proposed here (documented in LTSA 1996, 2000). Macroscopic models of the kind described are in principle capable of producing predictions every bit as good as the ones proposed in this paper. This is good for setting 'headline' national targets. What they cannot do is show how resources should be optimised, either spatially or between interventions.

### 2.3 An economic approach

Our approach to resource allocation is explicitly founded in the theory of welfare economics. But this is not the only approach to road safety management. The Dutch 'Sustainable road safety' and Sweden's 'Vision Zero' (Fildes and Langford 2002) are both founded on the behavioural proposition that human beings are fallible, and whatever we do to make them more alert, law-abiding, or competent, they will still make mistakes. Hence we must design and operate the network to accommodate human error, which we should do by rendering error impossible or its consequences acceptable.

We see no incompatibility in practice. The Dutch and Swedish approaches are engineering design philosophies that recognise, among other things, that one can never rely on behaviour modification alone to achieve complete safety. Nevertheless, as long as resources are limited, one will always have to choose between alternatives. The proposed management tool will assist planners to do so in the most advantageous way.

## 3 The theory

The theory underlying the model is described below, illustrated in Figure 1, and detailed in mathematical terms in the *Mathematical Appendix*.

**Figure 1: Logical flowchart of the model**

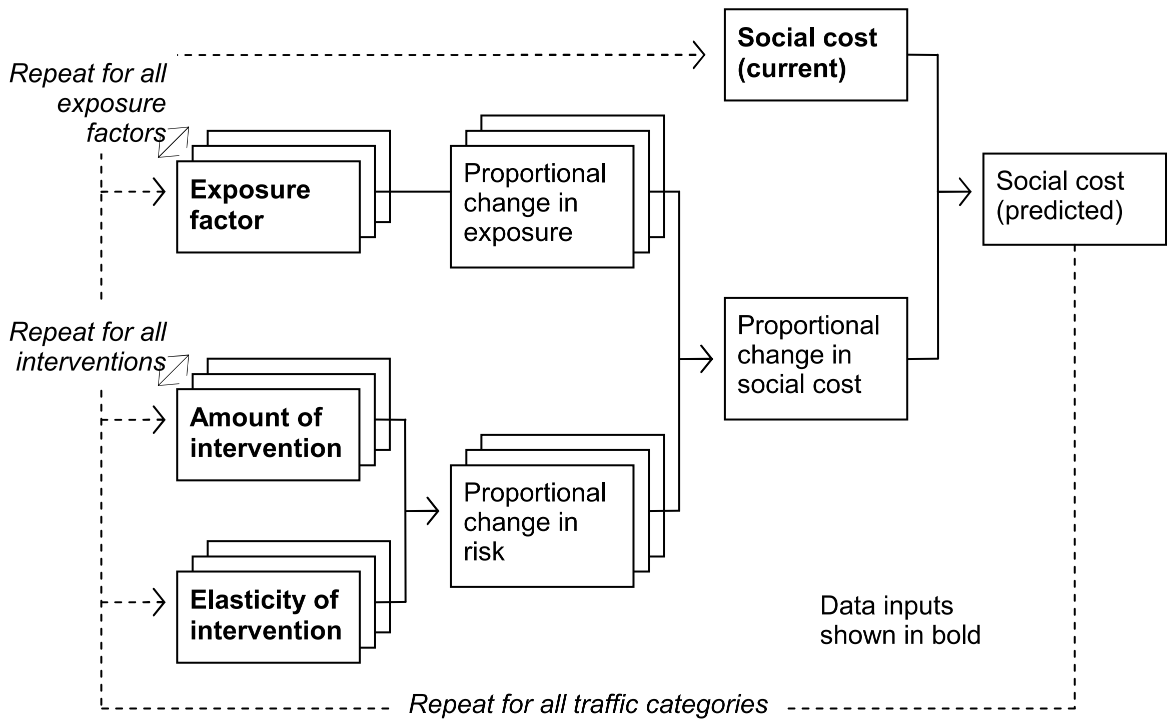


### 3.1 Current social cost

The social cost of crashes over a specified period (the 'current' period) is assigned to traffic categories. This shows where and to whom to direct the appropriate interventions, and partly determines their impact.

In the current model, traffic categories are effectively the same as road links; all traffic on a given link is treated as being in a single category. But a more detailed breakdown is possible, say, by type of controller or vehicle. This would allow a more precise targeting of interventions, but since the data are hard to obtain, this level of detail was not employed in the current study.

## 3.2 Exposure

All traffic in a given category is exposed to the particular level of risk associated with that category. Hence changes in traffic affect exposure, which in turn affects social cost.

In the current model, exposure is taken to be directly proportional to traffic volume, which in turn is directly proportional to factors such as population and to mobility. But more detailed and complex formulations are possible. For instance there is no need to assume that exposure varies in direct proportion to traffic; for as traffic grows, congestion (to the extent that there is any) will tend to reduce speeds and hence the severity of crashes; at the same time it will tend to increase their frequency. In the current model we assumed the two effects would roughly cancel out; in other circumstances they might not.

## 3.3 Risk

Road safety resources are applied to traffic categories in amounts specified by the analyst. They show where particular interventions take effect, who they will affect, and how much. The risk borne by each traffic category varies according to amount of road safety resources applied—the more resources, the greater the reduction.

In the current model, the link between the amount of resource and the level of risk—or 'dose-response' relationship—is mediated by an elasticity parameter.[1] We chose to specify the relationship in this way because elasticities are widely used by economists and engineers, and are computationally simple. But other formulations are possible and, like elasticities, may be estimated econometrically from empirical data or derived from engineering principles.

## 3.4 Predicted social cost

Social cost is predicted by multiplying current social cost by the proportional change in risk and exposure. Double the exposure (that is, double the volume of traffic in the current model) and you double the social cost; likewise double the risk, and you also double the social cost; double both and you quadruple social cost.

Since social cost is calculated separately for each traffic category, the results of the model can be disaggregated and reported by any of the attributes used to characterise traffic categories. Given that traffic categories always relate to road links, it is possible to report results by regions and police districts, which is helpful for setting and monitoring agencies' performance targets. If traffic categories also distinguish vehicle types, for instance, it is possible to set and monitor still more detailed road safety targets.

Naturally the level of disaggregation is limited by standard error of the prediction. Since crashes are random events, any measure of social cost is in fact a sample whose size depends on the amount of traffic to which the social cost relates. Our experience shows that it is rarely possible to obtain adequate estimates of social cost for spatial units smaller than municipalities, and even then three years' data are generally required.

## 3.5 Interactions between interventions

Interventions are never applied in isolation. A road safety plan typically specifies a suite of 20 or more complementary interventions applied with differing intensities in different places. Some relate to vehicles, some to road users, and some to infrastructure. Whatever the inter-

---

[1] An elasticity of minus $\eta$ indicates that a 1 percent increase in resource allocation will reduce risk by $\eta$ percent.

vention, it cannot save life or limb that have already been saved by another intervention; and the more lives that are saved, the harder it is to save those remaining.

The logic of the model accounts for these interactions. Interventions are made to combine multiplicatively where they affect the same traffic category, and additively where they do not.[2] Crash migration is one kind of interaction that is not handled by the model; however, this is unlikely to be a problem at the spatial granularity for which the model is designed.

Synergies between interventions, for instance that between education and enforcement, are likewise not endogenously handled by the model. However, these can be handled exogenously by bundling, so enforcement and education are treated as a single, composite intervention.

# 4 The practice

This section describes the model that was constructed for QT based on the theory described in the preceding section.

## 4.1 Application software and database structure

The modelling approach lends itself to a relational database such as SAS, Oracle or (for smaller models) MS Access. Table 1 lists the data tables and fields.

This database structure relates to the current model. A more detailed model might categorise traffic not just by road link but by other attributes as well, such as type of controller or vehicle. If so, the database structure would contain a data table of traffic categories instead of road links; it would be similar, except that the record for each link would be replicated for each type of traffic it carried.

Likewise, a more detailed model might employ more descriptors for both crashes and road links. There is in fact no theoretical limit to the number and type of descriptors incorporated into the model; and the only limit in practice is one of data availability.

## 4.2 Sources of input data

### 4.2.1 Road links

Roading data came from ARMIS, Queensland's official electronic database of the state-controlled road network (SCRN). Each link was defined so as to be as far as possible homogeneous; that is, constructed to roughly the same engineering standard throughout, located entirely within the same geographical unit, and with about same volume and mix of traffic along its length. The current model contains records of over 5000 road links.

### 4.2.2 Crashes and casualties

Crash and casualty data were extracted from RoadCrash, Queensland's official electronic database of road crash data. The current model contains records of over 50 000 crashes.

---

[2] Suppose two interventions A and B affect the same kinds of crashes, and each is capable of reducing the number of crashes by, say, 40% in isolation. Together, they reduce crashes not by 80% (= 0.4 + 0.4), but by 64% (=1–(1 – 0.4)(1–0.4)). The former is an additive interaction; the latter, a multiplicative one.

4.2.3 Elasticities

Is non-technical language, the elasticity of a given intervention is a measure of the responsiveness of risk to increases in the resources devoted to that intervention: it reflects the percentage change in social cost (or other outcome) produced by a one percent increase in resources. Elasticities are as yet rarely to be found in the literature, so additional work is needed to derive them. But to the extent that they depend on the laws of physics, they are independent of jurisdiction; so once derived they can be widely used by other road safety authorities.

**Table 1: Structure and content of the database**

| Data table | Fields | Comment |
|---|---|---|
| **Road links** *A table in which each record describes a link in the network* | Link identifier* | Required for linking to crash records |
| | Location (LGA*, Police District* etc) | May be geo-coded for flexibility |
| | Length* | Used to calculate traffic volume in veh-km/yr |
| | Traffic flow* | Used to calculate traffic volume in veh-km/yr |
| | Traffic volume* | Calculated from link length and traffic flow |
| | Descriptors (functional level*, seal, median, width, lanes, shoulders etc ) | One field for every descriptor used in the analysis |
| **Crashes** *A table in which each record describes a crash on the network* | Crash identifier* | Required for linking to casualty records |
| | Link identifier* | Required for linking to road link records |
| | Involvement: alcohol, excessive speed, young/old controller, heavy vehicle, pedestrian, bicycle, motorcycle etc* | Optional depending nature of analysis |
| | Descriptors (time of day, weather, lighting etc) | One field for every descriptor used in the analysis |
| **Casualties** *A table in which each record describes a casualty on the network* | Crash identifier* | Required for linking to crash records |
| | Severity* | Used to calculate social cost |
| | Social cost* | Calculated from crash severity and unit social cost |
| **Elasticities** *A table of one record listing elasticities* | Elasticity of intervention 1* | Used to calculate change in risk due to intervention 1, 2, 3 etc |
| | Elasticity of intervention 2* | |
| | Elasticity of intervention 3 etc* | |
| **Exposure factors** *One table for each exposure factor* | Proportional increase* | Used to calculate change in exposure |
| | Attributes of road link* | Those attributes relevant to the exposure factor in question |
| **Resource inputs** *One table for each intervention* | Proportional increase* | Used in combination with the appropriate elasticity to calculate change in risk |
| | Attributes of road link* | Those attributes relevant to the intervention in question |

**Note**: Asterisks show fields used in the current model; unmarked fields are optional and can be used as required in more detailed models.

### 4.2.4 Exposure factors

The current model employed two exposure factors, sourced as follows.

- *Population***.** We relied on ABS population projections, broken down by Statistical Division.

- *Mobility***.** We extrapolated from historical data to arrive at a statewide compound growth rate of 0.73 percent per annum.

### 4.2.5 Resource inputs

Unlike other data inputs, which are state variables or parameters outside government control, resource inputs are policy variables directly under government control. They are therefore not sourced but stipulated by the user of the model.

## 4.3 Outputs

### 4.3.1 Social cost and other measures of road safety outcome

The model computes social cost, crashes and casualties disaggregated by road link. These results can then be tabulated and cross-tabulated by any of the descriptors that characterise road links, for instance, by functional class of road or by geographical unit such as Police District or LGA.

In the current model, social cost is computed for three analysis cases:

- *'Before plan'*

- *'Business as usual'*—the outcome that would pertain if traffic grew as predicted but no changes were made in the mix or intensity of interventions

- *'After plan'*—includes the impact of the plan's proposed changes in resources.

The model is also capable of computing road safety outcomes for each intervention in isolation.

### 4.3.2 Risk and exposure

Besides computing safety outcomes, the model computes risk (expressed as social cost per vehicle-km) and exposure (expressed as vehicle-km). Risk might be relevant in its own right if, say, it were deemed that certain levels of risk were unacceptable on equity grounds. That, however, was not part of the current study.

### 4.3.3 Benefit–cost ratio (BCR)

Currently, the model expresses costs as proportional changes in physical resources. For instance, the model can show the saving in social cost (in other words, the benefit) resulting from, say, a 10% increase in the number of speed cameras. This is not a BCR in the formal sense. To calculate a formal BCR—one useful for ranking projects—cost must be expressed in monetary terms. This requires two additional data items to be incorporated into the model:

- the long-run marginal cost of each intervention (including the annualised capital cost of infrastructure improvements)

- the physical quantities currently being devoted to road safety.

The logic of the model can readily incorporate these data, and the logic for the calculations of BCRs, when the necessary data are available.

# 5 Illustrative hypothetical results

In this section we used actual Queensland road and crash data to generate illustrative results for a hypothetical plan to increase the resources devoted to the following nine interventions by 20%.:

| | |
|---|---|
| Publicity | Anti-skid surface |
| Roadside protection | Blackspot treatments |
| Police patrol | Random breath testing |
| Speed enforcement | Speed camera |
| Trauma management | |

Not all active interventions need be included in the model: exclusion implies only that an intervention continues to be implemented at the same level as before the plan, not that it is discontinued.

## 5.1 Statewide social cost disaggregated by case

Three cases were modelled (Figure 2). Before the plan, social cost is $2.05 billion. After the plan, social cost is predicted to be $1.83 billion. 'Business as usual' shows what would happen if traffic grows as predicted but road safety interventions remain unchanged at pre-plan levels: social cost is predicted to be $2.44 billion. This implies that social cost would increase by $390 million were it not for the impact of the plan. Instead, it is expected to fall by $220 million as the plan will reduce it by $610 million below what it would otherwise be.

**Figure 2: Actual and predicted social cost, by case**

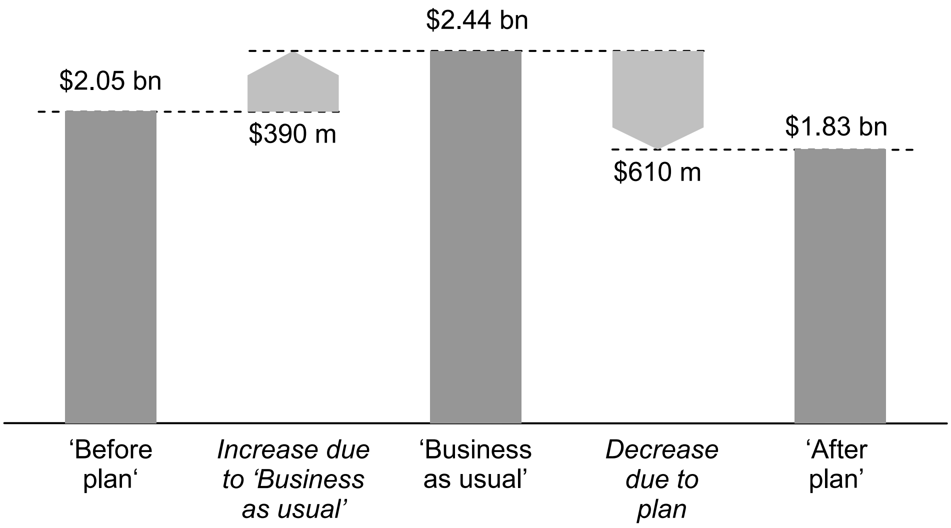

## 5.2 Risk disaggregated by functional class of road

Although risk declines for all functional classes of road as a result of the plan, the proportional change varies between classes (Figure 3) because interventions affect some classes of road more than others. As expected, the data show that road classes vary widely in risk. For instance, Class 1, which includes freeways, is a third less risky than Class 3, which is mainly undivided rural roads.

**Figure 3: Actual and predicted risk, by functional class of road**

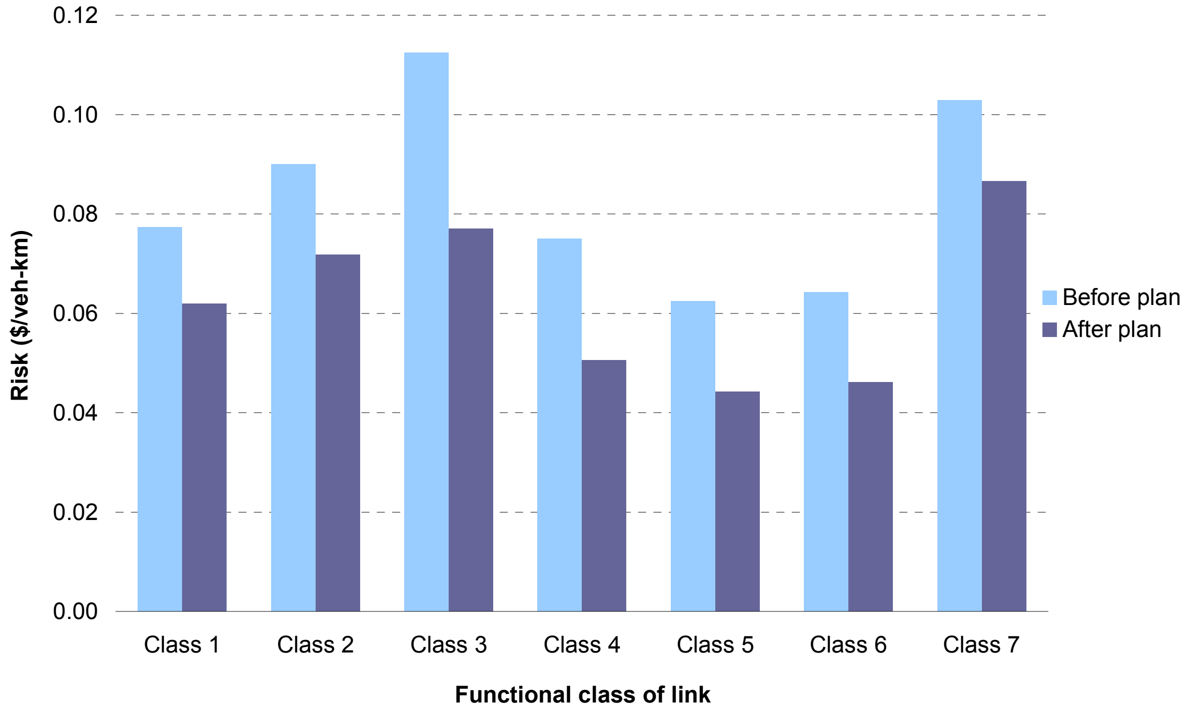


## 5.3 Social cost disaggregated by Police District

Because Police Districts differ in their mix of road classes, they are hard to compare. Districts with a lot of roads in a risky class would have an unduly high social cost, masking its true performance. To compare them on an equal footing their social cost was compared with a 'normalised' social cost, and the variance[3] calculated. A positive variance means the District's roads are more risky than average *after allowing for their functional composition*; a negative variance means the opposite.

Figure 4 presents variances in social cost for all Police Districts. In the current analysis, District variances change somewhat as a result of the plan, but there is no clear tendency for variances to decline. For that to happen, the plan would have to be targeted specifically at the 'worst-performing' Districts—those with the highest positive variances. This could be done by allocating resources spatially so that they were concentrated in the target Districts, or by directing resources at road types and/or crash types that are to be found disproportionately in the target Districts.

---

[3] Variance' in this context is the ratio of (1) the social cost calculated by the model, and (2) the normalised social cost. The normalised social cost is that which would pertain if the average risk for all roads of a given class in the District in question was the same as the average for all roads of the same class in Queensland as a whole.

**Figure 4: Actual and predicted variance in social cost, by Police District**

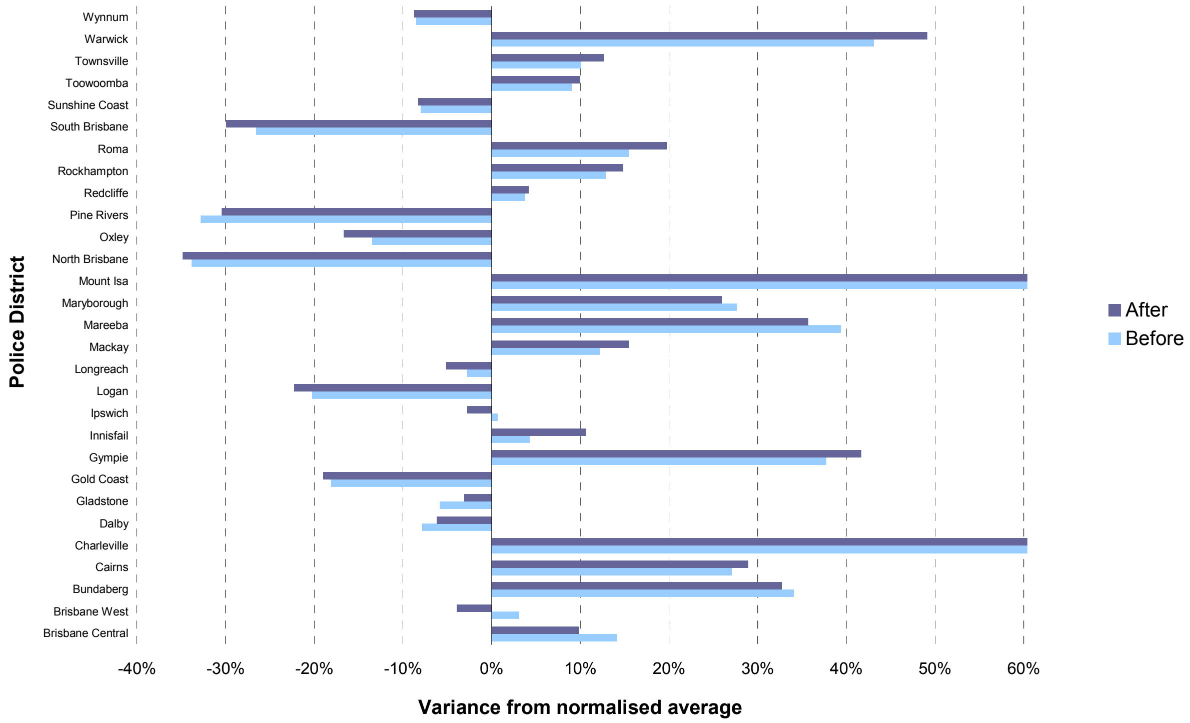


## 6 Conclusions

### 6.1 What we have achieved and what we have not

We have built a management tool that helps road safety decision-makers get the most from the resources at their disposal, and at the same time set demonstrably achievable (hence credible) targets and plan for them.

We have *not* quantified the 'dose-response' relationships between interventions and road safety outcomes. This task remains as hard as ever, and in some cases (for instance, education and advertising) largely intractable. But to expect the current work to do this is to misunderstand its purpose. The proposed management tool is not so much a model as a framework that embodies dose-response relationships established by others. Its value is that it marshals this corpus of knowledge into a coherent, consistent and comprehensive whole.

This has obvious advantages. Once the tool is in place, prediction will become straightforward, even routine. The tool itself will serve as a repository of knowledge. In this way it will clarify debate and so discipline the agenda. Disagreement will continue of course, but it will now be easier to focus on what the disagreement is about.

The tool will engender transparency. In a sometimes intensely political environment, it will make it easier to resist the pressure to set targets that are impossible to reach or budgets impossible to keep. It will also reveal the true cost of political imperatives: the government will still have the last word, but we will know at what cost.

### 6.2 Some possible criticisms and how we respond to them

A critic might argue that the task we have set ourselves is just too hard: one cannot quantify all the relationships needed to make the tool work.

While it is true that the tool is only as good as its component parts, which are often deficient, this is the wrong test. The correct test is: does the tool help? We believe it does, and for the reasons stated above. Besides, if our critics are right and road safety relationships are too complex to be quantified, then much of the road safety literature is without merit. We do not think so.

A critic might furthermore argue that road safety relationships are too complex to be captured in a simple parameter such as an elasticity. Again, this misses the mark. The tool can in principle accommodate any specification of its dose-response relationships; in the absence of better information we merely chose one that was widely used and understood. In any case, the fact that a specification is mathematically simple does not prevent its representing a complex relationship over a limited range.[4]

That said, there are policy areas that are notoriously hard to evaluate, none more so than education, advertising and the 'safety culture'. But this is not unique to road safety. It is beyond the scope of this paper to show how this problem can be tackled, other than to say that it can be modelled, if imperfectly, by such means as a secular trend to represent a gradual cultural change that affects all crashes everywhere.

### 6.3 Where to from here?

The tool could be extended in two main ways. First, our knowledge of elasticities and dose/response relationships in general could be greatly improved by adapting and augmenting existing knowledge to put it in an acceptable format. Second, as currently formulated the tool is a simulation model; optimisation must done by trial and error. Optimality conditions could be derived mathematically, and given sufficient data, could be used to find the optimal mix of interventions analytically.

[4] Economists have successfully done so for years. A road safety dose-response relationship is nothing more than a production function; and the specification adopted in the current model is functionally of a Cobb-Douglas kind—perhaps the most widely used specification in econometric models.

## References


Allsop R (2000). *British Road Safety Strategy To 2010: Target Setting, Monitoring And Research.* Paper presented at Conference on Road Safety Research, Policing and Education, Brisbane.

Australian Transport Council (2000). *National Road Safety Strategy 2001-2010.* ATSB, Canberra.

Bijleveld F, Commandeur J, Koopman S & van Montfort K (2002). *Multivariate nonlinear time series modelling of exposure and risk in road safety research.* SWOV Institute for Road Safety Research, Leidschendam, Netherlands and Department of Econometrics, Vrije Universiteit Amsterdam, Netherlands

Broughton, J, Allsop, R, Lynam, D, and McMahon, C (2000), *The numerical context for setting national casualty reduction targets.* Report 382, Transport Research Laboratory, Crowthorne, UK

Cameron M (2003). *Potential benefits and costs of speed changes on rural roads.* (Report CR 216). Australian Transport Safety Bureau.

Cameron, M, Haworth N, Oxley J, Newstead S, & Le T. (1993). *Evaluation of Transport Accident Commission road safety television advertising.* (Vol. Report 52): Monash University Accident Research Centre.

Fildes, B and Langford, J (2002). *Vision Zero and Sustainable Road Safety.* Monash University Accident Research Centre. In Austroads Australasian Road Safety Handbook: Volume 1. Report No AP-R234/03. Sydney.

Gaudry M (1984). *DRAG, un modèle de la Demande Routière, des Accidents et de leur Gravité, appliqué au Québec de 1956 à 1982*. Technical Report Publication CRT-359, Centre de recherche sur les Transports, et Cahier #8432, Département de sciences économiques, Université de Montréal.

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

LTSA (2000). *Predicting and costing road safety outcomes.* Working paper 6. Land Transport Safety Authority, Wellington, New Zealand.

Newstead S, Cameron M, & Narayan S (1998). *Further modelling of some major factors influencing road trauma trends in Victoria: 1990-96.* (Vol. Report No. 129): Monash University Accident Research Centre.

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.

Oppe S (1989). *Macroscopic models for traffic and traffic safety.* Accident Analysis and Prevention 21, 225–232.

Oppe S (1991). *Development of traffic and traffic safety: Global trends and incidental fluctuations.* Accident Analysis and Prevention 23 (5), 413–422.

Smeed R (1949). *Some statistical aspects of road safety research.* Journal of the Royal Statistical Society A 112 (1), 1–34.

Tsolakis, D, Rockliffe, N, and Cairney, P (2005). *Supply of road safety data analysis and issue identification framework.* ARRB Contract Report VC71002, ARRB Group Ltd, Melbourne.

Vulcan, P and Corben, B (1998), *Prediction of Australian road fatalities for the year 2010*, National Road Safety Summit, 1998, Canberra, Federal Office of Road Safety, Canberra.

# Mathematical appendix

## Specification

### Variables

The model links three variables that describe traffic characteristics. Exposure to risk, $E$ , is determined by the quantity of traffic, and is typically measured in units of vehicle-km, though other units are possible, such as person-km. Road trauma, $S$ , is a measure of the amount of trauma inflicted on a particular category of traffic, and is expressed in terms of crashes or casualties (possibly broken down by type), or social cost. Risk, $R$ , is defined as the amount of road trauma per unit of traffic. By definition, these variables are linked as follows:

$$S = R.E \tag{1}$$

All must be expressed in consistent units. For instance if exposure is expressed in vehicle-km and road trauma in dollars of social cost, risk is in terms of dollars per vehicle-km. Alternatively, if exposure is in person-km and road trauma in fatalities, risk is in terms of deaths per (million) person-km.

### Traffic categories

Different types of traffic grow at different rates and respond differently to interventions. For modelling purposes, therefore, all traffic on the network was disaggregated into disjoint categories as follows.

Let every unit of traffic on the network be characterised by n attributes, each of which can have a number of discrete values. Set $I$ of possible traffic categories is defined as the product set $I = I_1 \times I_2 \times ... \times I_n$ formed from the n sets of attribute values $I_1, I_2, ..., I_n$; and set i is defined as a unique set of *n* traffic attributes, where i a member of $I$ ( $i \in I$ ). Thus for every traffic category *i*, we define exposure $E_i$ , risk $R_i$ , and road trauma $S_i$ , where by definition

$$S_i = R_i .E_i \tag{2}$$

### Exposure

Since the factors that affect exposure to risk (that is, traffic volume) within the same traffic category generally combine multiplicatively, predicted exposure $E_i^*$ is given by

$$E_i^* = E_i \prod_k \left(1 + e_{ik}\right) \tag{3}$$

where $e_{ik}$ is the proportional change in exposure caused by factor *k* acting on traffic *i*.

### Risk

Since interventions within the same traffic category combine multiplicatively, predicted risk $R_i'$ is given by

$$R_i^* = R_i \prod_j \left(1 + r_{ij}\right) \tag{4}$$

where $r_{ij}$ is the proportional change in risk brought about by intervention j acting on traffic i.

The proportional change in risk is, however, determined by the proportional change in quantity of resources $q_{ij}$ of intervention *j* applied to traffic *i,* where $\eta_j$ is an elasticity parameter specific to intervention *j*:

$$\eta_j = \left( \frac{r_j}{q_j} \right) \tag{5}$$

Substituting (5) in (4):

$$R_i^* = R_i \prod_j \left( 1 + \eta_j q_{ij} \right) \tag{6}$$

**Road trauma**

From (1), the predicted amount of road trauma $S_i^*$ for traffic *i* is given by:

$$S_i^* = R_i^* . E_i^* \tag{7}$$

Summing across all traffic categories:

$$S^* = \sum_i \left( R_i^* . E_i^* \right) \tag{8}$$

Substituting (3) and (6) in (8):

$$S^* = \sum_i \left( R_i \prod_j \left( 1 + \eta_j q_{ij} \right)_i . E_i \prod_k \left( 1 + e_{ik} \right) \right) \tag{9}$$

Substituting (2) in (9):

$$S^* = \sum_i \left( S_i \prod_j \left( 1 + \eta_j q_{ij} \right)_i . \prod_k \left( 1 + e_{ik} \right) \right) \tag{10}$$