Methods for forecasting the impacts of connected and automated vehicles

Deliverable D3.2 – WP3 – PU

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 824361.

Methods for forecasting the impacts of connected and automated vehicles

Work package 3, Deliverable D3.2

Please refer to this report as follows:
Elvik, R., Meyer, S. F., Hu, B., Ralbovsky, M., Vorwagner, A., Boghani, H. (2020). Methods for forecasting the impacts of connected and automated vehicles, Deliverable D3.2 of the H2020 project LEVITATE.

Project details:	
Project start date: **Duration:** **Project name:**	01/12/2018 36 months LEVITATE – Societal Level Impacts of Connected and Automated Vehicles
Coordinator:	Pete Thomas, Prof – Prof. of Road & Vehicle Safety Loughborough University Ashby Road, LE11 3TU Loughborough, United Kingdom
	This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 824361.

Deliverable details:	
Version: **Dissemination level:** **Due date:** **Submission date:**	Final PU (Public) 29/02/2020 25/02/2020

Lead contractor for this deliverable:
Rune Elvik – Institute of Transport Economics

Report Authors:	Elvik, R., Meyer, S. F. (TOI, Norway), Hu, B., Ralboysky, M., Vorwagner, A. (AIT, Austria), Boghani, H. (Loughborough, United Kingdom)

Revision history

Date	Version	Reviewer	Description
03/02/2020	Preliminary draft 1	Joost de Winter	Accepted with reservation
03/02/2020	Preliminary draft 1	Bin Hu	Accepted with reservation
25/02/2020	Final draft	WP and task leader: Rune Elvik Coordinator: Pete Thomas	
25/02/2020	Final deliverable	Pete Thomas – Loughborough University ➔ EC	

Legal Disclaimer

Table of contents

List of figures

List of tables

Executive summary

This deliverable of the Horizon 2020 project Levitate presents methods for forecasting the impacts of connected and automated vehicles. The main objectives of the deliverable are:

1. To survey methods that can be applied to forecast the impacts of connected and automated vehicles and indicate which methods are best suited for specific impacts.
2. To apply the methods to develop forecasts of impacts and assess the uncertainty of the forecasts.
3. To quantify as many potential impacts as possible for inclusion in the policy support tool being developed in Levitate.
4. To identify impacts for which forecasts are currently difficult to make.

As a starting point, a list has been developed of potential impacts of connected and automated vehicles that have been identified in the literature. This list is used as a frame of reference for discussing methods that can be applied to forecast impacts. The ultimate objective is to develop a policy support tool to enable policy makers, for example a city government, to estimate the impacts of different policies influencing the introduction and use of connected and automated vehicles. To serve this end, it is desirable to quantify as many of the potential impacts as possible. Quantification is necessary in order to convert impacts to monetary terms and enable analyses of their costs and benefits. The conversion of impacts to monetary terms is decribed in deliverable D3.3 of Levitate. Cost-benefit analyses of connected and automated vehicles are described in deliverable D3.4 of Levitate.

Based on deliverable D3.1 of Levitate, a distinction is made between direct, systemic and wider impacts. Direct impacts are noticed by each road user on each trip. Systemic impacts are system-wide impacts occurring in the transport system, such as changes in traffic volume, travel time or route choice. Wider impacts are impacts that are noticed or occur outside the transport system. Accidents have been classified as a wider impact, since some of their impacts occur outside the transport system (hospitals, insurance companies, etc.)

A survey of methods that can be applied to study potential impacts of connected and automated vehicles found that several methods are relevant. A distinction is made between retrospective methods and future-oriented methods. A combination of methods in both categories is needed to gain a full overview of potential impacts of connected and automated vehicles. Among the future-oriented methods, various forms of traffic simulation have been widely applied to study impacts of connected and automated vehicles. Traffic simulation can indicate changes in road capacity or travel time directly. Changes in road safety are usually modelled in terms of surrogate measures, in particular traffic conflicts.

Estimates of potential impacts have been developed by combing the following methods:

1. Systematic literature reviews and (simplified) meta-analysis of studies
2. Surveys of willingness-to-pay for automation technology and of the valuation of travel time
3. System dynamics modelling, in particular with respect to the wider impacts of automation technology, illustrated for the case of moving house within parts of a city
4. Development of scenarios describing potential cyber attacks on connected and automated vehicles

The results of studies employing these methods are presented in this deliverable. It was not always possible to formally synthesise the results of several studies dealing with an impact by means of state-of-the-art techniques of meta-analysis. This is why the term simplified has been inserted above. The main findings can be summarised as follows:

1. Several studies indicate that connected and automated vehicles will increase road capacity, or utilise it more effectively.
2. An increase in road capacity or more effective utilisation of it will, all else equal, reduce travel time.
3. Connected and automated vehicles will reduce the number of road accidents.
4. Connected and automated vehicles will reduce emissions and air pollution.

There is consensus in the research literature about these impacts. Quantitative estimates of the impacts vary considerably. Quantitative estimates of impacts are therefore associated with considerable uncertainty. Traffic simulation studies enable a number of potential impacts of connected and automated vehicles to be quantified in terms of dose-response curves. A dose-response curve has the market penetration rate of automated vehicles as the dose, and the size of an impact, e.g. the percentage change in travel time, as the response. The dose-response framework is flexible and widely applicable. However, dose-response curves are still not available for all impacts of connected and automated vehicles. Moreover, some the currently available dose-response curves are highly uncertain. There is thus a need for developing more valid dose-response curves than those that are available now. Dose-response curves and estimates of the uncertainty of the curves have been developed for:

1. Impacts on capacity and mobility
    a. Lane capacity
    b. Junction capacity
    c. Delays on motorways
    d. Delays in roundabouts
    e. Delays in signalised junctions
    f. Travel time on motorways
    g. Delays in junctions (average of roundabouts and signalised)
    h. Travel time in cities (derived from changes in capacity and delays)
2. Impacts on safety
    a. Rear-end and lane change collisions (motorways)
    b. Accidents in signalised junctions
    c. Accidents in roundabouts
    d. Accidents in priority controlled junctions
    e. Cyclist and pedestrian accidents
    f. Accidents in urban junctions (average of signals, roundabouts, priority)

3. Impacts on fuel consumption (and emissions)
    a. Fuel consumption (and emissions, if these are assumed to be proportional to fuel consumption)
4. Other impacts and feedback impacts
    a. Diffusion curves for market penetration of automated vehicles
    b. (Induced) travel demand on motorways
    c. (Induced) travel demand in cities
    d. Behavioural adaptation to platoons by manually driven vehicles

In addition to the dose-response functions, the following impacts were quantified:

1. Willingness-to-pay for automation technology
2. Demand function for automated vehicles
3. Valuation of travel time in automated vehicles

A system dynamics model was illustrated for land use and moves between parts of a city, inducing longer commutes.

Some potential impacts of connected and automated vehicles are difficult to predict with any confidence. This applies to the following impacts:

1. Whether there will be a widespread transition from individual to shared mobility. There is no consensus on whether individual use of motor vehicles will continue at present levels or be replaced by various forms of shared mobility.
2. It is not clear what type of propulsion energy connected and automated vehicles will use. Some researchers expect the introduction of connected and automated vehicles to be associated with a transition to electric propulsion.
3. Connected and automated vehicles are vulnerable to cyber attacks. However, the risk of such attacks cannot be quantified. Potential scenarios can be described.
4. The costs of connected and automated vehicles are highly uncertain. It is not clear that connected and automated vehicles will be as affordable as current motor vehicles. The costs of automation technology may influence the level of inequality in access to transport.
5. Behavioural adaptation to connected and automated vehicles, in particular during the transition period before full market penetration. While some studies suggest various forms of behavioural adaptation, predicting its form and impacts is impossible or speculative.
6. Changes in employment are difficult to predict. While full automation will eliminate the need for drivers, other potential changes in employment are less known.

Evidence based on research is therefore not available for all potential impacts of connected and automated vehicles identified in Levitate. Research can therefore only inform policy makers about some potential impacts of connected and automated vehicles.

1 Introduction and objective

This deliverable gives an overview of methods for forecasting potential impacts of connected and automated vehicles. The importance of quantifying impacts is stressed. A main objective of Levitate is to develop a policy support tool. To develop this tool, it is necessary to establish quantified relationships between levels of implementation for connected and automated vehicles and the magnitude of impacts of such vehicles. This introductory chapter defines the objectives of the deliverable.

1.1 The Levitate project

LEVITATE (Societal level impacts of connected and automated vehicles) is a Horizon 2020 project which has the following main objectives:

1. To develop a range of **forecasting and backcasting scenarios** and baseline conditions relating to the deployment of one or more mobility technologies that will be used as the basis of impact assessments and forecasts. These will cover three primary use cases – automated urban shuttle, passenger cars and freight services.

2. To establish **a multi-disciplinary methodology** to assess the short, medium and long-term impacts of connected and automated transport systems (CATS) on mobility, safety, environment, society and other impact areas. Several quantitative indicators will be identified for each impact type.

3. To apply the methods and **forecast the impact of CATS** over the short, medium and long term for a range of use cases, operational design domains and environments and an **extensive range of mobility, environmental, safety, economic and societal indicators**. A series of case studies will be conducted to validate the methodologies and to demonstrate the system.

4. To incorporate the methods within a **new web-based policy support tool** to enable city and other authorities to forecast impacts of CATS on urban areas. The methods developed within Levitate will be available within a toolbox allowing the impact of measures to be assessed individually. A Decision Support System will enable users to apply backcasting methods to identify the sequences of CATS measures that will result in their desired policy objectives.

1.2 Work package 3 and objectives of this deliverable

This deliverable contributes to the second and third objectives. Developing methods for assessing and predicting the impacts of CATS consists of the following main stages:

1. Identification and classification of impacts of connected and automated transport systems

2. Description and measurement of impacts of connected and automated transport systems
3. Development of methods for backcasting and forecasting of impacts of connected and automated transport systems
4. Evaluation of comparability and amenability to monetary valuation of impacts of connected and automated transport systems
5. Developing methods for analysing costs and benefits of connected and automated transport systems

Deliverable D3.1 addressed the first point on this list and provided a taxonomy of potential impacts of connected and automated transport systems (CATS) and briefly discussed how to measure the impacts. This deliverable continues where deliverable D3.1 ended, by discussing methods for forecasting the impacts of CATS. The main objectives of the deliverable are:

1. To give an overview of methods that have been used or can be used in order to forecast impacts of connected and automated vehicles
2. To discuss criteria for the choice between methods or how different methods can best be combined
3. To develop quantified relationships between levels of implementation of connected and automated vehicles and as many potential impacts of such vehicles as possible based on available studies
4. To show how the quantified relationships can be applied to predict impacts of connected and automated vehicles.

The third objective is particularly relevant for the development of the Levitate policy support tool. This will contain an estimator module, allowing the user to estimate impacts of CATS in specific contexts, by specifying the technologies whose impacts are to be estimated and the type of traffic environment or type of transport where CATS will be introduced. It is necessary to specify impacts as functional relationships in order to perform cost-benefit analyses, which will also be an option in the policy support tool. The main research questions to be answered in this deliverable are:

1. Which methods can be used to forecast potential impacts of connected and automated vehicles?
2. What results have been obtained by means of different methods? Do different methods produce similar results?
3. How can functional relationships for the impacts of connected and automated vehicles be distilled from available studies?

The answers to all these questions are based on a literature review. To give a preview of the results, various forms of simulation have been widely used to model the impacts of CATS. Simulation will also be used in Levitate and will produce additional estimates of impacts not included in this deliverable. The policy support tool will include both the impacts based on the literature reviewed in this deliverable and further estimates of impacts based on simulation studies performed in other work packages of Levitate.

2 Overview of potential impacts

Potential impacts of connected and automated transport systems (CATS) were identified in Deliverable 3.1 of Levitate. For the sake of completeness, these impacts are also listed here.

Table 1 lists the potential impacts of CATS as identified in deliverable D3.1. A distinction is made between direct, systemic and wider impacts.

Table 1: Potential impacts of connected and automated transport systems

Impact	Description of impact
	Direct impacts (micro level)
Travel time	Duration of a trip between a given origin and a given destination
Travel comfort	Subjective rating of the level of comfort on a given trip
Valuation of time	Willingness-to-pay for reduced travel time
Vehicle operating cost	Direct outlays for operating a vehicle per kilometre of travel
Vehicle ownership cost	The cost of buying and keeping a vehicle
Access to travel	The opportunity of taking a trip whenever and wherever wanted
Route choice (individual)	Technology to support the best choice of route on a given trip
	Systemic impacts (macro level)
Amount of travel	Vehicle kilometres or person kilometres of travel per year in an area
Road capacity	The maximum number of vehicles that can pass a section of road per unit of time
Congestion	Delays to traffic as a result of high traffic volume
Infrastructure wear	The rate per unit of time at which a road is worn down
Infrastructure design	Equipping roads with technology for vehicle-to-infrastructure communication
Modal split of travel	The distribution of trips between modes of transport
Optimisation of route choice	Direction of vehicles to routes that minimise overall generalised cost of travel for traffic as a total
Vehicle ownership rate	Percent of households owning 0, 1, 2 etc vehicles
Shared mobility	Sharing a vehicle with others on a trip-by-trip basis
Vehicle utilisation rate	Share of time a vehicle is in motion (not parked); cabin factor (share of seats in use)
Parking space	Size of areas designated for parking
Traffic data generation	The availability of detailed trip data for transport planning

Table 1: Potential impacts of connected and automated transport systems

Impact	Description of impact
	Wider impacts
Trust in technology	Share of population indicating high trust in automation technology
Road safety	The number and severity of accidents
Propulsion energy	Source of energy used to move vehicles (fossil fuel or electric)
Energy efficiency	Rate at which propulsion energy is converted to movement; rate of loss due to conversion of energy to heat or noise rather than movement
Vehicle emissions	Emissions in micrograms per kilometre per vehicle (by chemical)
Air pollution	Concentration of pollutants per cubic metre of air
Noise pollution	Number of individuals exposed to noise above a certain threshold
Public health	Incidence of morbidity and mortality; subjectively rated health state
Employment	Changes in number of people employed in given occupations
Geographic accessibility	Time used to reach a given destination from different origins
Inequality in transport	Statistics indicating skewness in the distribution of travel behaviour between groups according to social status
Commuting distances	Length of trips to and from work
Land use	Density of land use for given purposes (residential, industrial, etc.)
Public finances	Income and expenses of the public sector

A total of 33 impacts are listed. The direct impacts occur at the micro level. The systemic impacts occur at the macro level and are an aggregation of the direct impacts. The direct impacts are impacts that are noticed by each road user on each trip. The systemic impacts are system-wide changes in traffic operations and safety.

Thus, changes in travel time are closely related to changes in road capacity and congestion. Change in travel comfort is related to trust in technology. Valuation of travel time is related to optimisation of route choice. Vehicle operation cost is related to vehicle utilisation rate. Vehicle ownership cost is related to vehicle ownership rate. Access to travel is correlated with amount of travel and inequality in transport. Individual route choice is related to optimisation of route choice.

Most of the studies that try to estimate and quantify impacts of CATS do so at the aggregate level and deal with systemic impacts. However, estimating systemic impacts will capture the typical impacts at the individual level, e.g. typical changes in travel time.

3 Methods for forecasting impacts

This chapter gives an overview of methods for forecasting impacts of connected and automated transport systems (CATS). The main objective is to identify which methods or combinations of methods are best suited for forecasting specific potential impacts. The possibility of combining results from studies employing the same method is discussed. Differences and similarities between methods are discussed. A distinction is made between historical or retrospective methods and future-oriented methods. Combining methods in order to develop quantified estimates of the impacts of CATS is discussed.

A common method for establishing knowledge about the impacts of transport innovations is to study what happened before and after their introduction. Did the new road reduce congestion and travel time? Are cars of model year n+1 safer than cars of model year n? Did the change in underground fares lead to a change in ridership? In all these, and very many other cases, the change made can be clearly specified; it occurred at a certain date or within a comparatively short period of time, and data relevant for evaluating impacts were available or could be collected both before and after the change.

The introduction of connected and automated vehicles does not fit this pattern. It is likely to be a prolonged and gradual process, evolving during a period of perhaps 10-30 years. There is no well-defined before- or after-period, except perhaps the period before any vehicles were connected or automated and the period after all vehicles are connected or automated. When the task is to assess the potential impacts of connected and automated vehicles, one clearly must do so before these vehicles have been widely introduced. The word "potential" is therefore essential. Any current assessment can only deal with potential impacts. Actual impacts can be studied much later.

This does not necessarily mean that traditional before-and-after studies or similar "historical" study designs are irrelevant when the task is to predict future impacts of systems not yet introduced. To give one example: Several studies predict that connected and automated vehicles can increase road capacity, or rather, utilise it more effectively, e.g. by driving with shorter headways, eliminating speed variance, change lanes with shorter gaps and select a speed that maximises throughput. If the increase in road capacity can be quantified, there are many "historical" studies of how increases in road capacity have affected travel demand. Surely, these studies are relevant when trying to predict a rebound effect elicited by the increase in road capacity associated with connected and automated vehicles.

Thus, one cannot rely on one type of study exclusively. The complexity and multidimensional character of the task of forecasting the impacts of connected and automated vehicles requires the use of a broad array of methods. It is therefore useful to list relevant methods and categorise them. This is done in the next section.

3.1 A classification of methods

A distinction is made between two main types of methods:

1. Historical or retrospective methods
2. Future-oriented methods

Within each group, a further classification of methods can be developed. The following list comprises the most commonly applied methods:

1. Historical or retrospective methods
    a. Longitudinal studies; time-series models
    b. Before-and-after studies (several versions exist)
    c. Epidemiological studies; comparative or retrospective risk analyses
    d. In-depth studies of accidents
    e. Meta-analysis
    f. Household travel surveys (to reconstruct actual travel)
    g. Travel demand modelling
    h. Naturalistic driving studies
2. Future-oriented methods
    a. Scenario analyses
    b. Delphi surveys
    c. Biomechanical modelling (of impacts involving future vehicles)
    d. Field operational trials
    e. Driving simulation; driving simulator studies
    f. Traffic simulation; mathematical modelling of traffic
    g. Mesoscopic simulation; activity-based-modelling
    h. Operations research
    i. System dynamics modelling
    j. Reliability engineering; prospective risk analyses
    k. Surveys (can be used for many topics)
    l. Willingness-to-pay studies (often surveys, but listed separately here)

Examples of knowledge that can be gained from the various methods are given below.

3.1.1 Historical or retrospective methods

Longitudinal studies and time series models may, for example, shed light on the time it takes for new technology to fully penetrate the market and factors influencing the length of this time. Motor cars were invented more than 100 years ago and have still not reached saturation levels in all countries of the world. However, in the highly motorised countries, the car fleet is now growing slowly, suggesting that in these societies, cars are close to full market penetration. Models of how long it takes for the car fleet to completely renew can be used to predict the duration of the transition phase from only manual vehicles to fully automated vehicles, assuming that historical turnover rates remain constant. For an example, relating to transition to electric vehicles, see Fridstrøm and Østli (2016)

There are no **before-and-after studies** of connected and automated vehicles. However, numerous before-and-after studies have been conducted to evaluate impacts that are relevant for connected and automated vehicles. There are, for example, studies evaluating impacts of level 1 and level 2 driver assistance and automation systems (see e.g. Høye 2011 on electronic stability control and Leslie et al. 2019 for active safety systems and advanced headlight systems). For some impacts, like new travel demand induced by increased road capacity, it appears reasonable to assume that historically observed relationships will apply to connected and automated vehicles (Hymel 2019).

Estimates of safety impacts of connected and automated vehicles can be obtained by **epidemiological studies** comparing the risks of automated vehicles to manual vehicles and **naturalistic driving studies** quantifiying the contributions made to accidents of driver distractions or hazardous actions that will be eliminated by connected and automated vehicles. Examples of these approaches and what one can get out of them are given later in the report.

Some estimates of the potential safety impacts of connected and automated vehicles (Fagnant and Kockelman 2015, Herrmann, Brenner and Stadler 2018) have relied on findings of **in-depth studies of accidents**. These studies typically conclude that more than 90% of accidents can be attributed to errors made by road users. The quoted studies assumed that these errors will be eliminated by connected and automated vehicles.

Meta-analysis is used to synthesise the findings of several studies dealing with the same topic. There are many techniques of meta-analysis, some of which may be relevant for summarising studies of potential impacts of connected and automated vehicles. There have, for example, been several studies of how truck platooning influences fuel consumption (Slowik and Sharpe 2018). It is of interest to summarise the results of these studies in terms of a mean estimate of effect and the uncertainty surrounding this estimate.

Household travel surveys are carried out regularly in many countries. These studies show current travel as well as, for example, current car ownership, and knowing this is essential when estimating potential changes in travel behaviour that may be produced by connected and automated vehicles. Travel surveys also often collect data on driving licences. One of the potential impacts of fully automated vehicles is to improve mobility for people without a driving licence.

Household travel surveys are often used as input in **travel demand modelling**. Models of travel demand are widely used when planning new transport infrastructure or predicting effects of, for example changes in traffic control or the introduction of road pricing. As connected and automated vehicles are widely believed to influence the generalised costs of travel, they are likely to also influence travel demand. Known elasticities of travel demand with respect to the generalised costs of travel can be applied to predict how changes in the generalised costs of travel brought about by connected and automated vehicles may influence travel demand.

Naturalistic driving studies (Dingus et al. 2016) apply inconspicuous cameras to record driving behaviour, including indicators of driver alertness and where the driver is looking. The cameras are very small and their pres ence in the vehicles is believed not to influence driver behaviour. An advantage of naturalistic driving studies is that all accidents are recorded. Complete recording of accidents eliminates an important source of error when comparing the safety of human driven vehicles and autonomous vehicles.

All accidents involving autonomous vehicles are recorded, whereas, except in naturalistic driving studies, this is not case for accidents involving cars driven by human drivers.

3.1.2 Future-oriented methods

The development and analysis of **scenarios** is a common method for studying possible future development, especially if no definite predictions can be made regarding these developments. Scenario analysis is eminently suited to the study of connected and automated vehicles, as there are too many deep uncertainties about the introduction of these vehicles to predict with any confidence when they will be introduced, how much safer they will be than vehicles driven by humans, whether they will affect car ownership rates, and so on and so forth. A scenario gives the analyst the opportunity to fix values for all these uncertain factors and then see what happens. If a set of scenarios are developed, and these are very different, they may encompass all possible outcomes of the introduction of connected and automated vehicles. Scenarios have been developed for cyber attacks on connected and automated vehicles.

A **Delphi survey** is a multi-stage survey of experts, often asking them about when a certain technology will be introduced and what its effects are likely to be. Answers in the first stage are usually reported back to respondents, giving them an impression of what others have answered and offering the opportunity to revise their answers in the first round. The idea is that after a few rounds, a consensus opinion will emerge. An example of the use of a Delphi survey is given by Høye, Hesjevoll and Vaa (2015).

Biomechanical modelling, applying open source human body models, is increasingly used to predict the effects of design elements or design changes in automobiles. Human body models are mathematical models and simulations of the behaviour of human bodies in accidents. While computationally intensive, such models are increasingly viewed as a better tool for analysing injury mechanisms than using crash test dummies. Mathematical models are more flexible and one may develop models of both adults and children, spanning a wide range of body sizes. The ongoing Horizon 2020 project VIRTUAL is developing a set of open source human body models that will be applied, among other things, to studying new seating postures in driverless cars.

A **field operational trial** (often abbreviated FOT) is a study made in a natural setting, either real traffic or a closed test track, to test the performance of drivers or of new technology. It has, for example, been used to evaluate the reaction times of drivers to take-over requests in SAE level 3 automated vehicles. Field operational tests are often used to test prototypes of new technology, to evaluate if it functions as intended.

Driving simulator studies have been widely applied to evaluate driver use of and adaptation to automation technology. Most advanced driving simulators can simulate fully automated driving. One may then compare automated driving to manual driving. Driving simulators are particularly useful for evaluating behavioural adaptation to automation technology.

Traffic simulation has been widely applied to estimate potential impacts of connected and automated vehicles. Since this method is quite commonly used to study potential impacts of new technology, it will be described in more detail than the other methods surveyed in this section. Traffic simulation is the mathematical modelling of transport systems through the application of computer software to better help plan, design and operate transport systems. It is a flexible tool that has proven to be valuable due to its applicability to ex-ante evaluations of new transport technologies (such as CATS). As it has been widely applied in Levitate, a more detailed description of it is given in section 3.1.3.

Mesoscopic simulation is a supplemental method within the group of simulation approaches, which emphasises the modelling of behaviours and choices of individuals. Such an activity-based-modeling (ABM) framework is realised by the mesoscopic traffic simulation tools of MATSim. "Mesoscopic" in this context underlines the fact that the method is less focussed on immediate interactions between road users, thus reducing the level and complexity of these details, but rather on the choices the simulated agents make to re-arrange their daily routes and schedules of activities. Each of the activities within a complete daily chain or "plan" are preferably reached in time by the means of transport available to each agent within the simulated area under investigation. The major conclusions that can be extracted from such models refer to changes in modal split, as well as differences in road network loads and vehicle utilisation. MATSim has been applied to a wide range of scenarios and locations and provides a rich set of results for comparison and transferability. For a recent application, see Ahanchian et al. (2019).

Operations research offers a complementary methodology toolset to simulation. It contains analytical methods for solving optimisation problems in the field of management of organisations – in this project particularly transport management. Among the methodologies, there are two classes. The first class consists of exact approaches which aim to solve problems to proven optimality – provided that they are given enough run-time and memory. Well known representatives are mixed integer programming or branch-and-bound. The second class are (meta-)heuristics which compute approximate solutions but usually require significantly less runtime. For this project and in practice the latter is more convenient since real-world problems are too complex for exact approaches. Of course there are also so-called hybrid methods that combines these two classes, trying to benefit from advantages of both sides.

System dynamics is a forecasting technique originating in meteorology. It has been adopted by the social sciences and became widely known when the book "The limits to growth" (Meadows et al. 1972) was published in 1972. It models processes of change by means of a set of differential equations. For a recent application to CATS, see Nieuwenhuijsen et al (2018). An illustration of system dynamics modelling is given in Section 3.1.4.

Reliability engineering aims to develop safety barriers in depth, meaning that if one of them fails, the system may continue to operate as normal or be brought to a controlled stop before an accident occurs. A key element in reliability engineering is redundancy. An early example is the use of dual master brake systems on cars. There is one system operating on both front wheels and one rear wheel, and a duplicate system operating on both front wheels and the other rear wheel. If one system fails, the car will still have a working braking system for three of four wheels.

Surveys, including stated preference surveys can be applied as a future-oriented method. There have been several surveys of how people view connected and automated vehicles and how much they are willing to pay for such vehicles. These studies may serve as a basis for assessing potential demand for connected and automated vehicles.

Studies of **willingness-to-pay**, although they are surveys, are listed as a separate method. Willingness-to-pay studies are relevant for assessing the potential impacts of connected and automated vehicles for at least two of these impacts. The first is whether there will be demand for connected and automated vehicles, i.e. are people willing to pay the additional costs that will be associated with connected and automated vehicles. If not, or if only the richest can afford such vehicles, their market penetration may be slow, at least at first. However, the costs of new technology tend to come down over time. The second potential impact of connected and automated vehicles where willingness-to-pay studies are relevant are impacts on travel time. Connected and automated vehicles may change both the length of time taken for a given trip and the quality of that time. Once you no longer have to drive yourself, you can spend time on productive activities or on relaxing, which is widely believed to reduce the value of travel time savings, as the time spent travelling may be felt as less "wasted" than today.

3.1.3 Traffic simulation

Simulating connected and automated vehicles (CAVs) is a multifaceted task. Each CAV is a complex entity consisting of multiple subsystems that need to be simulated in order to address the challenges arising from the different types of road network layout. The way that these subsystems are simulated (i.e. the tools and software used and the underlying assumptions) and the achieved level of detail are the criteria that lead to the categorisation of existing studies into two major groups:

- Studies using an architecture including traffic sub-microsimulation
- Studies using traffic microsimulation and an external component

A number of studies have used sub-microsimulation to simulate CAVs (Figueiredo *et al.*, 2009; Noort, Arem and Park, 2010; Pereira and Rossetti, 2012; Nitsche *et al.*, 2018). A sub-micro simulation tool such as PreScan and CarMaker, can simulate all the components of the vehicle accurately. This means, for instance, that the sensors of the vehicle are simulated individually and that their specifications (such as the scanning frequency, or the number of scanning beams) can be directly set. Moreover, they provide physical models for the car itself, such as the tyres, suspension and engine. The control algorithms for the "actors" of the simulation scenario (i.e. cars, traffic signals, etc.) are usually programmed in an external software such as Matlab/Simulink which communicates with the sub-micro simulation tool at every simulation step via a Transmissions Communication Protocol (TCP/IP), exchanging data.

However, this high level of detail comes at a price. A noteworthy drawback of this simulation approach is the high computational needs in order to run the tools of the framework simultaneously. This affects the size of experiments, which is relatively small compared to ordinary traffic microsimulation, and makes the comparison and generalisation of the results challenging. It is envisaged that an approach using a sub-microscopic simulator can be used to calculate the direct impacts (trip-by-trip basis) in a bottom-up simulation approach given that the results of the experiment can be transferred and generalised to a larger area such as a city by using scaling parameters. Connected and automated freight traffic based on the formation of platoons has some potential to impact the durability of roads and the reliability of bridges. While the effect on road durability can be avoided by variations in lateral positioning of trucks in the platoon, the effect on reliability of bridges is governed by distances between vehicles and their weights. The changed load pattern on bridges and its effect on the internal stresses in bridge components can be investigated by using bridge simulation models, like the finite element method or analytical models (in case of simple structures). Since the reliability of bridges is based on an evaluation of the probabilities that the bridge's internal forces will exceed its bearing capacity during the bridge's lifetime, probabilistic methods (O'Brien et.al., 2015) should be applied to investigate changes in extreme load effects due to truck platooning.

The second group of studies aims to address the aforementioned shortcomings regarding experiment size and computational needs by using a simpler simulation framework architecture. (Li *et al.*, 2013; Jeong, Oh and Lee, 2017; Rahman and Abdel-Aty, 2018; Stanek *et al.*, 2018; Papadoulis, Quddus and Imprialou, 2019). In most cases, a commercially available traffic microsimulation tool (such as AIMSUN, VISSIM, Paramics or SUMO) is used along with an external component. The microsimulation tool is applied to represent the infrastructure and creates the traffic in the predefined road system while the external component aims to simulate the CATS functionalities.

In both types of simulation framework mentioned above, there is a wide spectrum of parameters that can be adjusted (e.g. automation levels, market penetration rates, infrastructure-based parameters, CAV functionalities). These parameters could potentially formulate an immense number of scenarios. The selection of the ones that are most relevant to the project and the specification of their variation range will help reduce the computational effort.

3.1.4 System dynamics

Owing its heritage to control theory, system dynamics is a modelling technique where the whole system is modelled by modelling the sub-systems at component levels and aggregating the combined output. These sub-system components are defined by simple algebraic relationships and sometimes differential equations. These can be linear or non-linear in nature. The relationships can be continuous or discrete time events. Any feedback loops within the system (between the components) are captured through the sub-system equations that are defined. The behaviour of the whole system emerges from simulating the entire system level model that contains sub-systems. This technique is very useful in modelling management and social systems because their dynamics are generally of high order and non-linear and therefore cannot be represented by a single set of equations.

To illustrate the capabilities of system dynamics modelling, land use for housing was modelled in its simplest form.

Problem definition
What will be the effects of level 5 autonomous vehicles on land use for housing in a city that can be considered in zonewise sections?

Specifications

- Three zones were considered. Zone 1 was considered to be densely populated (city centre), Zone 2 was considered moderately populated (immediate zone outside city centre) and zone 3 was considered lightly populated (suburban area).
- Only Level 5 autonomous vehicles are thought to be able to drive themselves without any human intervention. This means that only those with access to level 5 AVs will consider relocating to a cheaper place for economic savings.

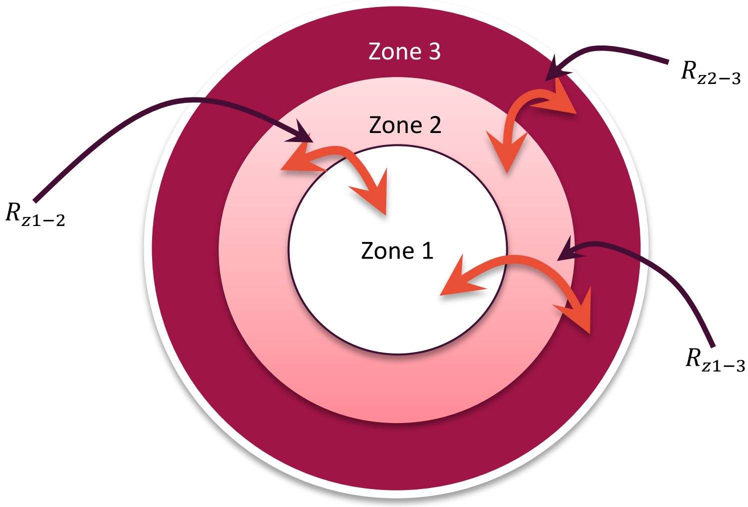

Figure 1: Schematic showing hypothetical city zones.

Assumptions and initial values

- Level 5 autonomous vehicles allow teleworking. However, the maximum time one is willing to spend inside the car is limited to certain level due to physical restrictions of movements inside the car. This maximum time was not considered at this stage.
- Everyone is assumed to be able to work during their commuting. However, only desk-based jobs can be done by telework.
- There is limited housing available and so the housing cost increases as population density increases.
- In practice, there is a delay between making the decision to relocate and actually relocating, which requires finding a house and arranging the necessary financial means. Here, it was assumed that there is no delay for simplicity.
- Usually, there is threshold in making the decision to relocate on the grounds of financial gain. Here, it was not implemented for now, again for simplicity.
- There is no effect of travel demand on increase in traffic. So, travelling time is not affected by change in population density in various zones.
- Total population is steady over the time. The increase in population due to net birth rate was not considered for simplicity.
- Initial population = 3138500, 1225200 and, 550000 for zone 1, zone 2 and zone 3, respectively.
- Net emigration (or immigration) in the city was considered to be 0.
- Working days = 260 days/year.
- Average salary = 25 euros/hour.
- Average housing costs = 15000, 12000 and, 10000 euros per year in zone 1, zone 2 and zone 3, respectively.
- Everyone is working in the city centre, so commuting distances are set accordingly to reflect this. Average commuting distances are set to 5, 10 and, 20

km for zone 1, zone 2 and zone 3, respectively. Also, the working population is thought to be steady, i.e., not ageing.

- Average travelling speed = 10, 15 and, 20 km/hr for zone 1, zone 2 and zone 3, respectively.
- Travelling cost per km = 0.2 euros/km.
- AVs are introduced at time = 10 years. This is done as a step change from 0 to X%.
- Total simulation is run for 20 years.
- Simulation time step is 1/52 year.

Model structure

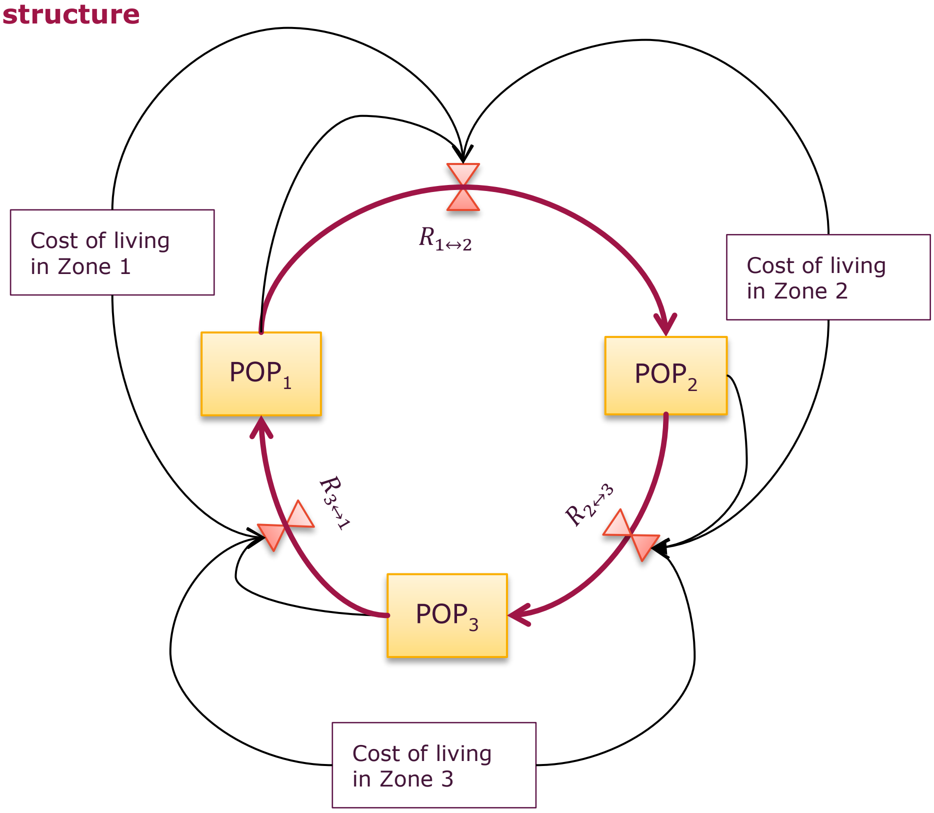

Figure 2: Model structure of zone wise population

Figure 2 shows the overall model structure of the population relocation within zone 1, 2 and 3. The rate at which the population is relocating is defined by $R_{i\leftrightarrow j}$. This rate is affected by the level of population such that if population is high, the rate of relocation is high since the AVs are introduced percentage wise. The rate of relocation is also affected by the cost of living, as this is the deciding step where comparison of costs is made and people are thought to move to a cheaper zone for economic gains. Here, it should be noted that in reality, there will be a threshold below which people would not relocate in order to save money. Savings of less significant amount will not be worthwhile considering the effort required for relocation. However, for simplicity, it is assumed that people start to relocate in the direction dictated by a positive or negative value of the relocation rate.

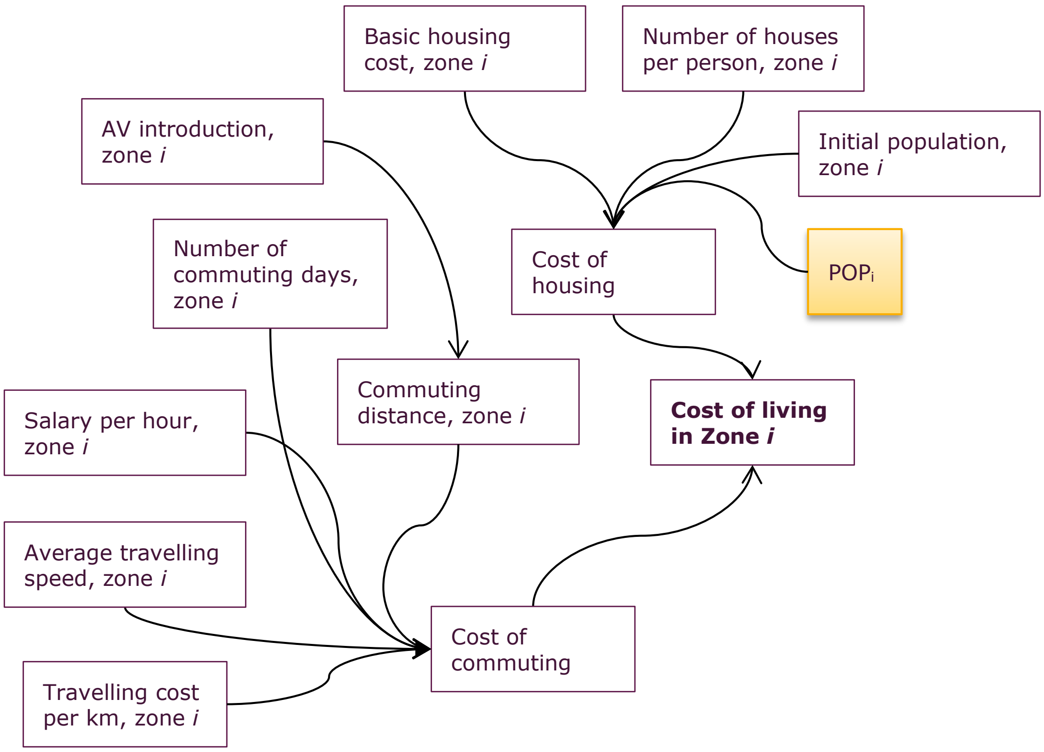

Figure 3: Flow diagram of cost of living

Figure 3 shows model structure of cost of living that is used in Figure 2. This is replicated for zone *i* = 1 to 3. It should be noted that the cost of housing has negative feedback from population such that increase in population in the zone will mean there will be an increase in demand for the house and therefore, house price will increase. The increased cost of housing means overall cost of living will also increase and will have a negative effect on population (since people are moving out) and therefore, the housing price will start to decrease again as a result. This is identified as a balancing loop.

It should be noted that population in each zone is divided into ones that own L5 AVs and ones that do not own L5 AVs. Following the above model structure, the following equations can be defined.

$$AV\ POP_{zone\ i} = POP_{zone\ i} \times AV\ fleet\ ratio;$$

$$nonAV\ POP_{zone\ i} = 1 - AV\ POP_{zone\ i};$$

$$(POP_{zone\ i})_{current} = (POP_{zone\ i})_{previous} + R_{i \leftrightarrow j} \times \Delta t;$$

$$Travelling\ cost\ AV\ owners\ per\ year = \left(2 \times number\ of\ commuting \frac{days}{year}\right) \times Average\ Commuting\ Distance \times Travelling\ Cost\ per\ km;$$

$$Travelling\ cost\ nonAV\ owners\ per\ year = \left(2 \times number\ of\ commuting \frac{days}{year}\right) \times \left\{(Average\ Commuting\ Distance \times Travelling\ Cost\ per\ km) + \left(\frac{Salary\ per\ hour \times Average\ Commuting\ Distance}{Average\ travelling\ speed}\right)\right\};$$

$$Housing\ cost\ per\ year = Present\ housing\ cost\ per\ year \times House\ per\ person \times \frac{(POP_{zone\ i})current}{(POP_{zone\ i})initial};$$

$$Cost\ of\ living\ per\ year\ for\ AV\ owners = Travelling\ cost\ AV\ owners\ per\ year + Housing\ cost\ per\ year;$$

$$Cost\ of\ living\ per\ year\ for\ nonAV\ owners = Travelling\ cost\ nonAV\ owners\ per\ year + Housing\ cost\ per\ year;$$

$$Relocation\ rate\ of\ AV\ owners = AV\ POP_{zone\ i} \times \left\{1 - \left(\frac{Cost\ of\ living\ per\ year\ for\ AV\ owners_{zone\ i}}{Cost\ of\ living\ per\ year\ for\ AV\ owners_{zone\ j}}\right)\right\};$$

$$Relocation\ rate\ of\ nonAV\ owners = nonAV\ POP_{zone\ i} \times \left\{1 - \left(\frac{Cost\ of\ living\ per\ year\ for\ nonAV\ owners_{zone\ i}}{Cost\ of\ living\ per\ year\ for\ nonAV\ owners_{zone\ j}}\right)\right\};$$

$$R_{i \leftrightarrow j} = Relocation\ rate\ of\ AV\ owners + Relocation\ rate\ of\ nonAV\ owners;$$

Here, i is zone number of the quantity being calculated and j is the zone number to which it is compared. Positive value of $R_{i \leftrightarrow j}$ means that the population is going from zone i to j and opposite is true if negative.

During simulation, the above equations were repeated for time, $t_{initial}$ to t_{final} in Δt steps.

In addition, as mentioned in assumptions and initial conditions, $AV\ fleet\ ratio$ was considered to be 0 before year 10 and x% after year 10. X% being 10%, 20%, and so forth up to 100%. The model was simulated for time frame of 20 years.

Results

As shown in Figure 4, the zone wise population was initially concentrating towards the city centre as evident by increase in population in zone 1 and 2 and decrease in population in zone 3. The L5 AVs were introduced after technology was mature. For this reason, year 10 was chosen. It should be noted that only 3 cases, L5 AV 0%, 50% and 100% are shown in Figure 4.

The model did not have any mathematical equation that suggested any sinusoidal behaviour in population dynamics. The effect seen in Figure 4 is purely due to two feedbacks, one from population to rate of relocation and other one from population to housing cost. It is possible to minimise oscillations by carefully choosing the values of the economic penalty such that the relocation is dampened and people do not move in and out of the zone frequently.

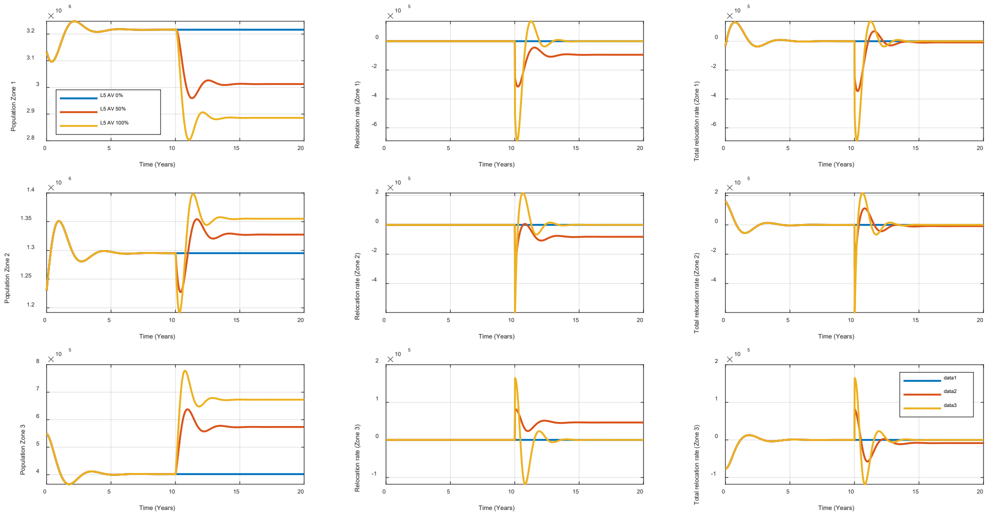

Figure 4: Zone wise population, their relocation rates due to L5 AVs and overall relocation rates.

Once a working model is obtained, one can extend that to other questions such as what would be the effect of zonewise road pricing on land use for housing, what would be the effect on emissions due to change in land use for housing due to L5 AVs, etc.

It is clear that zones 2 and 3 became more attractive to live in when L5 AVs were in the fleet. The value of time (salary considered as office time was invested in commuting) outweighed the travelling costs due to longer distance. So, as the L5 AV penetration increased the population in zone 1 progressively decreased and the population in zone 2 increased (Figure 5).

Population density was calculated by assuming that the zones are concentrically arranged as shown in Figure 5.

It must be noted that this model is only an illustrative example to demonstrate the techniques of system dynamics. It is far from realistic and no conclusions may be drawn from the results. It is intended only to show the potential for analysis by using system dynamics. To develop realistic forecasts several elements of the model would need to be refined. Currently, many land use and transport interaction (LUTI) models exist and some of them were created using system dynamics and validated using real city data (Pfaffenbichler, Emberger, & Shepherd, 2010; Shepherd, 2014). It would be possible to adapt those models to include AVs as an additional element and look at the effect on land use.

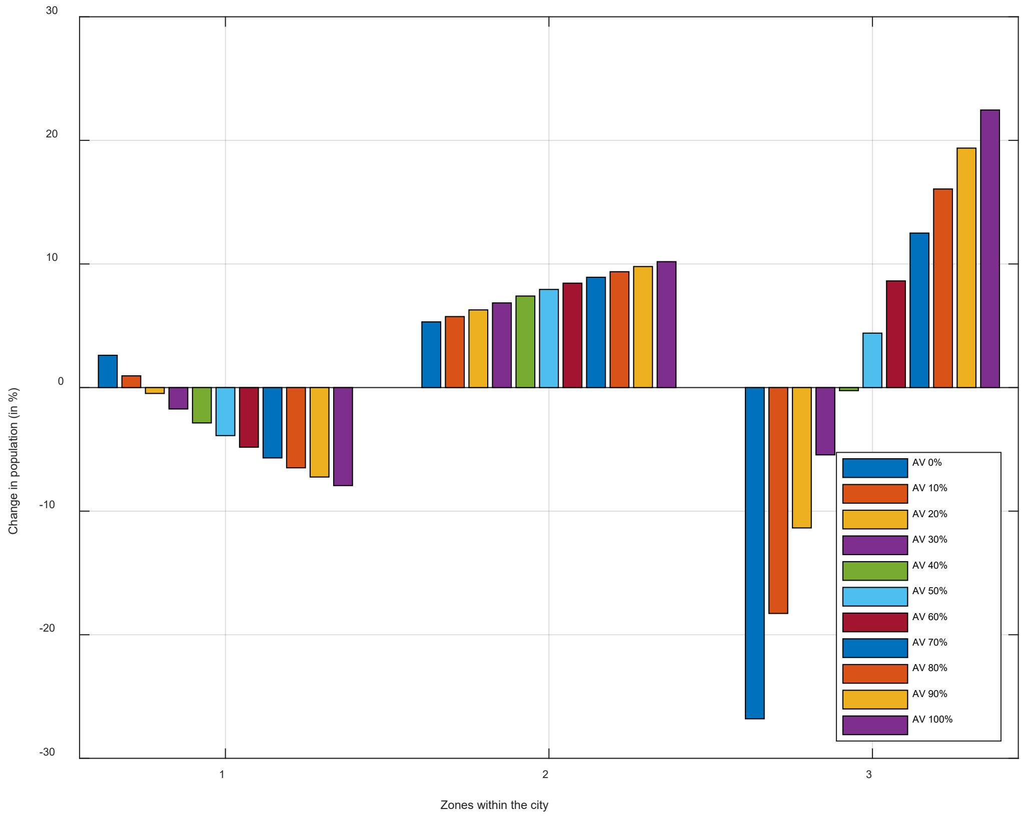

Figure 5: Percentage change in population in zone 1, 2 and 3 as an effect of introduction of L5 AVs.

3.2 Which methods are suitable for which impacts?

Surveys and studies of willingness-to-pay are well-suited for estimating the likely demand for automated vehicles and how fast they will penetrate the market. Surveys are also well-suited for studying trust in technology.

To the extent that studies applying traffic simulation have estimated various impacts of connected and automated vehicles, formal syntheses of such studies by means of meta-analysis is a suitable method.

Mesoscopic simulation and activity-based-modelling are well-suited for assessing modal split, road network loads and vehicle utilisation rate.

Operations research methods are widely used for freight transport, such as optimisation algorithms for fleet management and tour planning. For example, a typical parcel delivery problem requires to calculate the fastest tours from a given depot to service a set of customers. Freight transport is plannable and therefore it is reasonable to assume that freight operators optimise their business according to market-economic principles. In comparison, passenger transport is based on individual preferences. This makes it impossible to plan exactly, hence simulation is a more reasonable approach to assess the performance of the passenger transport system. Besides freight transport, fleet and tour optimisation can also be applied to urban shuttles applications. Like freight transport, shuttles are operated by, e.g., a public transport authority, which makes them plannable. With given requirements such as passenger demand and the area of operation, the ideal fleet and routes can be optimised.

Methods for fleet and tour optimisation primarily assess the direct impacts such as vehicle operating costs and freight transport costs. Furthermore, since the actual routes of freight vehicles are generated and evaluated, wider impacts such as air pollution and noise pollution can be obtained or upscaled from simulation.

3.3 Combining methods

All methods listed above are at least potentially relevant for predicting the impacts of CATS. An integrated method for prediction, by which is meant a method that is able to predict as many potential impacts of CATS as possible, confer the list of potential impacts given Chapter 2, can only be developed by combining several methods.

To give an example: Several studies have used traffic simulation to estimate the potential impacts of CATS on road capacity. The feasibility of combining the results of these studies can be assessed by means of an exploratory meta-analysis. The key criterion is whether the studies produce similar results. If they do, it makes sense to combine them and estimate a mean effect. The differences in results between different studies will then serve as an indicator of the uncertainty surrounding the mean effect.

However, a high level of heterogeneity in results by itself may not always prevent a formal synthesis; likewise, similarity in results can be coincidental. It is therefore necessary to examine every study critically. In particular, results of traffic simulation depend strongly on the assumptions made. The realism of the assumptions must always be assessed.

Few, if any, of the studies of the potential increase in road capacity associated with CATS have considered a potential rebound effect. To estimate the rebound effect, studies of the elasticity of vehicle kilometres of travel with respect to road capacity can be applied. Again, if these studies produce similar estimates, combining them makes sense. If estimates vary widely, a choice must be made regarding which studies are methodologically most rigorous.

Another way to utilise traffic simulation is to upscale partial results with operations research methodologies. As mentioned in the previous section, fleet and tour optimisation will compute the actual routes of each individual CAV if the operation is plannable. Therefore, for some use cases within freight transport such as e.g. automated urban delivery, it is possible to simulate the performance of an automated freight vehicle in small scenarios and upscale the results to the full route. These scenarios must be representative and cover the realistic use case as well as possible, e.g., a route of 3 km in a typical city center area or a suburb area and under certain traffic load. If tour optimisation computes a full route of 30 km for a freight vehicle, the overall performance can be obtained by combing several simulated scenarios. A flowchart of this type of combination is illustrated in Figure 6.

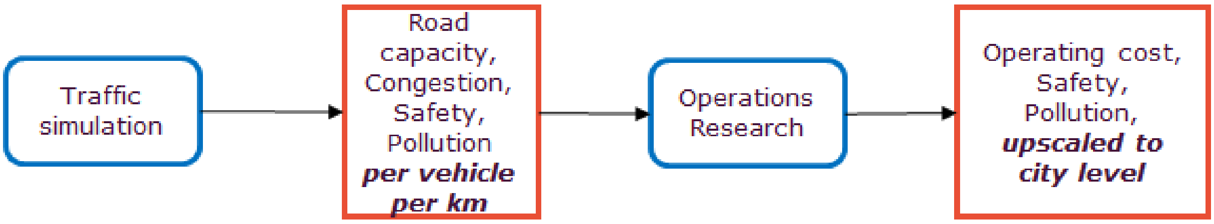

Figure 6: Flow chart for route optimization in freight transport.

The next chapter explains how quantitative relationships describing potential impacts of CATS can be developed.

4 Estimating and quantifying primary impacts

The objective of this chapter is to show how potential primary impacts of connected and automated vehicles can be estimated and quantified. Primary impacts are the intended impacts of automation technology; the benefits it is intended to bring about. The main approach taken in the chapter is to quantify impacts as dose-response curves using the market penetration rate of automation technology as the dose and the size of an impact as the response. The chapter presents dose-response curves estimated on the basis of studies identified in a literature review. Not all primary impacts can be represented by means of dose-response curves. Other approaches that can be used for quantifying impacts are discussed. The next chapter discusses secondary impacts (unintended behavioural adaptation) of connected and automated vehicles.

This chapter develops quantified relationships describing the impacts of connected and automated vehicles. It is shown how the following impacts can be quantified:

1. Road capacity (for sections and junctions)
2. Congestion
3. Travel time
4. Valuation of travel time
5. Individual access to travel
6. Use of shared mobility services
7. Road safety
8. Fuel consumption and emissions
9. Vehicle ownership costs
10. Willingness-to-pay for automation technology

Most of these impacts are quantified as dose-response curves, which have the market penetration rate of automated vehicles as the dose and the size of an impact as the response. All these impacts are uncertain, and the range of estimates found in available studies is an indicator of uncertainty. First, a framework for quantification is proposed.

4.1 A framework for quantifying impacts of CATS

Most statistical analyses designed to estimate the impacts of technology rely on real-world data, e.g. the actual accident involvement of cars with electronic stability control compared to cars without the system. Real-world CATS data are still in their infancy and do not reflect the system-wide impacts associated with a high market penetration rate of automation technology. An alternative to real-world data is to utilise simulated data and then apply statistical methods to forecast impacts. In Levitate, simulation can provide the data that are needed to forecast short, medium and long term CATS impacts.

The first step is to develop statistical models of the relationship between an impact category (e.g. safety) and its predictors (e.g. market penetration rate of CATS by level of autonomy, demand for CATS, types of new mobility services supported by CATS) for different types of traffic environment. These statistical models will be based on the simulated data obtained from a wide range of simulation scenarios. Table 2 gives an example of a hypothetical scenario describing how travel demand (vehicle kilometres of travel) and speed variation changes as a function of levels of implementation of automation technology.

In the current condition, the level of implementation of automation technology is set to zero, i.e. all vehicles are manually operated, but may have driver support systems already on the market and in wide use, like electronic stability control. To develop a scenario, assumptions are made regarding the mix of vehicles at different levels of automation in the short, medium and long term. In Table 2, it is assumed that in the long term, there will be vehicles at all SAE automation levels from 1 to 5 in an almost equal mix (bottom row of Table 2). It is further assumed that 70 % of vehicles will be connected. For both the short-, medium- and long term three forecasts are developed: average, pessimistic and optimistic. This gives a range of potential impacts.

Table 2: Potential impacts of CATS on travel demand and speed variation as a function of levels of implementation of automation technology.

Scenarios		Simulation Input Variables								
		Overall Traffic Demand	SAE level 0 MPR	SAE level 1 MPR	SAE level 2 MPR	SAE level 3 MPR	SAE level 4 MPR	SAE level 5 MPR	Network Speed Distribution	% Connected Vehicles
Baseline Model (present)	Current Condition	Real-world traffic demand data	100%	0%	0%	0%	0%	0%	Real-world data	0%
Short Term	Pessimistic	6% increased	90%	4%	2%	2%	1%	1%	5% less s.d.	5%
	Average	4% increased	80%	6%	4%	4%	3%	3%	10% less s.d.	10%
	Optimistic	2% increased	70%	8%	6%	6%	5%	5%	15% less s.d.	15%
Medium Term	Pessimistic	12% increased	60%	10%	8%	8%	7%	7%	20% less s.d.	20%
	Average	10% increased	50%	12%	10%	10%	9%	9%	25% less s.d.	30%
	Optimistic	8% increased	40%	14%	12%	12%	11%	11%	30% less s.d.	40%
Long Term	Pessimistic	18% increased	30%	16%	14%	14%	13%	13%	35% less s.d.	50%
	Average	16% increased	20%	18%	16%	16%	15%	15%	40% less s.d.	60%
	Optimistic	14% increased	10%	20%	18%	18%	17%	17%	45% less s.d.	70%

The relationships given in Table 2 may in principle be stated in the form of a predictive equation like this:

$$Impact_i = b_0 + b_1 * predictor_{1i} + b_2 * predictor_{2i} + \ldots + b_n * predictor_{ni} + u_i \quad (1)$$

In equation (1) impact i is the dependent variable. It can be the impact on safety, road capacity, travel time, etc. The predictors may include market penetration rate of SAE automation levels 1-5, or changes in traffic demand. The aim of the approach is to combine all the statistical models created for each road network element in order to create an integrated impact assessment toolbox that could be used by WP8 in the Policy Support Tool. Using this toolbox, a city or an organisation can input their own data (e.g. traffic demand data) for the corresponding market penetration rate that they want to investigate or forecast for their city and calculate the impact. The challenges arising when developing these models include the selection of the appropriate predictors as well as the selection of the appropriate functional form for the relationship between a predictor and an impact.

There is, as noted already, uncertainty about how fast automation technology will penetrate the market. Thus, many different assumptions can be made about this. To keep estimation of dose-response curves analytically tractable, the following guidelines are proposed:

1. The dose is stated as the percentage market penetration of automated vehicles, i.e. vehicles with at least SAE level 3 automation technology, but preferably level 5 automation technology.
2. The response is stated as the percentage change in an impact indicator associated with different percentage levels of market penetration of automated vehicles.
3. One impact is modelled at a time, i.e. each dose-response curve refers to a single impact, preferably a single indicator of an impact.

The next sections present the dose-response curves that have been developed so far. It is envisaged that the more dose-response curves than those presented in this report will be developed during the Levitate project.

4.2 Road capacity

4.2.1 Motorways

Simulation studies or mathematical models developed by Tientrakool et al. (2011), Fernandes and Nunes (2012), Shladover et al. (2012), Atkins (2016), Ye and Yamamoto (2018) and Shi, He and Huang (2019) have estimated changes in road capacity as a function of market penetration of high level automation technology. All these studies refer to the capacity of a driving lane on a motorway. The results of the studies vary and depend on the specific assumptions made about, among other things, distance between vehicles, vehicle dimensions, mean speed, and deceleration.

To illustrate the importance of the assumptions made for the results of the study, a couple of examples will be discussed in detail. Tientrakool et al. (2011) divide vehicles into three groups:

1. Manual vehicles
2. Vehicles with sensors but no technology to communicate with other vehicles
3. Vehicles with sensors and technology for communicating with other vehicles

The first type of vehicles represents current technology. The third type of vehicle represents a low level automation technology. Manual vehicles are assumed to follow each other with a (mean) time headway of 1.1 seconds. The following distance in metres then becomes:

$D_m = 1.1 \cdot V/3.6$

V is speed in kilometres per hour. Dividing this by 3.6 gives speed in metres per second. Thus, if speed in km/h is 60, the following distance becomes:

Following distance in metres = 1.1 · 60/3.6 = 18.3 metres. Following distance is the distance between the rear of vehicle 1 (the first in a row of vehicles) and the front of vehicle 2 (the second in a row of vehicles). The length of a vehicle was assumed to be 4.3 metres. Traversing a line across the road will take a vehicle 0.26 seconds at a speed of 60 km/h (at this speed, a vehicle travels 16.67 metres per second; traveling the length of the vehicle thus takes 4.3/16.67 = 0.258 seconds). Thus, serving one vehicle at the assumed following distance will take 1.1 + 0.26 = 1.36 seconds. Hourly capacity is 3600/1.36 = 2647 (3600 is the number of seconds per hour).

It is immediately seen that if a following distance of 1.5 seconds, a speed of 40 km/h and a vehicle length of 5 metres had been assumed, estimated road capacity would have been different (under these assumptions it would have been 1846). Different assumptions about the initial road capacity lead to different estimates of the increase in capacity that can be attained by spacing vehicles more densely or by assuming that they can communicate and react faster to changes in speed than a human driver.

Tientrakool et al. (2011) assume that for vehicles that can communicate, the following distance can be shortened and set equal to the distance traversed during the reaction time of the communication system. They assume a communication delay of 0.181 seconds. At a speed of 100 km/h, this translates into a following distance of 5.0278 metres. Tientrakool et al. (2011) estimate that lane capacity per hour can be increased from 2869 with only manual vehicles to 10721 with 100 % communicating vehicles. The paper by Fernandes and Nunes (2012) reaches similar conclusions about the potential for increasing road capacity by means of connectivity and platooning.

They estimate an increase in lane capacity of up to 175 % at a speed of 36 km/h for platoons of vehicles consisting of up to 20 vehicles per platoon, spaced by only 1 metre. Each vehicle was assumed to be 3 metres long, which is shorter than virtually all cars found in traffic today. It is likely that the potential increase in capacity would be smaller if there is a mix of vehicles of different lengths. At a speed of 72 km/h, the maximum potential increase in lane capacity was estimated to be 370 %, meaning that capacity increases almost by a factor of five. Fernandes and Nunes give the following formula for road capacity:

$$C = v\frac{n}{ns+(n-1)d+D}$$

C is capacity, indicated as vehicles passing per second (to obtain hourly volume, multiply by 3600), n is the number of connected cars in each platoon (1 if cars drive independently and are not platooned), s is vehicle length in metres (it can be the average value of vehicles differing in length), v is speed in metres per second, d is intraplatoon spacing in metres, and D is spacing between platoons in metres.

To give an example, assume that no vehicles are connected (n = 1), that they drive at 40 km/h (11.1 metres per second), that mean length is 6 metres, and that mean time headway is 2 seconds (equal to 22.2 metres). The factor d drops out (is set to zero), whereas D represents headway. We get:

$$C = 11.1\frac{1}{1 \cdot 6 + 0 + 22.2} = 0.3936$$

This corresponds to an hourly capacity of 1417 (0.3936 · 3600). Now assume that vehicles are connected in platoons each consisting of 15 cars (n = 15). The spacing of cars within each platoon is 2 metres (d = 2). The spacing between platoons is 55.6 metres (5 seconds), allowing pedestrians to use the gap between two platoons to cross the road. We now get:

$$C = 11.1\frac{15}{15 \cdot 6 + 14 \cdot 2 + 55.6} = 0.9591$$

This corresponds to an hourly capacity of 3452 vehicles. By varying the assumptions made, it is possible to estimate changes in capacity for different speeds, different vehicle lengths, different lengths of platoons, different spacing within platoons and different spacing between platoons. Thus, the framework provided by Fernandes and Nunes (2012) is very flexible.

Each of the studies listed above provides an estimate of road capacity for various levels of market penetration of connected and automated vehicles. These estimates can be plotted in a diagram with market penetration rate as abscissa (horizontal axis) and change in road capacity as the ordinate (vertical axis). Figure 7 shows a curve fitted to the results of the studies of Tientrakool et al. (2011) Fernandes and Nunes (2012), Shladover et al. (2012), Atkins (2016), Ye and Yamamoto (2018) and Shi, He and Huang (2019).

The curve was fitted to the data points by examining which of several functional forms best fitted the data points. The following functional forms were compared: linear, exponential, logarithmic, and polynomial. For the data points in Figure 7, a second degree polynomial was found to best fit the data points. All data points have the same weight. It was not possible to assign a unique statistical weight to each data point based on its statistical precision, as estimates of the precision of data points are not consistently reported in the primary studies. It is nevertheless possible to assess informally whether summarising the data points in terms of a dose-response curve makes sense and is a representative summary of the data points.

First, does the fitted curve pass through the range of data points, i.e. does it have both positive and negative residuals? It is seen that the curve fitted to the data points in Figure 7 satisfies this criterion. Second, if there are enough data points, their distribution can be studied by means of tools for exploratory analysis in meta-analysis, such as stem-and-leaf plots (Elvik and Ramjerdi 2014). A stem-and-leaf plot can give an indication of whether estimates are symmetrically distributed and whether there are outliers. Some of the data points in Figure 7 appear to be outliers, as indicated by the dashed area. These data points were included when the function was fitted.

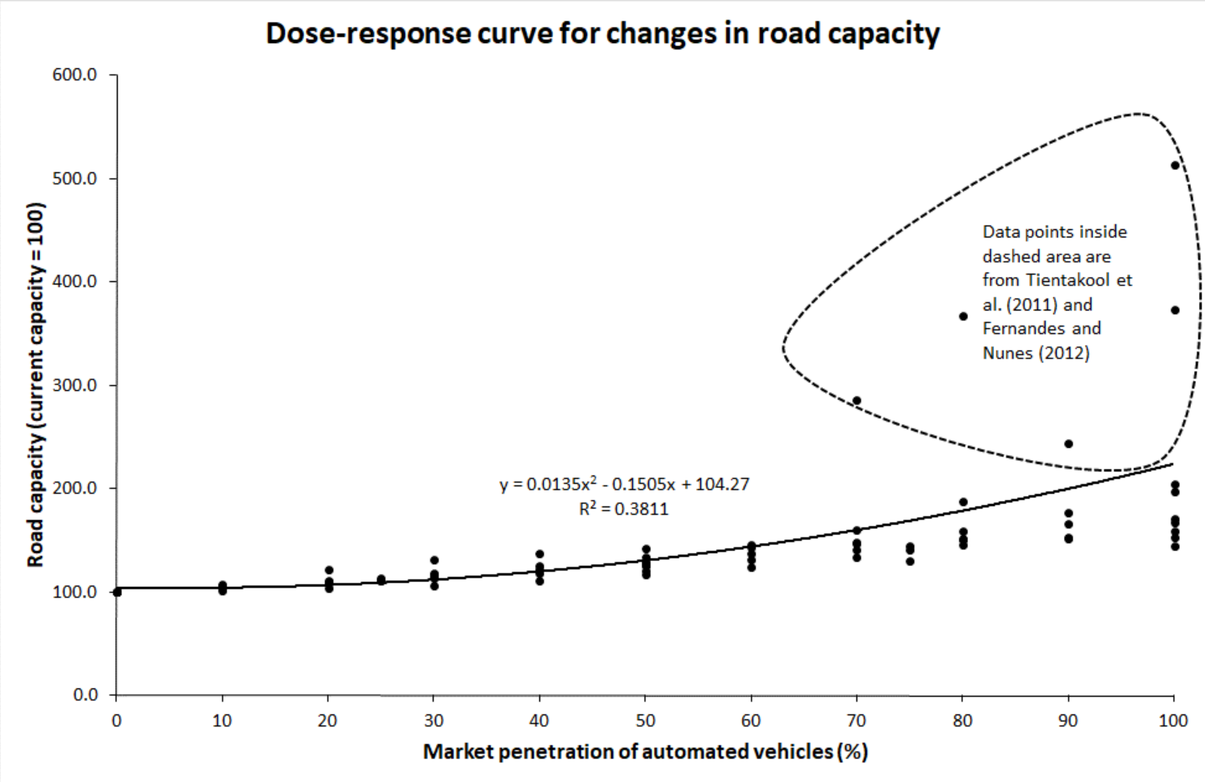

Figure 7: Dose-response curve for changes in road capacity as a function of market penetration of high level automation technology

It is seen that results are perfectly consistent, in that all studies predict an increase in road capacity. The dispersion of the data points is also quite small up to 60% market prenetration rate for connected and automated vehicles. At higher market penetration rates, the early studies of Tientrakool et al. (2011) and Fernandes and Nunes (2012) predicted a considerably larger increase in road capacity than subsequent studies. The dispersion of the data points is an indication of uncertainty, which includes the impacts on study results of differences between studies with respect to the assumptions made about speed, acceleration, deceleration, platoon length, etc. It is seen that uncertainty increases at higher rates of market penetration of connected and automated vehicles. The function fitted predicts a reduction of road capacity at low penetration rates of automated vehicles.

It is important to note that all dose-respanse curves, like the one fitted in Figure 7, apply only when all else is equal (commonly referred to as ceteris paribus). This means, for example, that there is no change in road layout and no change in traffic volume. The curve only shows that it is possible to serve a larger number of connected and automated vehicles than the number of human driven vehicles.

All studies included in Figure 7 refer to lanes on motorways. However, the most common type of road in most countries is a rural, undivided two-lane road. On this type of road, there will be a transition period during which a mixture of manual and automated vehicles will share the road. Automated vehicles will then have to adopt headways that allow them to stop, or at least greatly reduce speed when the vehicle in front of them, possibly a manually operated vehicle, brakes or stops. The choice of headway in this situation can be treated as a tradeoff between capacity and safety, to which an optimal solution in principle exists (Liu et al. 2020). If long headways are adopted, road capacity is reduced and there may be delays and congestion. If short headways are adopted, there will be rear-end collisions. Liu et al. (2020) assigned monetary values to delays and rear-end collisions in order to determine the following distance that miminised the sum of the costs of delays and collisions. They assumed that automated vehicles will have a reaction time of 0.2 seconds and decelerate at 5 metres per square second. Using data on accidents, traffic speed and following distances from two highways in the United States, they found that optimal following distances were between 63 and 89 feet (19 and 27 metres). Lane capacity would then be between 1888 og 2080 vehicles per hour.

Various scenarios were developed, but the shortest following distances were in no case shorter than about half of the distances found in the baseline scenario. This study suggests that a doubling of road capacity is likely to be largest attainable effect if it is assumed that automated vehicles must keep sufficient following distance to be able to stop or minimise accident severity in a given situation. As long as there is a mixture of automated and manual vehicles, the choice of following distance will be a tradeoff between safety and capacity. It is only when 100% market penetration of connected and automated vehicles has been reached that this tradeoff can be resolved by prioritising an increase in road capacity exclusively.

4.2.2 Junctions

Atkins (2016) modelled changes in capacity in signalised junctions and roundabouts as a function of the market penetration of automated vehicles. The dose-response curves fitted to their results are shown in Figure 8.

Junction capacity was estimated to increase by about 40% in roundabouts and nearly 30% in signalised junctions. It seems reasonable that the increase in junction capacity will be smaller than the increase in lane capacity, as connectivity and platooning are unlikely to be applicable in junctions. Le Vine, Zolfaghari and Polak (2015) note that travel comfort in automated cars may be experienced as low if they drive like a human driver. Passengers in cars experience acceleration, braking and turning as more violent and less comfortable than drivers do.

If automated cars are programmed to have acceleration and deceleration equal to that of light rail or long distance rail, the capacity of signalised junctions will be reduced.

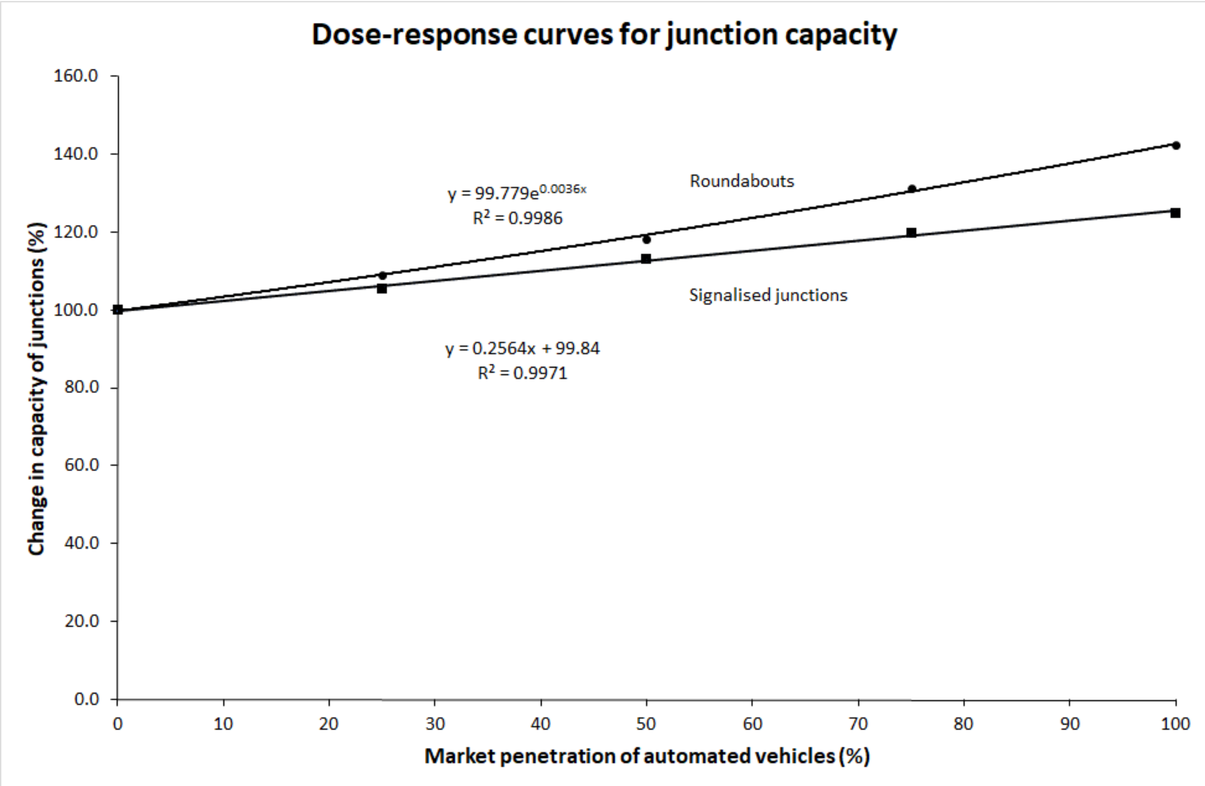

Figure 8: Dose-response curves for changes in junction capacity as a function of market penetration of automated vehicles

4.3 Congestion

4.3.1 Road sections

When the capacity of road sections and junctions increases, congestion may be reduced. Potential impacts of connected and automated vehicles on congestion delays have been studied by Atkins (2016) and Kockelman et al. (2018). A dose-response curve fitted to the average results of their studies is shown in Figure 9.

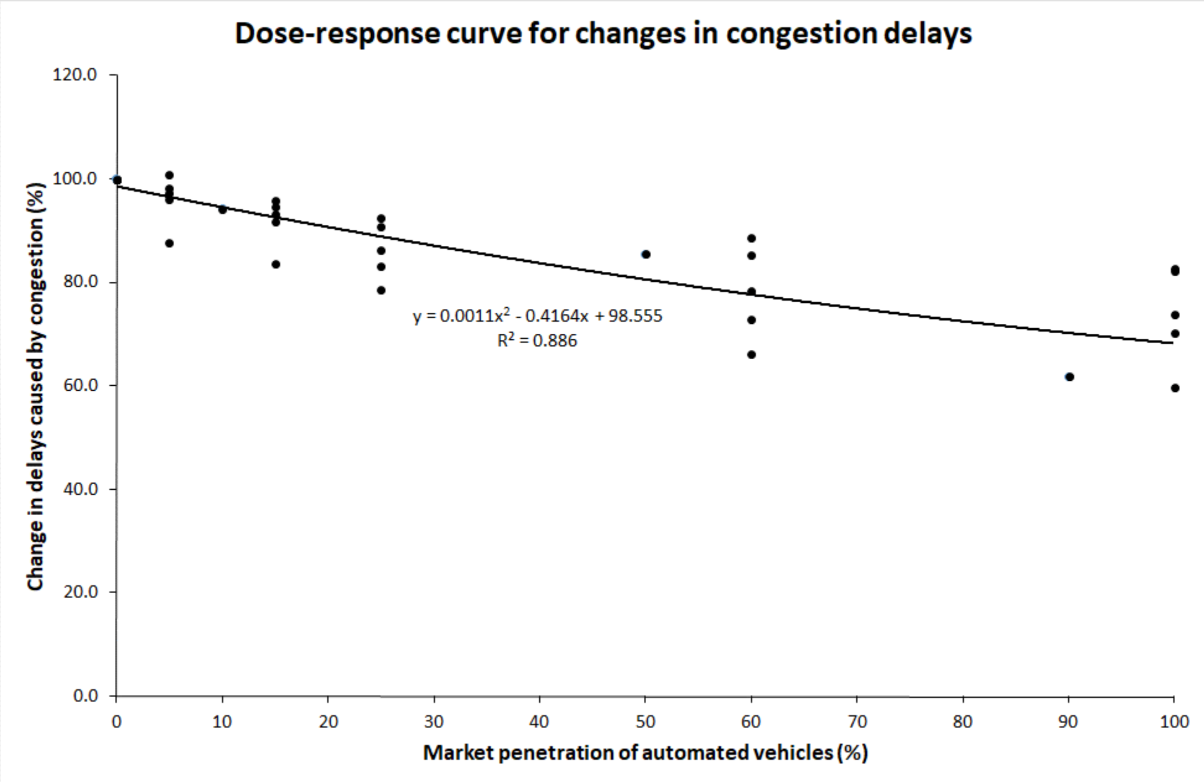

Figure 9: Dose-response curve for changes in delays caused by congestion

The dose-response curve passes through the data points and is a representative summary of the results. At 100% market penetration of automated vehicles, congestion delays can be reduced by about 25%. The dispersion of data points around the curve fitted to them reflects differences in the precise definition of congestion and in congestion severity. On roads that regularly become congested, usually about 20-30% of traffic is congested.

4.3.2 Junctions

When junction capacity increases, one should expect delays to be reduced, all else equal. Studies by Atkins (2016) and Morando et al. (2018) find that this is the case, but their results differ substantially and vary according to whether junctions are signalised or roundabouts. Results for roundabouts are shown in Figure 10.

Both studies find that delays will be reduced in roundabouts. According to Atkins (2016) the maximum reduction of delays is a little more than 20%. According to Morando et al. (2018), maximum reduction of delays is 50%. The shape of the dose-response curves also differs. The curve fitted to the study by Atkins suggests an accelerating reduction of delays as market penetration of automated vehicles increases. The curve fitted to the study of Morando et al. suggest that reductions in delays will be fast when automated vehicles start to penetrate the market, but then slow down as these vehicles reach 100% market penetration. Note that this refers to the level of market penetration only, and says nothing about how fast or slow in time the market penetration of automated vehicle will be.

Despite these differences, the studies agree that delays are likely to reduced.

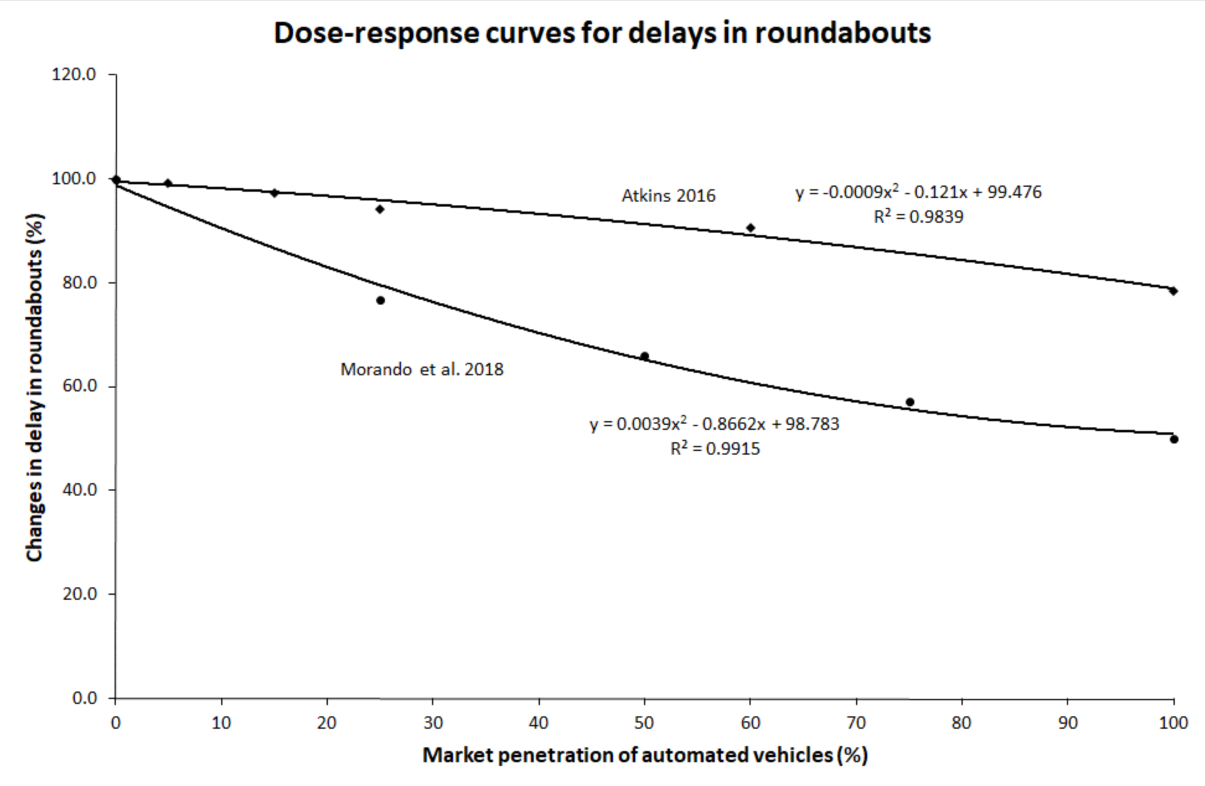

Figure 10: Dose-response curve for changes in delays in roundabouts

Figure 11 shows results for signalised junctions. Both studies conclude that delays will be reduced, but the reductions are much smaller than in roundabouts. According to Atkins (2016) the maximum reduction in delays is about 5%. According to Morando et al. (2018) the maximum reduction in delays is about 16%. The shape of the dose-reponse curve is the same for both studies: a second degree polynomial suggesting that reductions in delays are greatest at the beginning and then slow down.

One may wonder whether signalised junctions will be needed when all vehicles are automated. Surely, the vehicles will have vehicle-to-vehicle communication systems which can decide which vehicle has priority. However, it may still be necessary to stop all vehicles to enable pedestrians to cross the road.

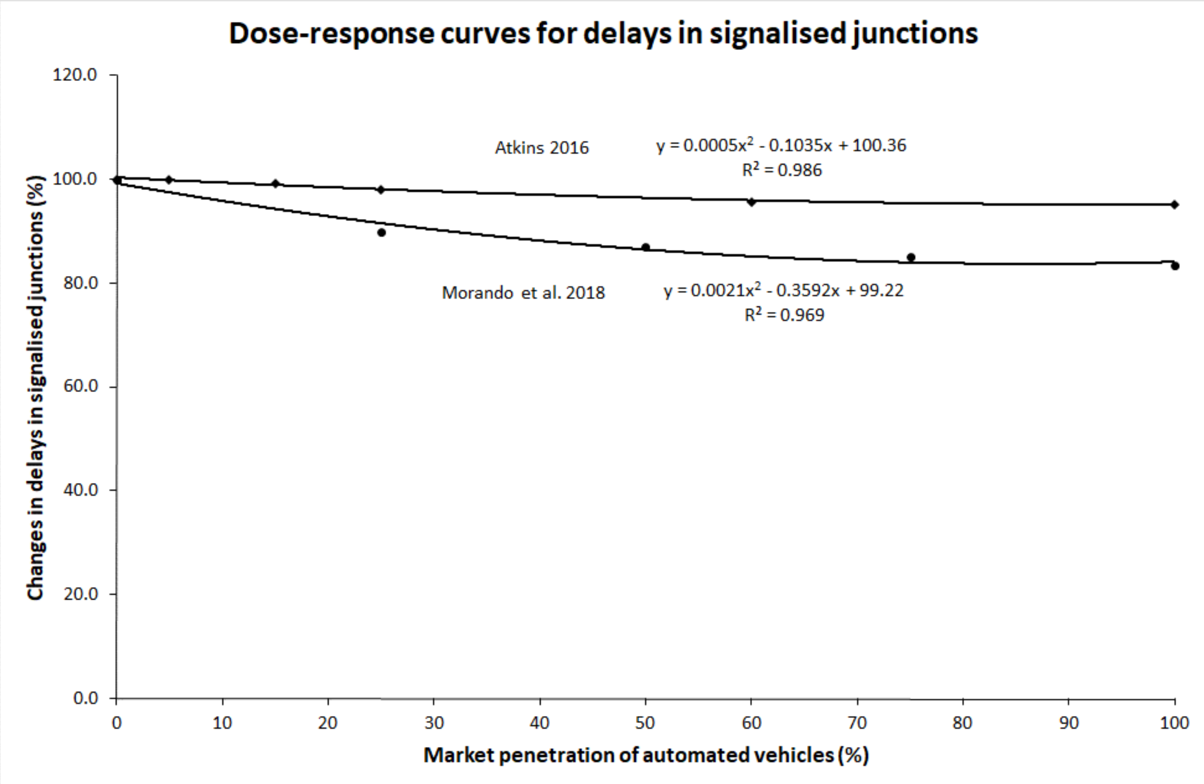

Figure 11: Dose-response curves for delays in signalised junctions

4.4 Travel time

Changes in mean travel time as a result of vehicle automation has been studied by Atkins (2016) and Booz Allen Hamilton et al. (2018). The results of these studies, and a dose-response curve fitted to average results, are shown in Figure 12.

Estimates are a mixture of those applying to congested conditions and those applying to non-congested conditions. Combining estimates referring to these conditions reflects the reality that part of traffic is in congested conditions, part of it in non-congested conditions. One might apply weights, by assuming, for example, that 25% of traffic is congested and 75% non-congested. However, at the current stage of analysis, emphasis is on exploring whether different studies give sufficiently similar results that averaging and combining them makes sense. From that point of view, the function fitted to the data points in Figure 12 makes sense. It passes right through the data points. It suggests that a maximum saving in travel time of about 17% is possible. This might seem surprisingly small compared to the potential increase in road capacity shown in Figure 7. However, it should be remembered that most of the time, there is spare capacity on most roads. Most of traffic is not delayed by congestion.

Besides, if it assumed that automated vehicles will comply with speed limits, these limits represent an upper bound on how large savings in travel time it is possible to accomplish. Against this background, it does not seem unrealistic that the savings in travel time are likely to be less than 20%.

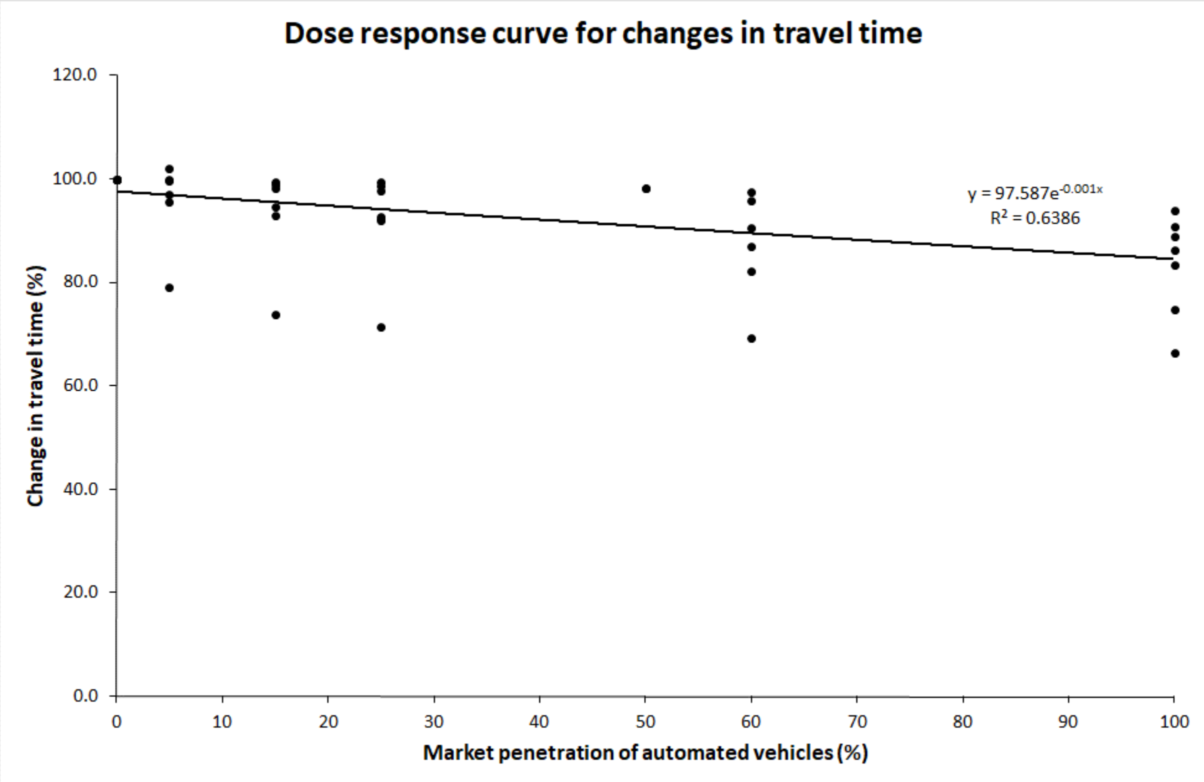

Figure 12: Dose-response curve for change in mean travel time.

4.5 Valuation of travel time

Fully automated vehicles are expected to lead to a lower valuation of travel time, since travel time can be used productively or for relaxing. It will then be felt as less wasted than today, when a car driver has to concentrate of the driving task and has few opportunities for engaging in other activities.

There are few empirical studies of the valuation in travel time in automated vehicles, and the few that exist obviously only produce hypothetical findings, as automated vehicles are still not in widespread use. Table 3 summarises the results of delphi surveys and stated preference studies of the valuation of travel time in automated vehicles.

Willumsen and Kohli made a Delphi survey of 45 transport experts and asked them about changes in the valuation of travel time associated with vehicle automation. They made a distinction between commuting trips, trips made as part of work and other trips (leisure trips). As can be seen from Table 3, the mean values of travel time given by the experts indicated a relative reduction of about 15% regardless of trip purpose.

Kolarova et al. (2018) studied the value of travel time savings in conventional and automated vehicles. They found that the value of time in a private automated vehicle was about 45-55% lower than in a conventional private car. In a driverless taxi, the value of time was about 15-35% lower than in a conventional private car.

Table 3: Relative valuation of travel time in automated vehicles. Current valuation = 100.

Study	Vehicle use	Current valuation of travel time = 100			
		Commute	In work	Leisure	Mean
Willumsen, Kohli 2016	Individual	84.1	87.1	85.2	85.5
		Low inc	Middle inc	High inc	Mean
Kolarova et al. 2018	Individual	45.4	44.3	57.8	49.2
	Shared	69.0	67.3	87.7	74.7
		Low inc	Middle inc	High inc	Mean
Steck et al. 2018	Individual	69.4	69.5	69.5	69.5
	Shared	90.0	90.0	90.0	90.0
		Driver	Passenger		Mean
Flügel et al. 2019	Individual	69.6	44.6		57.1
	Shared	79.0	88.1		83.8
All studies	**Individual**				**65.3**
	Shared				**82.8**

The study by Steck et al. (2018) is very similar to the study by Kolarova et al. (2018), but the results differ slightly. A reduction of the valuation of travel time in automated vehicles of about 30% was estimated for individual use of the automated vehicles. For shared use, the reduction in the valuation of travel time was 10%. Flügel et al. (2019) made a stated preference study. The mean reduction in the value of travel time savings was a little more than 40% for individual use of automated vehicles and around 15% for shared use of automated vehicles.

The results of studies comparing individual and shared use of automated vehicles are consistent. All studies show a larger reduction of the value of time for individual use, suggesting that individual use of automated vehicles will be felt as more comfortable than shared use.

The studies suggest that, as a first approximation, one may assume a 35% reduction of the value of travel time saving for individual use of automated vehicles, with a lower limit of 15% and an upper limit of 50%. For shared use of automated vehicles, a mean reduction of the value of travel time savings of 15% is suggested, with a lower limit of 10% and an upper limit of 25%.

4.6 Individual access to travel

A potential advantage of fully automated vehicles is that they make individual travel accessible to groups who are not allowed to drive, either because they are below the age limit (children), have impairments that exclude them from obtaining a driving licence, or have simply opted not to acquire a licence. It is possible to estimate the potential increase in travel enabled by automated vehicles by using statistics on the share of the population who has a driving licence. Such statistics are often collected in household travel surveys. To give an example of how these statistics can be used to assess the gain in access to travel, a Norwegian household travel survey will be used (Hjorthol, Engebretsen and Uteng 2014).

Table 4 shows population by age groups, the share of the population in each group who has a driving licence and the estimated potential increase in travel if 100 % had a driving licence. This indicates the maximum potential increase in travel when fully automated vehicles become accessible to everybody, either on an individual basis or as a form of shared mobility.

Table 4: Potential increase in travel by universal access to fully automated vehicles.

Age group	Population	With licence (%)	Daily travel (km)	Max daily travel (km)
18-24	473,917	75	29.7	39.6
25-34	683,399	90	37.0	41.1
35-44	719,134	95	39.3	41.4
45-54	697,683	97	51.5	53.1
55-66	711,040	95	45.2	47.6
67-74	343,612	93	39.1	42.0
75-+	355,110	83	26.2	31.6
All	3,983,395	91	39.8	43.5

If all individuals in the 18-24 age group had a driving licence, mean travel could increase to $(100/75) \cdot 29.7 = 39.6$ km. This is an absolute maximum, as it assumes that those who do not have a driving licence do not travel at all and thus do not contribute to the current mean length of 29.7 km per day. This is not realistic. It is, however, a simple estimate that can be developed from summary statistics making a minimum of assumptions. To develop more detailed estimates, one would need to have more detailed data and need to make more detailed assumptions, the validity of which cannot be tested today. The very simple approach illustrated above should be widely applicable and indicates an order of magnitude. Given the imprecision of many other estimates of impacts of CATS as these can be made today, such an estimate is judged as adequate.

4.7 Use of shared mobility

There are two types of shared mobility, both of which are expected to increase as a result of vehicle automation. The first type is car sharing, which is found already in many cities. It can be viewed as a form of car rental service in which you book a car as it is needed. The second type is ride sharing. Ride sharing can be viewed as an informal type of public transport. Several individuals share a vehicle, but the vehicle does not run according to a fixed route or fixed time table as current public transport does.

Shared mobility has a great potential for reducing the number of vehicles needed to serve a given volume of travel. A number of the studies reviewed by Soteropoulos et al. (2019) have estimated potential impacts of shared mobility in terms of how few vehicles would be needed to serve current travel volume. In all these studies, it was assumed that integration of automated vehicles with information technology available on smartphones would enable travelers to book trips with a minimum of waiting time. Furthermore, all studies relied on algorithms for optimising the utilisation of shared vehicles, i. e. routines designed to maximise both the service time of the vehicles (the share of each day they were moving; currently around 1 hour out of 24 for an individually owned car, but possibly as much as 16 out of 24 hours for a shared vehicle with the same service hours as, for example, scheduled buses) and their occupancy rate (cabin factor in airline terminology). The results of studies reviewed by Soteropoulos et al. (2019) are presented in Figure 13.

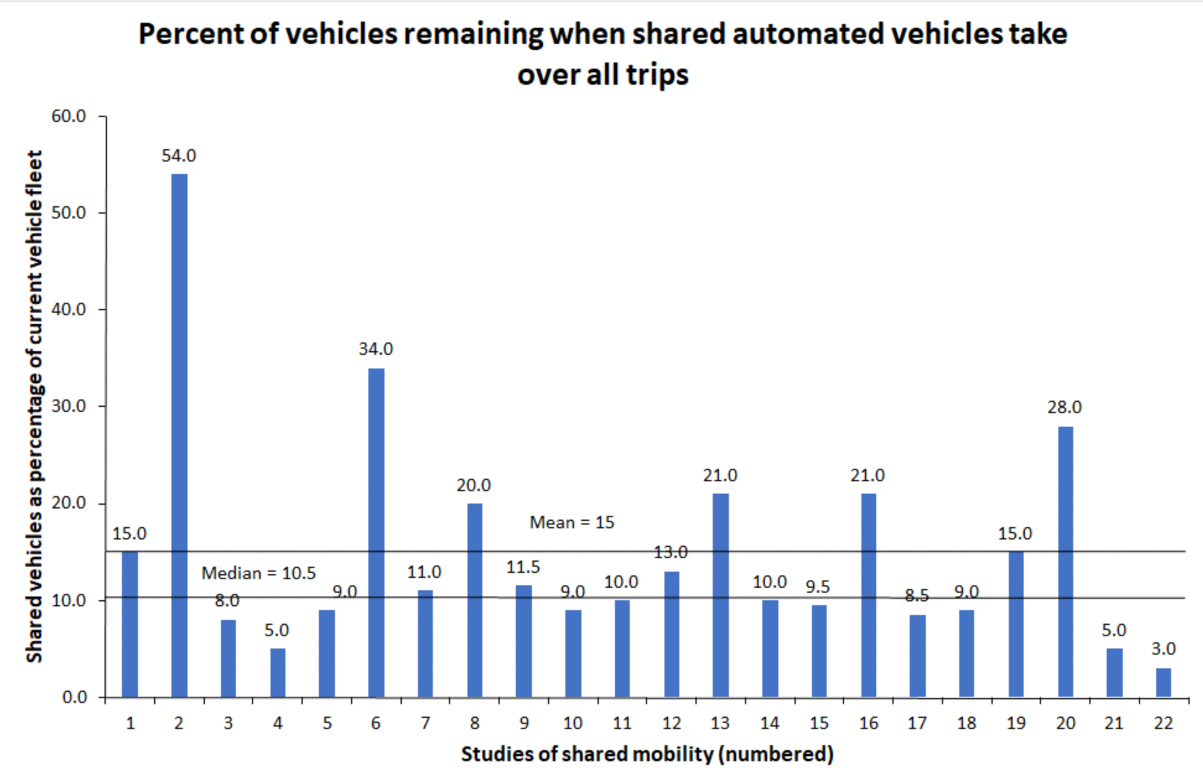

Figure 13: Number of shared vehicles needed to serve current travel volume as a percentage of current vehicle fleet

The results vary, but most studies suggest that a reduction of the vehicle fleet by 90% is possible. This illustrates the enormously inefficient use of individually owned cars today. They are driven mostly on short trips with a single occupant or perhaps a driver and a passenger. They remain parked for 23 out of 24 hours. The studies of shared mobility show that, in theory, one could reduce the number of vehicles by 85-90% and still serve all currently made trips.

It would obviously require very efficient logistics and advanced information systems to accomplish the very efficient use of vehicles modelled in these studies. Another uncertain factor is whether people will prefer shared mobility to individual mobility. Clearly, you can save lots of money by not owning a car, but Grush and Niles (2018) argue that individual ownership will remain widespread even when vehicles have become fully automated.

When predicting impacts of connected and automated cars, it may ultimately not matter much whether these cars are assumed to be owned and operated individually as cars are today or as shared vehicles. Why not? The principal causal pathway of all impacts of connected and automated vehicles discussed so far is traffic volume. It is changes in traffic volume that determines net effects on safety, mobility and environmental factors.

If shared mobility is as successful as some studies indicate, the number of vehicles will be greatly reduced, but traffic volume may not be reduced. Rather than having 100 cars being in traffic 1 hour per day, we may have 10 cars running in traffic 10 hours per day. It is the product of the number of vehicles and the distance (and time) they travel that generates traffic volume. Shared vehicles can only serve the current number of trips by being in operation a substantially greater part of each day than an average car. To serve peak demand by a smaller number of cars, extensive ride sharing would be needed. Some of the studies reviewed by Soteropoulos et al. (2019) have estimated both changes in the number of vehicles and changes in vehicle kilometres. In studies stating changes as intervals, the midpoint value was used. The results of these studies are summarised in Figure 14.

It is seen that even in studies assuming a reduction of the vehicle fleet of around 90%, vehicle kilometres of travel are expected to increase by some 10-20%. In studies assuming a reduction of the vehicle fleet of around 80%, predicted increase in vehicle kilometres of travel ranges from 11 to 67%, on the average 35%. Thus, if it is assumed that increased road capacity and shorter travel time will induce more traffic (see the discussion of secondary impacts in Chapter 5), such an assumption would seem to be equally valid both for the case in which automated mobility remains predominantly individual and for the case in which shared mobility becomes widespread.

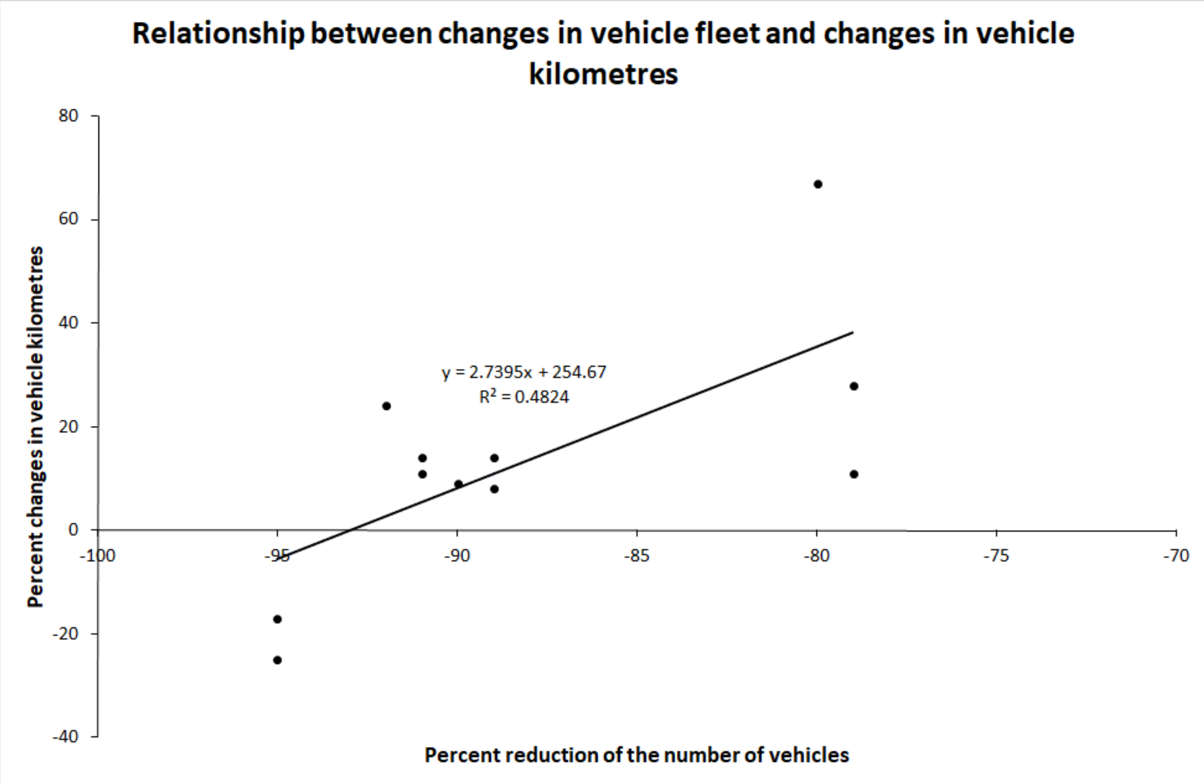

Figure 14: Changes in the number of vehicles and changes in vehicle kilometres

4.8 Road safety

Connected and automated vehicles are expected to improve road safety. Simulation studies support this expectation, but the realism of their results should be assessed by reference to estimates derived by means of other methods.

4.8.1 Rear-end and lane-change accidents on motorways

Studies made by Kockelman et al. (2016), Olia et al. (2016), Li et al. (2017), Rahman et al. (2018), Yang et al. (2018), Papadoulis et al. (2019) and Rahman et al. (2019) have used traffic simulation to estimate the effects of CATS-technology on traffic conflicts involving potential rear-end and lane change collisions on motorways.

The predicted number of traffic conflicts at 0% market penetration was set to 100 for each study; the estimated number of conflicts at higher market penetration levels show percentage changes compare to 0% market penetration. Figure 15 shows the results.

It is seen that all studies predict a reduction of the number of traffic conflicts. The estimates nevertheless vary considerably. A function has been fitted to the data points. The best fitting function was a second degree polynomial. The function passes through or near to the median data points at penetration rates of 10%, 25%, 50% and 100%. Despite the wide dispersion of data points around the fitted curve, it is an unbiased summary of the main tendency of the studies. Note that at 100% market penetration, most of the data points are located near the bottom of the diagram, indicating a reduction in conflicts of 90-95%. The data points located further up in the diagram are outlying.

The dose-response curve in Figure 15, and similar curves presented below, applies when all else is equal. This means that it applies to current traffic volume. If vehicle automation generates more traffic, one should adjust for the effect of this on the number of accidents by multiplying the dose-response curve by an adjustment factor accounting for the effect of increased traffic volume on the number of accidents. This may possibly result in a curve with a different shape from the one shown in Figure 15 and subsequent figures.

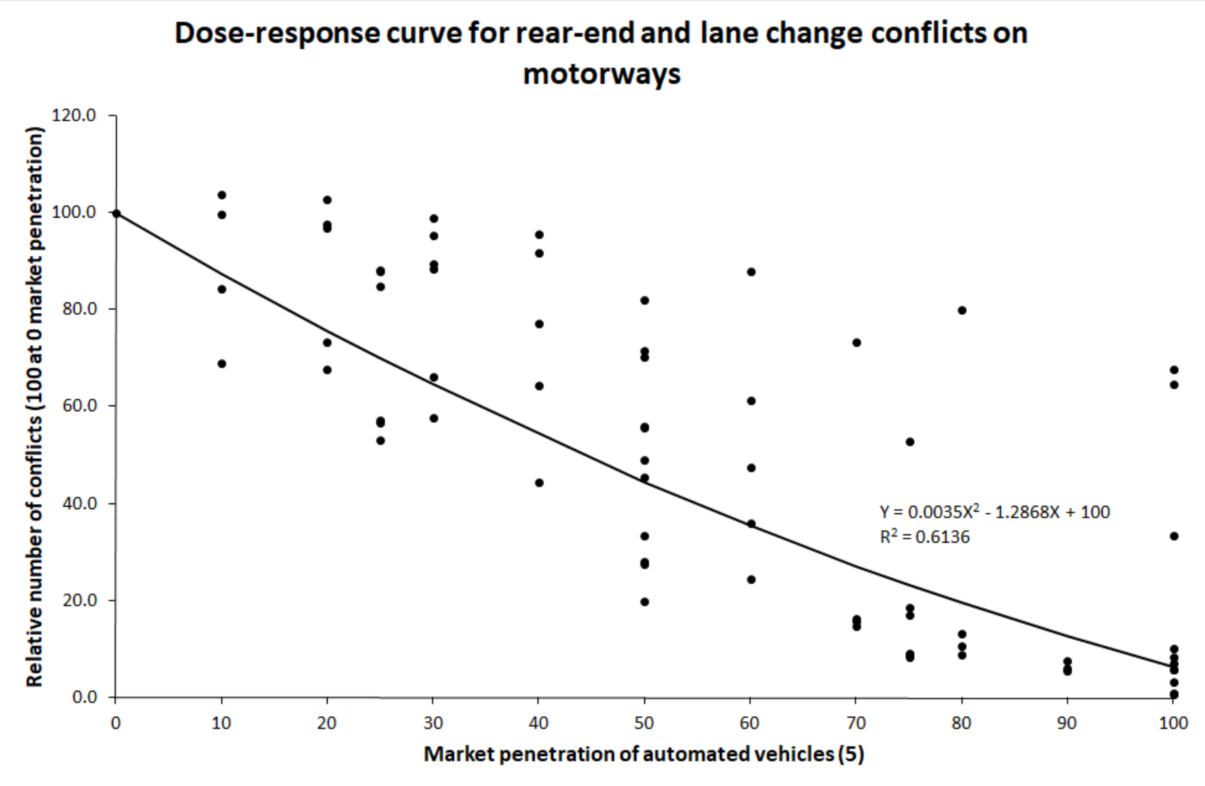

Figure 15: Dose-response curve for rear end and lane change collisions as indicated by the number of traffic conflicts estimated in simulation

4.8.2 Accidents in signalised junctions

Three studies have applied traffic simulation to estimate potential safety effects of connected and automated vehicles in signalised junctions (Kockelman et al. 2016, Morando et al. 2018, Virdi et al. 2019). The results of the three studies differ substantially. In Figure 16, separate dose-response curves have therefore been fitted to the results of the three studies.

It is seen that Kockelman et al. (2016) predict a large reduction of the number of conflicts. However, the study provides only three data points. Morando et al. (2018)

predict a smaller reduction of the number of conflicts. Virdi et al. (2019) predict an increase in the number of conflicts at low market penetration rates for connected and automated vehicles, but then a decline, ending close to the decline at 100% market penetration predicted by Morando et al. (2018).

It would require a detailed examination of the assumptions made in each study to explain why their findings differ so much. The studies agree that at high levels of market penetration, connected and automated vehicles will reduce the number of conflicts in signalised junctions. The curve fitted to the study by Morando et al. (2018) is, for most penetration rates, located between the curves fitted to the studies by Virdi et al. (2019) and Kockelman et al. (2016). This curve can thus be regarded as a mid-point between the two more extreme curves.

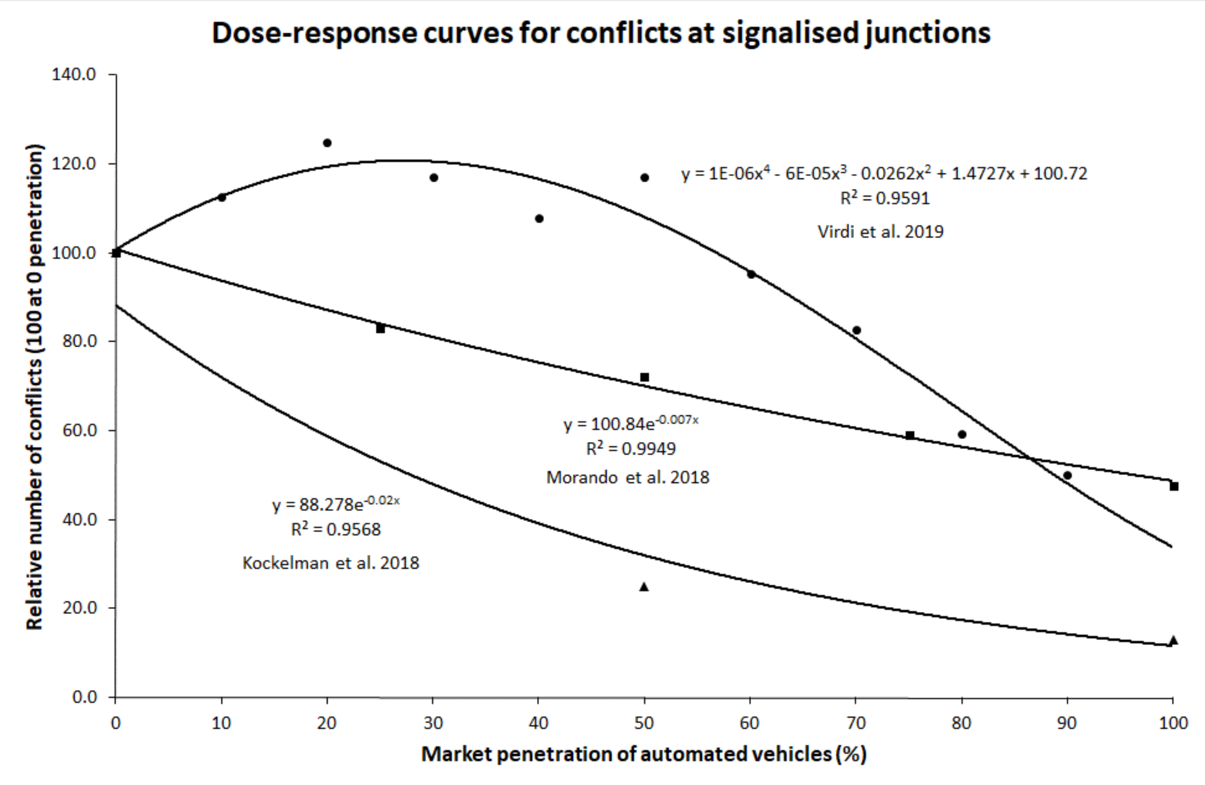

Figure 16: Dose-response curves for accidents in signalised junctions as indicated by the number of traffic conflicts estimated in simulation

4.8.3 Accidents in priority or stop controlled junctions

Kockelman et al. (2018) applied traffic simulation to estimate the number of conflicts in two junctions controlled by stop signs as a function of market penetration of automated vehicles. Estimates were presented for 0, 50 and 100% market penetration. The results for the two junctions differed somewhat. With only three data points, a second degree polynomial fitted the data points perfectly. Virdi et al. (2019) also applied traffic simulation to estimate changes in the number of conflicts as a result of increasing market penetration of automated vehicles. The function fitted to their results is very different from the functions fitted to the results of Kockelman et al. (2018). Figure 17 compares the functions.

When the results of different studies vary as much as in this case, at least three questions arise: (1) Why do the studies reach so different results? (2) Does it make sense to estimate average results based on the studies? (3) If estimating an average does not make sense, should one of the studies be preferred?

To answer the first question, a detailed examination of the assumptions made in traffic simulation would be needed. Such an examination might not necessarily resolve the inconsistency, as different choices of assumptions could all be justified.

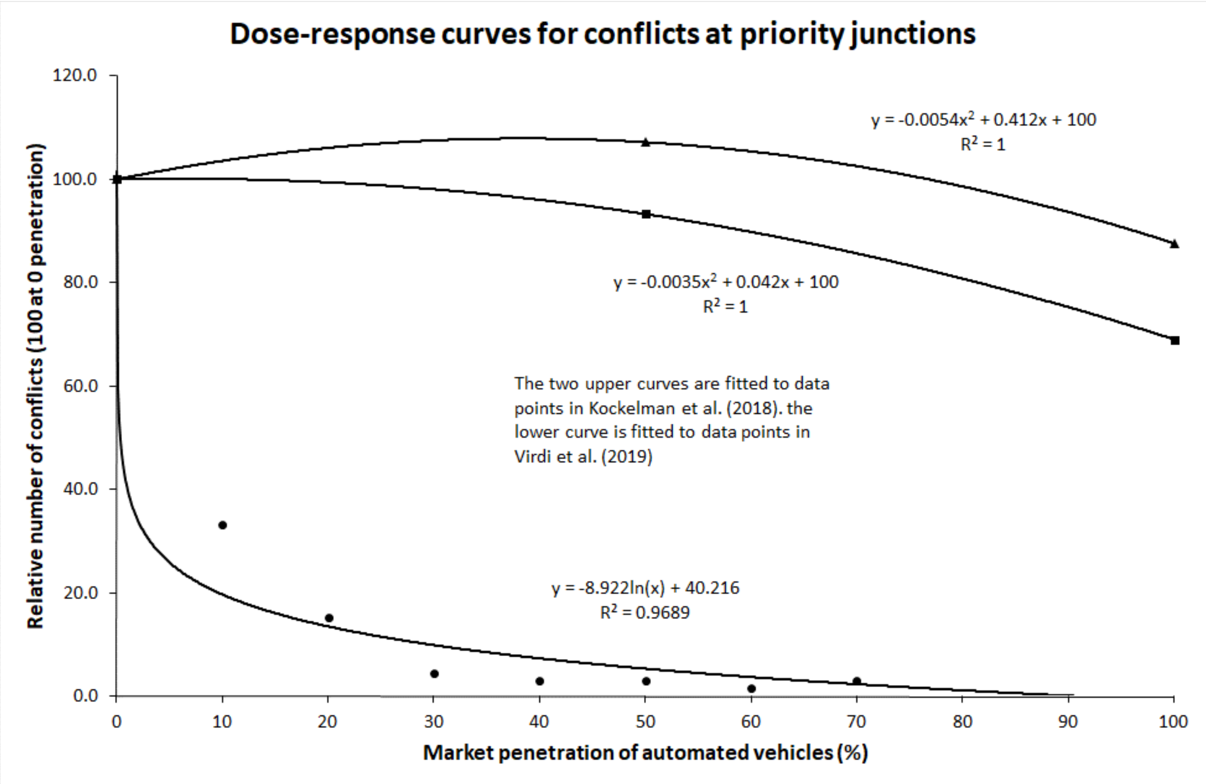

Figure 17: Dose-response curve for accidents in priority or stop controlled junctions as indicated by the number of traffic conflicts estimated in simulation

Virdi et al. (2019) write: "The approach used in this study in applying highly refined CAV operations, such as precise vehicle co-ordination, minimal headways and complete co-operation may be optimistic. … CAV operation may not match the assumptions made in this study."

In other words, the assumptions made by Virdi et al. (2019) may be too optimistic. It has long been known that interaction between motorists in priority-controlled junctions is strongly influenced by informal norms (Ebbesen and Haney 1973). If traffic on the priority road is highly congested, moving at walking pace, cars entering from minor roads will often be permitted to "force" their way into the priority road, as those who are stuck in the queue on that road will understand that if cars without the right of way are not let in, they may have to wait unacceptably long before there is a sufficient gap in traffic to enter the priority road. Even in less congested conditions, drivers on a priority road will let drivers from the minor road enter the priority road by slightly slowing down if they notice that there is a long line of cars waiting to enter from the minor road. It is challenging to make automated cars understand such informal conventions, and during the transition period when there is a mix of human operated and automated cars, misunderstandings may occur if human drivers assume that the automated cars will behave as a human driver would do. Uncertainty about the capabilities of technology to adopt informal conventions suggests that a conservative interpretation of available studies should be adopted. In a case of very different results, as shown in Figure 17, the most conservative estimates will be preferred. No attempt will be made to synthesise the findings of studies providing as different estimates of the impacts of connected and automated vehicles as shown in Figure 17.

The two curves fitted to the results of Kockelman et al. have been synthesised by means of the following polynomial:

Response = $-0.001\text{Dose}^2 - 0.12\text{Dose} + 100$

With Dose stated as percentage market penetration and the fitted value of the Response at 100 % market penetration being 78.

4.8.4 Accidents in roundabouts

Morando et al. (2018) and Virdi et al. (2019) used simulation to estimate the number of traffic conflicts in roundabouts as a function of the market penetration of automated vehicles. Their results are presented in Figure 18.

The functions fitted for the two studies are very different and it makes little sense to estimate an average of them. The question is therefore once more which of the two functions is regarded as most realistic. Morando et al. (2018) specify nine parameters that were used in simulation These parameters were specified for human drivers and two versions of automated vehicles. The curve fitted in Figure 18 is based on version 1 of the automated vehicles, which can be regarded as somewhat more “aggressive” in their behaviour than version 2 of automated vehicles.

Virdi et al. (2019) use the same parameter values for human drivers as Morando et al. (2018). Both studies base these parameters on the Wiedemann (1974) car following model. They do not list the parameters used for connected and automated vehicles, but state that the VCCP algorithm (Virdi CAV Control Protocol) was used to control the behaviour of connected and automated vehicles with respect to acceleration, lane changing and gap acceptance. It was further assumed that connected and automated vehicles would have access to information about all nearby vehicles, and would thus, as an example, only change lane when a sufficient gap was available.

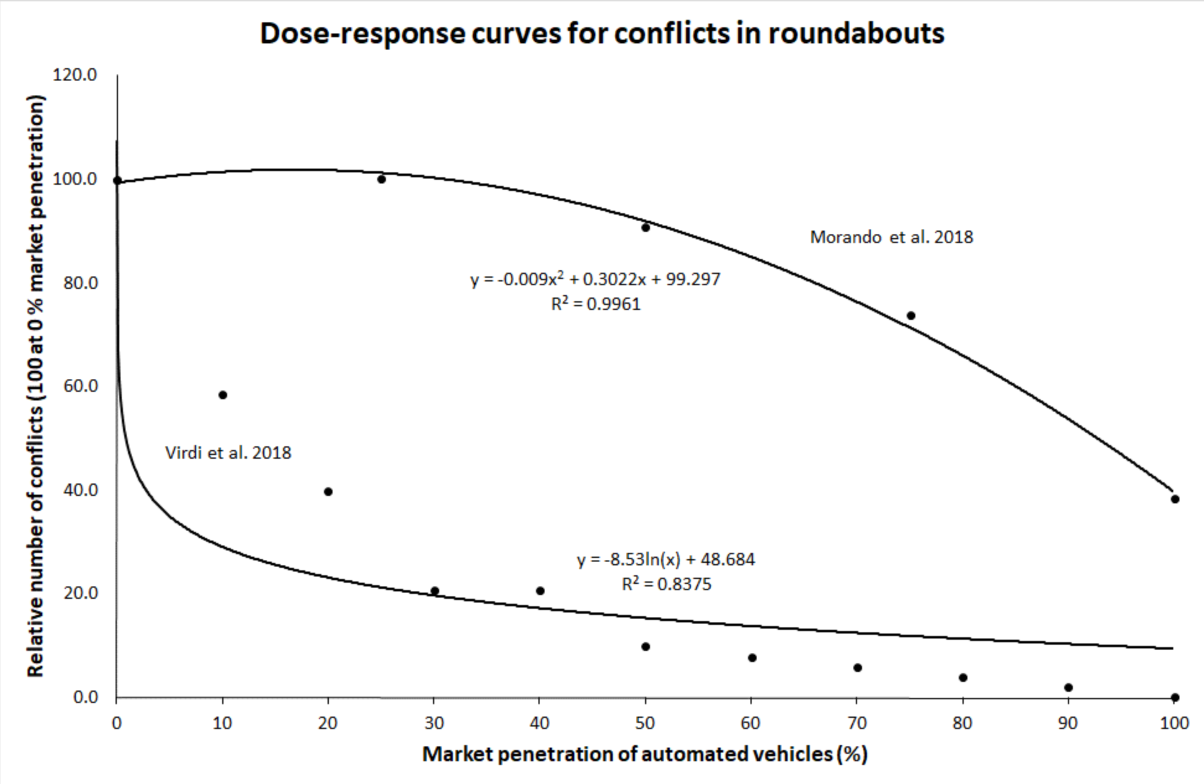

Figure 18: Dose-response curve for accidents in roundabouts as indicated by the number of traffic conflicts estimated in simulation

Given the assumptions made, both studies are probably "correct" in their findings, in the sense that these findings follow mathematically from the assumptions made. Yet, even in roundabouts there will be interactions between automated vehicles and non-automated road users, which cannot fully be described in terms of fixed parameters. Caution therefore once more suggests conservatism and relying on the function associated with the smallest safety benefits.

4.8.5 Cyclist and pedestrian accidents

Two studies have been found that have tried to estimate potential impacts of new vehicle technology on cyclist or pedestrian accidents (Combs et al. 2018; Kovaceva et al. 2019). Both studies used accident data as a basis for evaluating potential impacts of automation technology design to detect pedestrians or cyclists and apply automatic braking to avoid or reduce the severity of a collision. A dose-response curve fitted to the results of the study is shown in Figure 19.

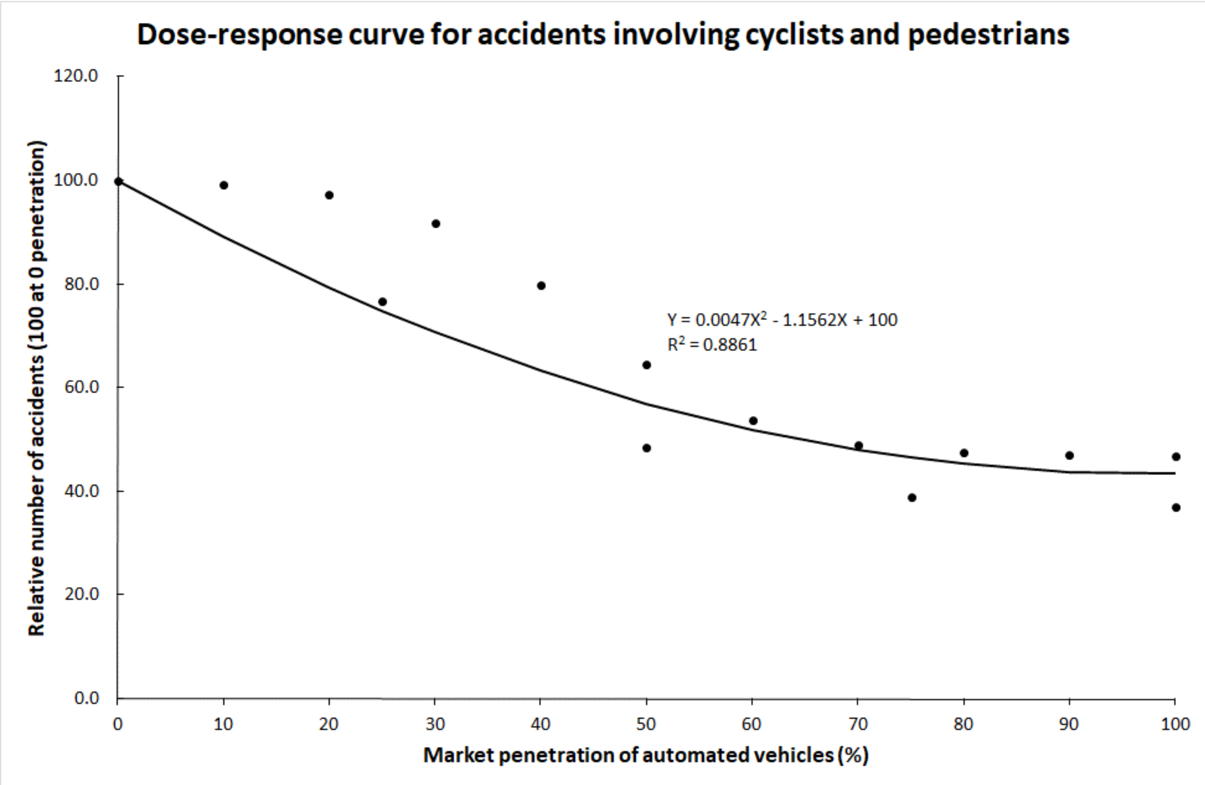

Figure 19: Dose-response curve for pedestrian accidents using Radar and Lidar as indicators of new technology

To fit the dose-response curve, various combinations of Radar and Lidar studied by Combs et al. (2018) were judgmentally converted to the market penetration scale used in the other dose-response curves. The potential for reducing cyclist and pedestrian accidents when all vehicles have detection technology and emergency braking systems was 57%.

4.8.6 Summary of estimated potential impacts on safety

The maximum potential impacts on road safety of connected and automated vehicles according to the dose-response curves in Figures 15-19 are:

Type of accident	Potential reduction (%)
Lane change and rearend on motorways	-94
Accidents in signalised junctions (conservative estimate)	-52
Accidents in priority or stop controlled junctions	-22
Accidents in roundabouts	-61
Cyclist and pedestrian accidents	-57

Estimated maximum potential reductions are seen to vary considerably between the different types of accidents. The estimates currently found in the literature do not include all types of accidents. No estimates have been found for head-on accidents or single vehicle accidents. Coverage of different types of traffic environment is also incomplete. The simulation studies refer, broadly speaking, either to motorways or to urban traffic environments. No studies refer explicitly to a rural traffic environment. This is a significant omission, as most vehicle kilometres in most countries are driven on rural, undivided two-lane roads. These roads differ from urban roads by having fewer junctions and fewer pedestrians and cyclists.

The incompleteness of the dose-response curves, as well as the fact that some of these curves are based on a single study, means that one cannot base methods for predicting impacts of connected and automated vehicles on these dose-response curves only. To develop a library of dose-response curves covering the needs of stakeholders, it is necessary to supplement the dose-response curves presented above with additional curves, obtained either by additional simulation studies, by relying on other approaches for estimating potential impacts of connected and automated vehicles, or, as a last resort, by generalising the dose-response curves presented above to types of accidents or types of traffic environment not explicitly covered by them.

The next section discusses other approaches to the estimation of the potential safety impacts of connected and automated vehicles.

4.8.7 Combining simulation with other approaches

The relevance and applicability of the following approaches for estimating potential impacts on safety of connected and automated vehicles will be discussed:

1. Estimates based on findings of in-depth studies of accidents (The in-depth approach)
2. Estimates based on epidemiological studies of the contributions of risk factors to accident (The epidemiologic approach)
3. Estimates based on evaluations of specific technologies already on the market or close market launching (The technology extrapolation approach)
4. Estimates based on comparative analysis of the reliability of human operators and automated systems (The comparative reliability approach).
5. Estimates based on comparing the actual accident rates of automated vehicles to the accident rates of manual vehicles (The accident rate comparison approach)

4.8.7.1 The in-depth accident investigation approach

In-depth studies of accidents typically identify human error as a major contributor to accidents. Examples include lack of observation (e. g. caused by distractions), misunderstandings of situations (erroneously thinking that another road user would give way), and misjudgements (erroneously believing stopping distance was adequate). Some kind of human error is usually classified as having contributed to about 90% of the accidents. Based on this estimate, it is tempting to conclude that by eliminating human drivers, one may eliminate all accidents to which human error contributed. For an example in the literature, see Fagnant and Kockelman (2015).

While there is little doubt that various forms of human error make an important contribution to accidents, assessing potential changes in safety associated with the gradual elimination of these errors is impossible. Driver error will cease to be an issue only when all vehicles are fully automated and drivers eliminated. However, in the transition phase, it is by no means clear what shape an imagined gradual "error-elimination" curve (as a function of market penetration of automated vehicles) might have. Automation at a certain level, particularly SAE level 3, may induce new errors. Platooning of a certain percentage of vehicles, involving shorter headways than human drivers currently adopt, may affect the headway chosen by cars still operated manually (see discussion in Chapter 5). Little is known about such behavioural adaptations. Therefore, it is impossible to derive a dose-response curve merely on the basis of knowing that human error contributes to a high percentage of accidents.

4.8.7.2 The epidemiologic approach

Epidemiological studies aim to quantify the contribution of various risk factors to accidents and injuries, and the potential reduction of the number of accidents or injuries that can be attained by eliminating the increase in risk associated with a risk factor. The potential for reducing accidents or injuries is generally stated as attributable risk, which is defined as follows:

Attributable risk = $\dfrac{PE(RR-1)}{(PE(RR-1))+1}$

PE is the proportion of exposure (vehicle kilometres) exposed to a certain risk factor. RR is the relative risk associated with a risk factor. Thus, if 30% of traffic is in darkness and the risk is 50% higher than in daylight, PE = 0.30 and RR = 1.50. The risk attributable to darkness is 0.13 (0.3 · 0.5)/(1 + (0.3 · 0.5)). Thus, if the risk associated with darkness could be eliminated, the total number of accidents could be reduced by 13%. Estimates of attributable risk show potential accident reductions. In practice, it is rarely possible to fully eliminate a risk factor. Accident reduction will in most cases be smaller than an estimate of attributable risk suggests.

In a study of the prospects for improving road safety in Sweden, Elvik and Amundsen (2000) estimated attributable risk for about 20 risk factors. The results, taken from Elvik (2012) are reproduced in Table 5. The table shows first-order attributable risks, i.e. the attributable risk estimated for each risk factor without considering its correlation or interaction with other risk factors. It is seen that risk factors in each of the main groups identified all make a substantial contribution to fatalities and injuries.

To give an example of the interpretation of the estimates, consider poor car crashworthiness. The attributable risk for fatalities is estimated at 0.341. This means that if all cars had the best crashworthiness performance of cars in their weight class, the number of fatalities could be reduced by 34%.

Attributable risks sum to more than 1, in fact to 2.36 for fatalities and 1.56 for injuries. This does obviously not mean that one can eliminate fatalities two times over. It shows, basically, the same as in-depth studies, which is that more than one factor contributes to each accident.

It is reasonable to expect fully automated and autonomous vehicles to eliminate or greatly reduce most of the attributable risks listed in Table 5. Age and other characteristics of the driver – which are major risk factors today – would no longer matter if the vehicle is driverless. They might continue to matter if automation was partial only and the driver would be asked to take control in certain situations.

Table 5: First order attributable risk associated with various risk factors in Sweden. Based on Elvik and Amundsen (2000); reproduced from Elvik (2012).

Risk factor	Reference condition	First-order attributable risk	
		Fatal injuries	All injuries
A: System design			
A1: Travel in an urban traffic environment	Travel in a rural traffic environment	–0.131	0.214
A2: Substandard roads	Roads with good standard	0.034	0.010
A3: Unprotected roadside obstacles	Clear roadsides	0.167	0.039
A4: Erroneous highway signs	Correct highway signs	0.010	0.015
A5: High-risk junctions	Low-risk junctions	0.036	0.022
A6: Poor car crashworthiness	Best car crashworthiness in each weight class	0.341	0.066
A7: Involvement of heavy vehicles in crashes	No involvement of heavy vehicles in crashes	0.106	0.005
A. Total		**0.563**	**0.371**
B: Environmental Factors			
B1: Darkness	Daylight	0.165	0.107
B2: Winter and wet roads	Summer and dry roads	0.196	0.146
B3: Animals	No animals	0.015	0.038
B. Total		**0.376**	**0.291**
C: Vulnerability of road users			
C1: Unprotected road users	Protected road users (car occupants)	0.320	0.266
C2: Children	Adults	0.016	0.022
C3: Young drivers	Middle-aged drivers (safest age groups)	0.086	0.101
C4: Older road users	Middle-aged road users (safest age groups)	0.206	0.068
C. Total		**0.628**	**0.457**
D: Unsafe road user behavior			
D1: Speeding	Keeping speed limits	0.376	0.210
D2: Drinking and driving	Driver sober	0.072	0.048
D3: Not wearing seat belts	Wearing seat belts	0.084	0.032
D4: Other violations	No violations	0.092	0.061
D5: Excessive driving in towns	No excessive driving in towns	0.005	0.015
D. Total		**0.629**	**0.366**
E: Rescue services			
E1: Not state-of-the-art rescue service	State-of-the-art rescue services	0.167	0.071
E. Total		**0.167**	**0.071**
F: Total of all risk factors		**2.363**	**1.556**

Unfortunately, it is not easy to estimate the combined effects of eliminating all or most of the risk factors listed in Table 5. The principal problem is that many of the risk factors are highly correlated. Eliminating the contribution of one of a pair of highly correlated risk factors would then go some way towards eliminating the contribution of the other. Unless correlations among the risk factors are accounted for, there may be serious double counting of potential benefits of eliminating the risk factors. At this point, limits of current knowledge impose a serious restriction on the epidemiological approach. Elvik and Amundsen (2000) tried to account for correlations among the risk factors. Without going into detail, adjusting for correlations reduced the sum of attributable risk from 2.56 to 1.96 for fatalities and from 1.56 to 1.23 for injuries.

The combined effect of eliminating all attributable risks, adjusted for overlap and correlations between them, was estimated by using the method of common residuals. The basic logic of this method can be explained as follows:

Suppose the risk attributable to factor A is 0.2 and the risk attributable to factor B is 0.1. The combined effect of eliminating the risk factors in then not 0.3 – the sum of their attributable risks. If risk factor A is eliminated, it leaves a residual of 0.8 (1 – 0.2). Eliminating risk factor B can only act on the residual, i.e. it reduces risk by 0.1 of 0.8, or 0.08. The combined contribution is then 0.28.

By applying the method of common residuals, one ensures that the total reduction of fatalities and injuries resulting from the combined elimination of a set of attributable risks can never exceed the value of 1. It is mathematically impossible as indeed it should be. Elvik and Amundsen (2000) estimated that by eliminating all risk factors listed in Table 6, fatalities could be reduced by 89 % and injuries by 73 %. Elvik (2009) subsequently developed a more conservative version of the method of common residuals, the dominant common residuals method, which tries to account for correlations among risk factors.

The epidemiological approach is flexible, as it can be used both to estimate potential effects of full automation and potential effects of partial automation, depending on the risk factors included in the estimation. Ideally speaking, estimates of attributable risk should be customised to each use case. This will almost surely not be possible for all risk factors listed in Table 5. Another limitation is that the shape of any gradual risk-factor elimination curve is unknown, i.e. estimates of the potential gain in safety by eliminating or reducing a certain combination of risk factors cannot easily be converted to a dose-response curve having market penetration of automated vehicles as the independent variable. Risks are non-linear and interact. Thus, if exposure to a risk factor is reduced by 50 %, it does not follow that the risk attributable to that factor is reduced by the same percentage. An inventory of attributable risks therefore cannot be converted to a dose-response curve showing impacts of connected and automated vehicles, but it can serve as a benchmark for assessing the plausibility of potential impacts estimated by simulation.

4.8.7.3 The technology extrapolation approach

Vehicle automation is widely thought of as a gradual process. First, advanced driver support is introduced. Next, low level automation is introduced for selected driving functions or defined domains of operation. Then, domains of operations are extended and levels of automation increased. Finally, fully automated vehicles are developed that can operate on any road.

Suppose that, as one benchmark, one has estimates of the potential safety effects of fully automated vehicles at the highest level of automation. This benchmark is the upper limit of attainable safety impacts (the endpoint of a dose-response curve). Suppose further, that estimates of safety impacts of non-automated driver support systems are available. These estimates show the maximum safety gains that can be attained without automation, or the lower bound of the number of accidents that automation can influence.

Thus, if, for a given type of accident full automation (SAE level 5) has the potential to reduce the number of accidents by 80% and advanced driver support systems has the potential to reduce accidents by 40%, the progression from lower to higher levels of automation can be assumed to have impacts between 40% and 80% accident reduction.

Advanced driver support systems can reduce the number of accidents considerably. A recent study by Leslie et al. (2019) reported that camera-based forward collision warning and automatic emergency braking reduced rear-end accidents (as striking vehicle) by 46%. Front pedestrian braking reduced pedestrian accidents by 13%. Lane departure warning combined with lane keep assistance reduced lane departure accidents by 20%. Blind zone alert, combined with lane change alert and rear camera mirror reduced lane change accidents by 37%. Various systems supporting safe backing, reduced backing accidents by 81%. Intelligent headlights reduced accidents involving animals at night by 35%. The study clearly shows that accident reductions increase as more driver support systems are combined.

Wang et al. (2020) made a similar study. They reported a meta-analysis of studies evaluating the safety effects of autonomous cruise control, automatic emergency braking, blind spot warning, electronic stability control, forward collision warning, lane change warning, lane departure warning, pedestrian collision mitigation and intersection movement assistant. Except for autonomous cruise control and automatic emergency braking, these are all advanced driver support systems that do not remove the driver from the driving task, but supports it by providing information or issuing warnings. The combined effect of all systems was estimated for six countries, using official collision statistics. Combined effects varied between 41% and 54% accident reduction, with a simple average of 47%. This shows that driver support systems currently on the market can reduce the number of accidents almost by half even before the transition to full automation starts.

However, it is unknown if these advanced driver support systems reduce the potential additional effects of vehicle automation. Clearly, for some driver support systems, it is reasonable to assume that they reduce the effects of automation. Thus, if mandatory intelligent speed adaptation (ISA) and alcolhol ignition interlock were to be installed in all driver operated vehicles, automation would no longer bring the benefit of reducing accidents attributable to speeding or drinking and driving. Thus, it is important to be clear about which driver support systems are assumed to be implemented in the baseline situation, before automation starts.

All the systems evaluated by Leslie et al. (2019) are already implemented and will most likely reach full market penetration before automation starts. ISA and alcohol ignition interlocks, on the other hand, may not be widely implemented before automation starts to spread. In a Norwegian Delphi survey (Høye, Hesjevoll and Vaa 2015), respondents expected 60% of new cars in 2030 to have automatic lane keeping (i. e. this specific function would be automated, but the car would still be operated by a driver), but only 12% of new cars in 2030 to have alcohol ignition interlock and 6% of new cars to have ISA.

It is likely that continued introduction and improvement of advanced driver support system will reduce the potential safety effects of vehicle automation, meaning that any safety benefits of automation would start to develop from a lower baseline number of accidents than the current number.

4.8.7.4 The comparative reliability approach

Humans are very reliable as operators of motor vehicles; thousands of kilometres are driven between each accident. To improve safety, automation technology must operate at a higher level of reliability than human drivers. Potential safety impacts of automation technology may be assessed by comparing its reliability to the reliability of human drivers.

To take a very simple example: Most of the time, drivers are able to adapt their speed to the vehicle in front of them. Most cases of braking do not end in a rear-end collision. While the exact "success rate" in avoiding rear-end collisions is unknown, let us assume it is 999,999 out of 1,000,000. This rate of reliability is the minimum a technical replacement of a human must attain in order to improve safety. It such a level of reliability possible?

Probably the most detailed and representative data that can shed light on human reliability are data from naturalistic driving studies. Dingus et al. (2016) provide data based on the SHRP 2 naturalistic driving study. As of April 2015 (which was probably when the paper was written) these data represented more than 56 million kilometres of driving, during which more than 1,500 accidents had happened. (This number is mentioned in the text, but the detailed study reported was based on 905 accidents; it is not entirely clear whether these 905 accidents refer to the same time period as the 56 million kilometres driven; below, that will be assumed). Since all events are continuously monitored, it is reasonable to assume that accident data are complete. Mean accident rate was 26.8 per million kilometres driven, corresponding to a reliability level per kilometre driven of 0.9999732 (i.e. the probability that a randomly selected kilometre was accident free).

Human reliability varies considerably. The figure given above is the average for all drivers who took part in the SHRP 2 study. The reliability of automation technology is currently not known. Hence, quantifying it and comparing it to human reliability is not feasible and no dose-response curve can be derived from such a comparison.

4.8.7.5 The accident rate comparison approach

Noy, Shinar and Horrey (2018) mention studies that have found automated cars to have higher accident rates (accidents per million kilometres driven) than manually driven cars. These studies are based on very few accidents and must therefore be regarded as inconclusive. A considerable amount of data is needed in order to reliably estimate any difference in safety between manual and automated cars (Kalra and Paddock 2016).

The state of California requires all incidents involving automated cars to be reported. This is a requirement in the law permitting automated cars to operate on public roads, mixing with non-automated road users. Based on these data, Favaro et al. (2017) have estimated the risk of accident involvement for automated vehicles. Data collected from September 2014 to March 2017 include 26 accidents and 1 088 453 miles travelled (1 751 321 kilometres). Favaro et al. estimate an accident rate of 23.89 per million vehicles miles of travel. They indicate that the mean accident rate for driver operated vehicles in the United States is 2.0 per million vehicle miles of travel, suggesting that the accident rate for automated vehicles is more than ten times higher than the accident rate for manual vehicles.

This comparison is meaningless. The only fair basis for a comparison is data from naturalistic driving studies. It is only data from naturalistic driving studies that can reasonably be assumed to have equally complete coverage of accidents as data for automated vehicles. Based on the SHRP 2 naturalistic driving study, the mean accident rate for manual vehicles can be estimated as 16.07 accidents per million vehicle kilometres. The corresponding rate for automated vehicles is 14.85 accidents per million vehicle kilometres, or 7.6 % lower than driver operated vehicles. Automated vehicles are already a little safer than driver operated vehicles.

This simple comparison overlooks the fact that risk varies, both for driver operated and automated vehicles, depending, among other things, on the distance driven per year. Table 3 in the paper by Favaro et al. (2017) allows four groups of automated vehicles to be formed, differing in terms of annual distance driven. Corresponding data for driver operated cars can be taken from a paper by Antin et al. (2017), which, unfortunately, only includes drivers above 65 years. These drivers had a mean accident rate of 14.74 accidents per million vehicle kilometres; slightly lower than the overall risk for drivers participating in the SHRP 2 naturalistic driving study. Figure 20 shows how accident rate depends on annual distance driven for driver operated and automated cars based on the studies of Favaro et al. and Antin et al.

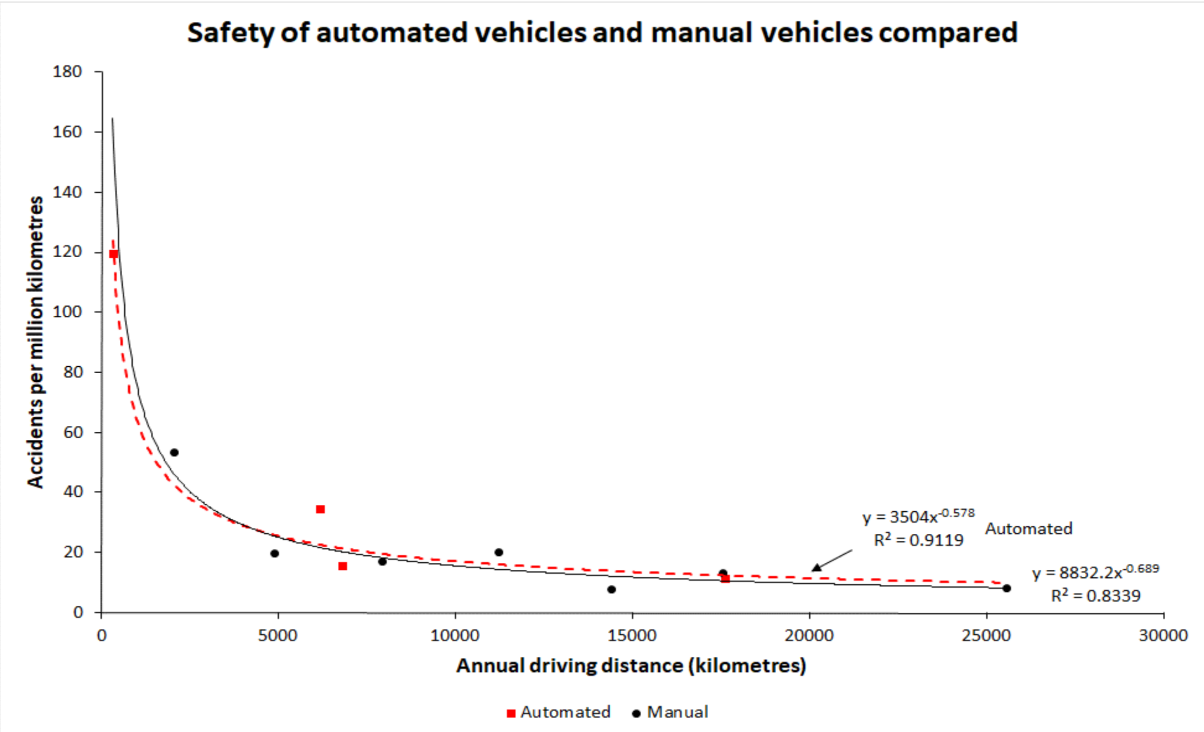

Figure 20: Variation of accident rate by driving distance for driver operated and automated cars. Derived from Favaro et al. (2017) and Antin et al. (2017).

The curves are almost on top of each other. Given the fact that the mean accident rate estimated for the SHRP 2 drivers was 16.07, not 14.74 accidents per million vehicle kilometres, it is likely that the risk curve for automated vehicles would be below that for human drivers if all SHRP 2 drivers had been included, not just those aged 65 and above.

An updated comparison can be made on the basis of data provided by Boggs et al. (2020). They list mileage data and crash data for 36 companies operating automated vehicles on public roads in California. Mean crash rate was 21 per million vehicle kilometres, but crash rate varied considerably between the companies. Most of them had zero crashes. To err on the side of caution, crash rate has only been estimated for companies that had at least 1 crash; the long-term crash rate of the other companies is obviously not zero. Figure 21 shows a comparison between automated cars and human drivers under these assumptions.

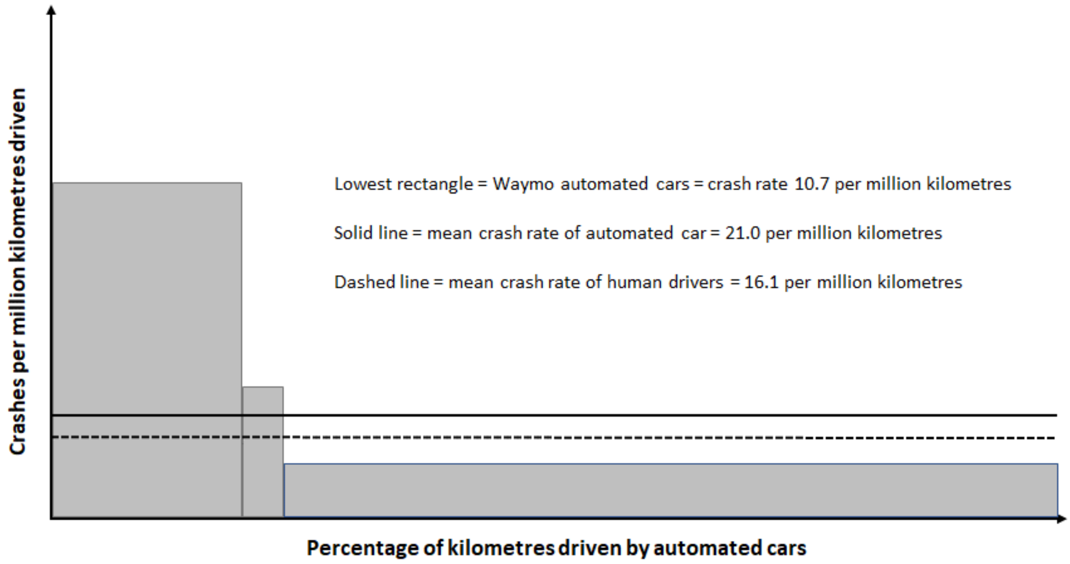

Figure 21: Comparison of crash rates for automated cars and human drivers. Derived from Boggs et al. 2020

The horizontal axis shows the percentage of kilometres driven by automated cars for three groups of cars. One group had a high crash rate of about 67 per million kilometres and represented 18.6% of the total number of kilometres driven. This group is shown by the rectangle to the left. A small group had an average crash rate of 27 per million kilometres and represented 5.2% of the total number ogf kilometres driven. Finally, the Waymo cars of Google had a crash rate of 10.7 per million kilometres and contributed to 76.2% of the total number of kilometres driven. Thus, most of the driving performed by automated cars had a lower crash rate than both a human driver (16.1) and the mean for all companies operating automated cars (21.0).

Sparrow and Howard (2017) argue that once it has been established that automated cars are safer than driver operated cars, (human) driving should be banned. Remarkably, we may already be at that point, although automated vehicles must still be regarded as an immature and experimental technology. In a similar vein, Kalra and Groves (2017) speak about "The enemy of good", by which they refer to the costs to society of waiting for automated vehicles to become perfect, rather than switching quickly to them once they are only a little safer than human drivers.

4.8.7.6 Conclusions about other approaches

As noted above, it is at this time only possible to develop a few dose-response curves showing potential safety impacts of connected and automated vehicles based on simulation. These dose-response functions do not include all types of accident and do not apply to all types of traffic environment. They do therefore not provide a complete basis for predicting the potential impacts on safety of connected and automated vehicles.

In this section, the question was asked whether the dose-response curves based on simulation can be supplemented by dose-response curves estimated by means of other approaches. Unfortunately, the answer is negative. One may at best get a rough estimate of the maximum potential safety effects of automation technology if the technology succeeds in eliminating all known risk factors contributing to accidents. Such an estimate may function as a benchmark, but does not provide a basis for developing dose-response curves. A benchmark indicates the endpoint of a dose-response curve, i.e. the number of accidents remaining when the highest level of automation has reached full market penetration. However, between the starting point and the end point, a dose-response curve can have many shapes, as shown below:

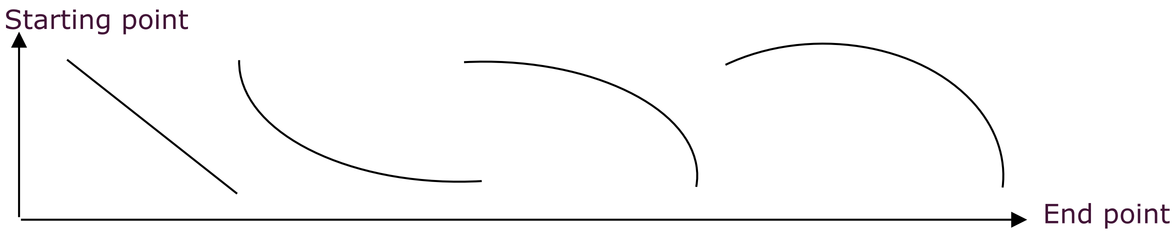

It is important to know the shape of the dose-response curve in order to estimate impacts at all levels of market penetration of automation technology from 0 to 100 %. If only the starting point and the end point are known, the total impacts on safety cannot be estimated. Thus, a curve initially falling steeply and then flattening close to its lowest level will, during the time it takes to traverse the function from start to end, save more lives than a curve which initially increases and starts to fall sharply only when it gets close to full market penetration. The area under the curve falling fast at first will be smaller than the area under the curve first rising and then falling sharply.

Knowing the shape of the dose-response curve has great policy relevance. If the curve is shaped like the rightmost curve above, it may be wise to first introduce a technology with a dose-response curve like the second curve from the left above. This curve falls sharply even at low levels of market penetration, and once it has reached a low level, the technology associated with the rightmost curve above can be introduced, because the initial rise of this curve will start from a lower initial number of accidents than it would if the technology giving quick gains in safety had not been introduced before.

The dose-response curves emerging from the simulation studies quoted above had different shapes. It therefore seems clear that the shape of the dose-response curves can be different for different types of accident, different automation technologies and different traffic environments. Is it possible at all to say anything about the expected shape of the dose-response curves for those types of accident etc. not included in the curves presented above?

4.8.8 Event-based exposure as potential conflicts – implications for dose-response curves

In a series of papers, Elvik et al. (Elvik, Erke and Christensen 2009, Elvik 2010, 2015) have explored some implications of defining traffic exposure as events. Examples of events include: encounters (cars passing each other in opposite directions), simultaneous arrivals at junctions, braking, lane changing. Common to all events is that the event is a potential conflict. The event has two outcomes: no conflict or conflict. A conflict in turn has two outcomes: accident or no accident.

These studies are relevant for assessing the shape of dose-respose curves, as closed-form solutions relating the number of events to the number of conflicts exist for many types of events. To give an example: each encounter is an opportunity for a conflict or head-on collision. The number of encounters is:

$$\text{Number of encounters} = \left(\frac{\text{Number of vehicles in both directions per unit of time}}{2}\right)^2$$

If the assumption is made that automation technology can prevent all head-on collisions, this is analogous to reducing traffic volume from a given initial level to zero. The dose-response curve from market penetration from 0 to 100%, and an initial level of 100 accidents will then be:

Dose-response curve (encounters) = $0.01x^2 - 2x + 100$

This curve has an initial value of 100 and a final value of 0. It falls sharply at the beginning and becomes flatter closer to 100% market penetration. Should the technology be less than 100% effective, for example expected to prevent only 80% of head-on collisions at 100% market penetration, the curve can be shifted so that its initial value is 100 and its final value is 20 rather than 0, while preserving the shape of curve. This theoretically derived dose-response curve can be used for head-on collisions.

Closed-form expressions were also derived for the probability of conflict in three-leg and four-leg junctions as a function of the number of simultaneous arrival from potentially conflicting directions. A simultaneous arrival was defined as within the same 1 second. If two vehicles arrive within the same 1 second, there is clearly a small safety margin and a high probability of conflict. If it is again assumed that automation technology is 100% effective, i.e. it has vehicle-to-vehicle communication systems that can reliably detect other arriving vehicles and take appropriate action to prevent a conflict, then the dose-response curve can again be modelled as a reduction of the number of simultaneously arriving vehicles from a given initial number to zero. The following dose-response curve perfectly fits (for hourly entering volumes between 50 and 2000, all approaches combined):

Dose-response curve (conflicts in junctions) = $0.00003x^3 + 0.004x^2 - 1.6928x + 100$

Within the range between 0 and 100, this polynomial produces meaningful values. It is suggested to use it for conflicts in junctions not covered by the simulation studies, e.g. conflicts between cars and cyclists. Again, should the technology not eliminate conflicts, but reduce them by 80%, the curve can be shifted to a final value of 20 (rather than 0), while preserving its shape.

Closed-form curves also exist for lane changes and braking in a row of cars following each other in the same direction. One type of accident that cannot easily be linked to a specific event is running off the road. No dose-response curve has been found for automation technology intended to help prevent running off the road. Vehicle technology intended to serve such a function may possibly need to be supported by electronically readable road markings that help keep the car within its lane. The dose-response curve will then be a diffusion curve, in which accident reduction will be proportional to the share of vehicle kilometres driven by automated cars having the technology.

4.9 Fuel consumption and emissions

The principal source of uncertainty when trying to estimate the impacts of vehicle automation on fuel consumption and emissions is whether connected and automated vehicles will continue to use fossil fuels or become electric. Energy efficiency is expected to continue to improve even if vehicles are powered by combustion engines. This means that, for a given traffic volume, less energy will be consumed and less emissions produced. If electric vehicles take over, they will still use energy for movement, but this will not be associated with emissions. How much "cleaner" electricity is than fossil fuels, depends on how it is produced. Within the Levitate project, we have chosen not to go into how electricity is produced. The main difference between fossil fueled vehicles and electric vehicles is that the latter are zero emission.

Studies by Wadud et al. (2016) Olia et al. (2016), Vasebi et al. (2018) and Chen et al. (2019) allow a dose-response function to be developed for fuel consumption. These studies all assume that automated vehicles will continue to have internal combustion engines. A curve fitted to the data points provided by these studies is shown in Figure 22.

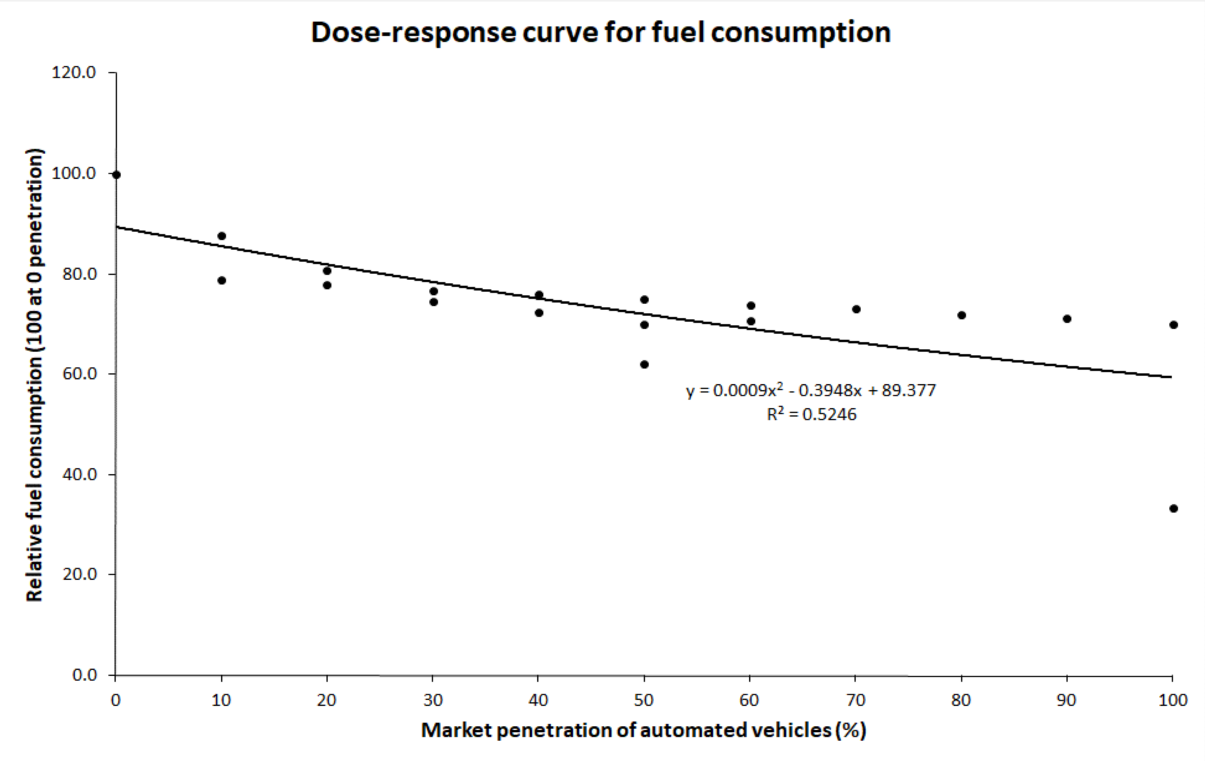

Figure 22: Change in fuel consumption associated with vehicle automation

While the studies produce estimates that are close to each other up to 50% market penetration for automated vehicles, there is less consistency at higher levels of market penetration. Only one study included penetration levels of 70, 80 or 90 percent. Two studies provided estimates for 100 percent market penetration, but these estimates were far apart.

Slowik and Sharpe (2018) summarise studies of the effects of truck platooning in fuel consumption. In studies reporting changes as an interval, the midpoint of the interval is used. In ten studies, the mean saving in fuel consumption was 11.8%, with a range from 2.8% to 22%. Presumably, all estimates refer to savings in the types of fuel (mainly diesel) currently used by trucks. Estimates are interpreted as showing maximum effects, i.e. savings if all trucks are in a platoon (100% penetration of platooning).

4.10 Vehicle ownership costs

Car ownership costs include the costs of buying and keeping a car and the costs of operating it. Studies have been made to determine whether vehicle automation will change these costs. The studies do not easily lend themselves to a dose-response curve, but will be referred to in chronological order.

Fagnant and Kockelman (2015) estimate that an automated vehicle initially will cost 10,000 US dollars more than a conventional vehicle, but that the additional cost will fall to 5,000 US dollars at 50 % market penetration and to 3,000 US dollars when automated vehicles reach 90 % market penetration.

Anderson et al. (2016) estimate that an average US household could save about 6000 US dollars per year by not owning a car and joining a car-sharing program.

Bansal and Kockelman (2017) estimate the cost of level 3 automation to 15,000 US dollars in 2015, dropping to 11,607 dollars in 2020 and 3,220 dollars in 2045. The cost of level 4 automation was estimated as 40,000 US dollars in 2015, reducing to 30,951 dollars in 2020 and 8,586 dollars in 2045.

Wadud (2017) reported a cost of ownership analysis for automated vehicles. He estimated current costs of owning and operating cars by income quintiles (shares of one fifth) in Great Britain. He estimated that buying an automated vehicle in 2020 would cost 16,600 US dollars more than a conventional vehicle, but that cost would fall by 5 % per year. Total costs of ownership were estimated to increase by 29.6% in the lowest income quintile, falling to 16.3%, 7.9% and 2.9% in the next income quintiles. Only the upper fifth income quintile (the top 20% income earners) would save costs. For this group, a reduction of car ownership costs of 6.4% was estimated. For commercial vehicles, cost savings ranging from 15 to 30% were estimated. Thus, automated vehicles can save costs in commercial transport and for the richest private car owners. For 80% of the population, automated vehicles would be more expensive than current vehicles, even allowing for savings attributable to lower fuel consumption and more productive use of travel time. It should be noted that no savings in insurance costs was assumed. Wadud argued that although automated vehicles will have a lower accident rate than conventional vehicles, they will be more expensive to repair or replace, resulting in an unchanged insurance premium.

Bösch et al. (2018) compared the costs of owning and operating conventional cars to automated cars. For private cars, the cost per kilometre of an automated car were found to be 4 % higher than the cost of a conventional car. The main contributor to the increased cost was depreciation, which was higher because the automated car was assumed to be more expensive than the conventional car. Insurance costs were assumed to decline by 50%. The per kilometre cost of a taxi was estimated to be 85% lower for an automated taxi than for a taxi with a driver.

Shabanpour et al. (2018), in a stated preference study, assumed a purchasing price for an automated vehicle of either 40,000, 50,000 or 60,000 US dollars.

Slowik and Sharpe (2018) estimated the additional cost of level 5 automation technology for a truck to be 23,400 US dollars. The technology included sensors, communication systems and processing software.

Tirachini and Antoniou (2020) give cost estimates for fully automated vehicles of 29,490 Euro for a car, 43,433 Euro for a van, 281,234 Euro for a minibus, 419,429 Euro for a standard bus, and 627,696 Euro for an articulated bus. Presumably, all these estimates refer to the extra costs of an automated vehicle compared to an otherwise identical non-automated vehicle.

There is no consensus about what the additional cost of an automated vehicle will be, with estimates ranging from a low of 10,000 US dollars to a high of 40,000 US dollars. As will be seen in the next section, trying to determine the cost more precisely is important to support the interpretation of studies of willingness-to-pay for automated vehicles. If the highest cost estimate is assumed, almost nobody is willing to pay that much for automated vehicles. If, on the other hand, the lowest cost estimate is assumed, there will at least be some who are willing to pay the additional cost of an automated vehicle.

4.11 Willingness to pay for automation technology

4.11.1 Existence of willingness to pay

Most studies of the willingness to pay for vehicle automation proceed in two stages. The first stage is to ask a sample of individuals whether they are willing to pay anything at all. If the answer is "yes", respondents proceed to the second stage in which they are asked how much they are willing to pay. We will first present the results of the first stage of studies of willingness to pay, then the results of studies of the second stage.

Table 6 shows the findings of studies of whether a positive willingness to pay for automated vehicles exists or not. The results of the studies vary a lot, but some studies find that a majority are not willing to pay anything for automation technology. In China, a clear majority of respondents indicate a positive willingness-to-pay for automation technology. In Australia and Japan, the opposite is the case. Results for the USA vary and are not always clearly presented, leaving readers to guess what the authors intended.

No attempt has been made to formally synthesise the results presented in Table 6.

Table 6: Existence of willingness to pay for vehicle automation in surveys.

Study	Country (N)	Not willing to pay (%)	Willing to pay (%)
Payre, Cestac, Delhomme 2014	France (421)	22	78
Schoettler, Sivak 2014	China (610)	22	78
	India (527)	30	70
	Japan (585)	68	32
	USA (501)	55	45
	Great Britain (527)	60	40
	Australia (505)	55	45
Kyriakidis et al. 2015	109 countries (4886)	22	78
Bansal et al. 2016	USA (347) SAE 4	19	81
Bansal, Kockelman 2017	USA (2167) SAE 3	55	45
	USA (2167) SAE 4	59	41
Bansal, Kockelman 2018	USA (1088) SAE 4	29	71
Cunningham et al. 2019	Australia (6133)	57	43
Liu, Yang, Xu 2019	China (441)	31	69
Liu et al. 2019	China (1355)	26	74
Shin, Tada, Managi 2019	Japan (188089)	53	47

4.11.2 Size of willingness to pay

The amounts people are willing to pay for automation technology, if positive, vary considerably. This is not unusual. Studies of willingness-to-pay are almost always hypothetical (i.e. no actual payments are made) and the results are heavily influenced by study design; indeed, so much so that one may doubt whether the results really reflect any objective reality. By objective reality in this context is meant that actual payments would equal stated payments in the surveys.

Be this as it may. Discussing methodological problems of willingness-to-pay studies can easily become a lengthy exercise. For the moment, only the results of the studies will be presented without going into detail regarding the methods employed in each study. Table 7 shows summary estimates of willingness to pay reported by various studies. The summary estimates are either mean values or median values. Not all studies report the information needed to estimate both mean and median values.

Estimates of mean willingness to pay for a fully automated vehicle are between US dollars 1740 and 7589. This range is quite small, but even the highest estimate is lower than the lowest estimate of the cost of an automated vehicle, 10,000 dollars (see above).

Table 7: Estimates of willingness to pay for automation technology.

Study	Technology	Respondents	Mean WTP	Median WTP
Payre et al. 2014	Fully automated	With WTP > 0	EURO 1624	
Schoettle, Sivak 2014	SAE 4	All, China		USD 1600
	SAE 4	All, India		USD 160
	SAE 4	All, Japan		USD 0
	SAE 4	All, USA		USD 0
	SAE 4	All, Great Britain		USD 0
	SAE 4	All, Australia		USD 0
Kyriakidis et al. 2015	Fully automated	All, 109 countries		USD 3800
Bansal et al. 2016	SAE 3	Not clear	USD 3300	
	SAE 4	Not clear	USD 7253	
Bansal, Kockelman 2017	SAE 3	All, USA	USD 2438	USD 0
	SAE 4	All, USA	USD 5857	USD 0
Daziano et al. 2017	Partial automation	All, USA	USD 3538	
	Full automation	All, USA	USD 4917	
Bansal, Kockelman 2018	SAE 2	Not clear	USD 2910	
	SAE 3	Not clear	USD 4607	
	SAE 4	Not clear	USD 7589	
Liu, Yang, Xu 2019	Fully automated	All, China	USD 2782	USD 563
Liu et al. 2019	Fully automated	All, China	USD 3520	USD 752
Shin et al. 2019	Fully automated	All, Japan	USD 1740	

There is a systematic pattern in the results, in that mean willingness to pay increases with level of automation. In studies allowing both mean and median willingness to pay to be estimated, median willingness to pay is considerably lower than mean willingness to pay. This is a common finding in valuation studies.

4.11.3 Demand curves for vehicle automation

A few studies provide enough detail to allow a demand function to be estimated. One of these studies is Liu, Yang and Xu (2019). shows the demand curve derived from their study. The curve is very skewed, meaning that it has a long almost flat part and a part where it rises very steeply. A curve with this shape arises when a small minority is willing to pay very much for a good, whereas the large majority is not willing to pay very much. This suggests that the introduction of automated vehicles may be associated with an increase in inequality in transport.

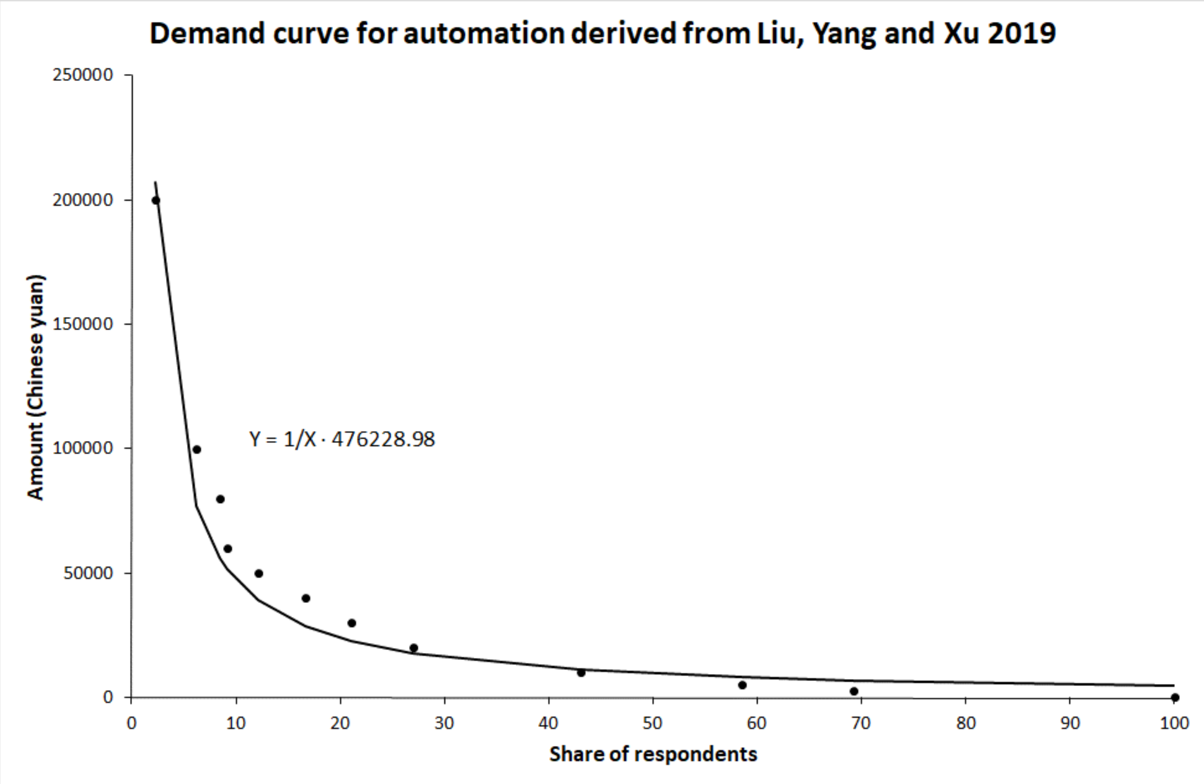

Figure 23: Demand curve for automated vehicles derived from Liu, Ynag and Xu 2019

4.12 Other societal impacts

4.12.1 Needs for parking space

Some of the studies reviewed by Soteropoulos et al. (2019) have estimated changes in the need for parking space. It should be noted that all these studies refer to shared mobility, i.e. they refer to a situation in which individual car ownership has to a very high degree been replaced by use of shared mobility.

Estimates vary from -80 % to -94 %, with a mean of -89 %. These estimates refer to a situation when shared mobility has completely replaced individual mobility, i.e. they represent the endpoint of a dose-response curve.

4.12.2 Inequality in transport

Inequality in transport denotes differences in access to transport resources (driving licence, car, other means of transport) and in the consumption of transport services between groups with different social status. There are many indicators of social status, but income is frequently used.

There are few studies of the potential impacts of automated vehicles on inequality in transport. A study by Cohn et al. (2019) relied on data for census tracts. It made the assumption that: “With ubiquitous implementation of automated vehicles, it is assumed all households would have access to a vehicle.” This is not very realistic in view of current estimates of the cost of automated vehicles. It is perfectly obvious that not all households could afford an automated vehicle. The only way all households could have access to one, was if they were given away freely. Besides, if you base the analysis on an assumption that everybody has access to the same transport resources, the results are a foregone conclusion: You will of course find that inequality has been reduced.

Towards the end of the paper, Cohn et al. admit that: “These assumptions are generous and aspirational and may not reflect initial pricing or accessibility of automated vehicles to low-income communities.” The analysis of Wadud (2017) is more realistic and shows what is, at least initially, likely to be a fact, namely that only the richest households can afford automated vehicles and benefit from them.

4.13 Summary of impacts that can be quantified

Table 8 lists potential impacts of connected and automated vehicles and indicates for which of these impacts it is judged feasible to develop a dose-response curve showing impacts as a function of the market penetration of automated vehicles.

It is seen that quantification of potential impacts by means of dose-response curves is possible only for a minority of the impacts listed in Table 8. For some impacts the maximum potential impacts can be estimated. For impacts in this category, it may be possible to develop dose-response curves by assuming that impacts follow the diffusion curve of automated vehicles into the market. The diffusion of new technology typically follows an S-shaped curve.

For some impacts, the likely direction of an impact can be indicated; yet developing a dose-response curve is fraught with large uncertainty. It is, for example, not clear whether automated vehicles will run on fossil fuel, be electric or be a mixture of electric and fossil fueled vehicles. What does seem likely in any case, is that fuel efficiency will improve. For some impacts, a dose-response curve based on technology diffusion appears reasonable. A case in point is effects on employment.

Finally, for some impacts no meaningful quantified estimate can be developed at the current state of knowledge. Some of these impacts may differ in the short and long term. Inequality in transport, for example, may initially increase as only the wealthy can afford automated vehicles. However, as prices come down, more and more people can afford automated vehicles, and, ultimately, the universal use of such vehicles can make travel more accessible for those who are currently limited in their travel opportunities. Thus, in the short run inequality may increase, in the long run it may decrease. Based on current knowledge, however, any quantification of impacts will be highly speculative.

4.13.1 Possibility of quantifying impacts

Table 8: The feasibility of quantifying potential impacts of connected and automated vehicles.

Impact	Possibility of quantifying impact
Direct impacts	
Travel time	Yes, by means of dose-response functions
Travel comfort	Not possible at the current state-of-knowledge
Valuation of time	Yes, as a percentage change in value
Vehicle operating cost	Yes, as a change in mean cost per kilometre driven
Vehicle ownership cost	Yes, as additional cost of buying an automated vehicle
Access to travel	Yes, as potential increase in travel if access is universal
Route choice (individual)	Not possible at the current state-of-knowledge
Systemic impacts	
Amount of travel	Yes, by means of dose-response curves
Road capacity	Yes, by means of dose-response curves
Congestion	Yes, by means of dose-response curves
Infrastructure wear	Yes, by means of a loading model
Infrastructure design	Not possible at the current state-of-knowledge
Modal split of travel	Not possible at the current state-of-knowledge
Optimisation of route choice	Not possible at the current state-of-knowledge
Vehicle ownership rate	Not possible at the current state-of-knowledge
Shared mobility	Maximum sharing can be estimated, not how to get there
Vehicle utilisation rate	Maximum utilisation can be estimated; pathway not known
Parking space	Maximum potential reduction can be estimated

Table 8: The feasibility of quantifying potential impacts of connected and automated vehicles.

Impact	Possibility of quantifying impact
Wider impacts	
Trust in technology	Not possible at the current state-of-knowledge
Road safety	Yes, several dose-response curves
Propulsion energy	It is not clear what source of energy will be used
Energy efficiency	Direction of impact can be indicated; possible dose-response
Vehicle emissions	Yes, if assumed proportional to fuel consumption
Air pollution	Direction of impact can be indicated
Noise pollution	No studies have been found
Public health	Not possible at the current state-of-knowledge
Employment	Direction of impact can be indicated
Geographic accessibility	Too few studies for dose-response curve; local variations
Inequality in transport	Not possible at the current state-of-knowledge
Commuting distances	Direction of impact can be indicated
Land use	Direction of impact can be indicated
Public finances	Not possible at the current state of knowledge

The next section discusses currently available dose-response curves and possible additions to these curves in greater detail.

4.13.2 Available dose-response curves

The dose-response curves that have been presented above are listed in Table 9 in their order of presentation. The possibility of adding dose-response curves is discussed.

The dose-response curves for impacts on accidents are all plausible. When applying the functions, a distinction should be made between three types of traffic environments:

1. Urban environments
2. Rural environments
3. Motorways

The dose-response curves for accidents in junctions will apply mostly to an urban traffic environment. Head-on collisions and running off the road are more common on rural roads. Rearend and lane change collisions are common on motorways, but running off the road is also an accident type that may occur on motorways. It will be assumed that there are no accidents involving cyclists or pedestrians on motorways (strictly speaking, this is not correct; there are some very few freakish pedestrian accidents involving people who have left their car and started walking along the motorway or trying to cross it; it is assumed that it will require an effort to identify such accidents in official statistics; hence they will be ignored).

Table 9: Dose-response curves for potential impacts of connected and automated vehicles.

Impact	Dose-response curve
Mobility impacts	
Road capacity	Second degree polynomial
Junction capacity	Second degree polynomial
Congestion delay	Second degree polynomial
Delays in roundabouts	Second degree polynomial
Delays in signalised junctions	Second degree polynomial
Changes in travel time	Exponential
Amount of induced travel	Second degree polynomial
Reduction of car fleet with shared mobility	Endpoint of dose-response curve only
Safety impacts	
Rearend and lane accidents on motorways	Second degree polynomial
Accidents in signalised junctions	Exponential
Accidents in priority junctions	Second degree polynomial
Accidents in roundabouts	Second degree polynomial
Cyclist and pedestrian accidents	Second degree polynomial
Head on collisions	Second degree polynomial
Accidents in junctions (in general)	Third degree polynomial
Single vehicle accidents; running off the road	Can be derived from diffusion curve (discussion in text)
Environmental impacts	
Fuel consumption	Second degree polynomial

The dose-response curves for impacts on mobility are also all quite plausible. The only impact of those listed in Table 9 for which a dose-response curve could not be derived is the reduction of the car fleet as a result of a complete transition to shared mobility. Such a transition must be regarded as improbable; it is altogether more reasonable to expect individual car ownership to continue on a wide scale. Be that as it may; as discussed above it is not necessary to make specific assumptions about whether mobility will be individual or shared, as most studies of shared mobility predict that vehicle kilometres of travel will increase, in line with the predictions that can be made based on knowledge of the travel demand induced by increased road capacity.

The dose-response curve for fuel consumption is highly uncertain. It is based on just three studies that provided a few data points. Another source of uncertainty is whether automation will be associated with a transition to electric vehicles.

4.13.3 Diffusion curves for new technologies

Can a dose-response curve be developed by assuming that it has the same shape as a diffusion curve for a new technology? Figure 24, taken from an ITF-report (ITF 2017) shows diffusion curves for a number of new technologies.

The diffusion curve for cars (auto in Figure 24) is the ragged green line starting between 1900 and 1915, going up and down between 1930 and 1945 and then increasing, but becoming flat at a level of about 90 %. Clearly, this diffusion curve is unlikely to describe how connected and automated vehicles will spread through the market. The diffusion of connected and automated vehicles is more likely to resemble the diffusion curves for colour TV, computers or the Internet.

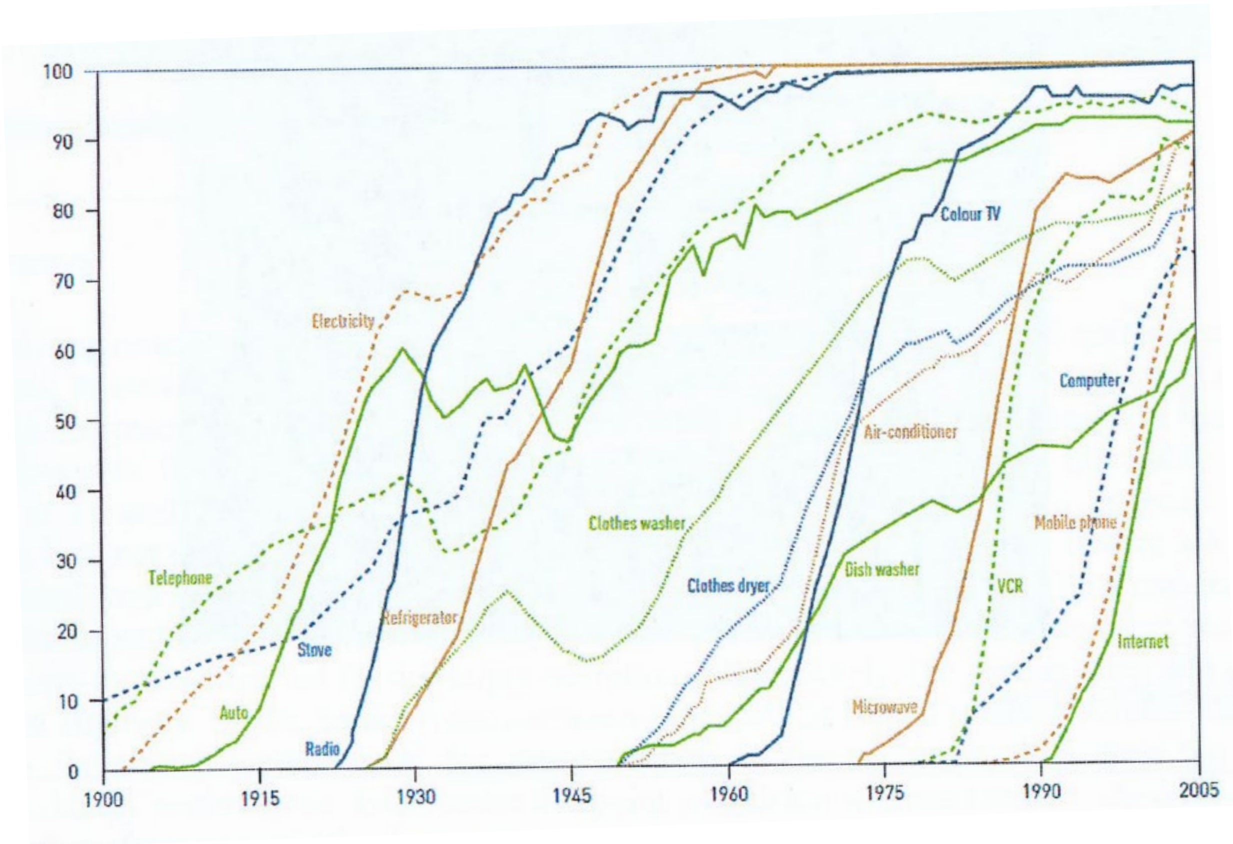

Figure 24: Diffusion curves for new technologies. Source: ITF 2017

The diffusion curves in Figure 24 are, however, all misleading in the sense that they only show the percentage of households or individuals who own an item, e.g. the percentage who own a colour TV, own a computer or have an Internet link. The curves do not show how much the products are used. It is reasonable to assume that, for example, you will watch a colour TV more when it is new than when it gets older. Likewise, a new computer will be used more than an old one. The same goes for cars: new cars are driven longer distances per year than old cars. Figure 25 shows mean annual distance driven by cars in Norway as a function of car age (smoothed curve).

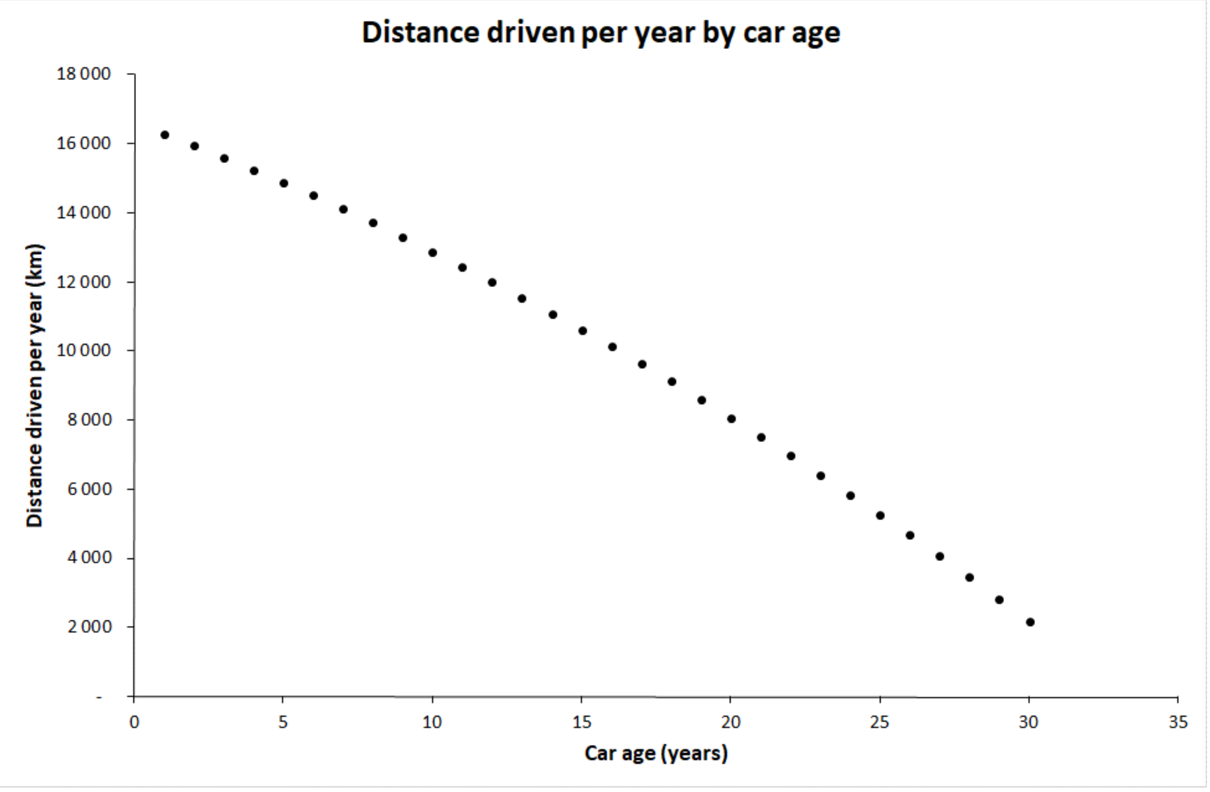

Figure 25: Mean annual distance driven as a function of car age

It is reasonable to think that initially, only a small share of new cars will be automated, say 10% of car sales the first year. If these cars are driven equally long as manual cars, they will represent 10% of the total distance driven by cars that are up to 1 years old. Next year, automated cars may represent 20 % of the sales of new cars, and, again assuming equal driving distance, 20% of the total distance driven by cars that are up to 1 years old. In addition, automated cars that are between 1 and 2 years old will drive 10% of the total distance of cars of this age. As automated cars increase their share of the sakes of new cars, the process will generate an S-shaped diffusion curve.

5 Estimating and quantifying secondary impacts

The objective of this chapter is to identify and discuss potential secondary impacts of connected and automated transport systems (CATS). A secondary impact is a rebound or feedback impact, or behavioural adaptation, which goes in the opposite direction of the primary impacts of CATS. As an example, if connectivity and platooning increases road capacity, there will be an adaptation to this in terms of induced travel demand. There has been less research on the secondary impacts of CATS than on the primary impacts. Therefore, although many secondary impacts can be identified and are likely to occur, estimating and quantifying them is more difficult than estimating and quantifying primary impacts.

The following potential forms of behavioural adaptation to CATS are discussed in this chapter:

1. Induced travel demand as an adaptation to increased road capacity, less congestion and possibly lower value of travel time,
2. Voluntary or involuntary reduction of cognitive effort made by drivers of conditionally automated vehicles; adaptation to time allowed for taking over control of vehicles,
3. Contagion of reduced safety margins of automated driving, in particular shorter time headways, to manual driving, increasing the risk of accidents,
4. Incomplete formalisation of interactions between automated and non-automated road users resulting in misunderstandings and gambling on safety margins,
5. Increased load on bridges, and potentially overload, as a result of platoons of heavy vehicles,
6. Cyber risks, i.e. the potential for hackers to attack CATS and block traffic or provoke accidents.

The only one of these impacts that can be quantified is induced travel demand. For the other secondary impacts, too little is known about their potential magnitude to try to quantify them.

5.1 Travel demand

The relationship between increase in road capacity and increase in traffic volume is a topic of long controversy. Today, however, there is wide (if perhaps not universal) agreement that a positive relationship exists. What continues to be discussed is whether so much new traffic is induced that it completely fills up the additional road capacity. Different studies have come to different results at this point.

It would detract from the main objective of this deliverable to delve into an extensive review of studies of induced travel demand. The issue is nevertheless worth discussing in some detail. The discussion will be based on studies published after 2000, both review studies and evaluations of the impacts of specific road projects. The review does not aim to be complete, but to give an impression of what different studies have found.

Noland and Lem (2002) reviewed evidence of induced travel demand in the United States and United Kingdom. Six estimates of the long-term elasticity of vehicle kilometres of travel with respect to lane kilometres (miles) per capita varied between 0.707 and 1.160, with an average value of 0.880. Note that if the elasticity is 1, all additional road capacity is consumed by increased traffic, meaning that there will still be congestion (although, as a result of traffic growth, more cars will be stuck in traffic) and no saving in travel time. In models using instrumental variables, to control for endogeneity in the supply of road capacity (meaning that roads are expanded when they become too congested), elasticity estimates varied between 0.289 and 1.944, with a mean value of 0.997; very close to 1. In a table listing elasticity estimates in previous studies, values varied between 0.29 and 1.1. For some studies, an interval was given. The mean value of the elasticity of vehicle kilometres with respect to lane kilometres, using the midpoint value in studies giving an interval (e.g 0.65 if the interval was 0.5-0.8), was 0.595.

Cervero (2003) conducted a path analysis of the long-term impacts of 24 road expansion projects in California. The long-term elasticity of traffic volume with respect to speed was 0.637.

Elvik and Amundsen (2004) evaluated road safety effects of four major road construction projects in the city of Oslo. As part of the evaluation, data on traffic volume was collected. Controlling for changes in traffic volume in Oslo in general, the four projects generated induced traffic of 6, 15, 20 and 40%. The corresponding increases in road capacity, crudely estimated, were, respectively, 20, 50, 50 and 100%. This gives elasticities of 0.30 (6/20), 0.30 (15/50), 0.40 (20/50) and 0.40 (40/100). It should be noted that these are short-term elasticities, which tend to be lower than long-term elasticities.

Elvik et al. (2017) evaluated road safety effects of a new motorway in Norway. Traffic increased by 82%, compared to 28% in Norway as a whole during the same period. Thus, net induced traffic was 42% (182/128). Road capacity increased by about 100%, making for a short-term elasticity of vehicle kilometres with respect to road capacity of 0.42 (42/100).

Tennøy et al. (2019) evaluated two cases of urban road expansion in Norway. In one case, expansion from two to four lanes was associated with a traffic growth of 53%, compared to 30% for Norway as a whole. If expansion from two to four lanes represents a 100% increase in capacity, short-term elasticity in this case was 0.18. In the other case, a road was expanded from two to four or five lanes. Traffic increased by 78%, compared to 22% in Norway in general. Again, if the increase in capacity is roughly estimated as 100%, elasticity of vehicle kilometres with respect to road capacity becomes 0.46.

Finally, Hymel (2019) developed ten estimates of the long-term elasticity of vehicle kilometres with respect to urban lane kilometres per capita, ranging between 0.703 and 1.063. A simple mean of the estimates is 0.811. He also listed previous studies. A simple mean of the elasticities estimated in those studies is 0.817. All were positive.

It is clear that an increase in road capacity is likely to induce more traffic. The summary estimates of long-term elasticities presented above are fairly consistent: 0.880 (Noland and Lem), 0.997 (Noland and Lem), 0.595 (Noland and Lem), 0.637 (Cervero), 0.811 (Hymel) and 0.817 (Hymel). The mean of these estimates is 0.789.

It is therefore suggested that an elasticity of 0.75 is applied, with an upper bound of 1 and a lower bound of 0.5. This means that vehicle kilometres of travel are predicted to grow at a rate of 75% of the increase in road capacity, in other words a parallell line to the dose-response curve, starting at the same point (100) and ending at a value of 190 at 100% market penetration, compared to a final value of 220 for the road capacity dose-response curve.

5.2 Shifting between automated and manual driving

SAE level 3 vehicles are referred to as conditional automation. These vehicles may operate automatically in certain operational domains. An operational domain may be defined as, for example, a certain type of road or certain environmental conditions. When the vehicle leaves an operational domain, the driver will be requested to take over control of the vehicle. Several studies have shown that prolonged driving in an automated mode makes drivers adapt to automation in ways that make them ill-prepared for taking over control.

In an early driving simulator study, Jamson et al. (2013) found that in automated mode drivers:

- Spent less time driving with a small safety margin (5.82% vs 8.26%), suggesting a gain in safety,
- Had their eyes closed more of the time (3.8% vs 1.8%), suggesting a loss of safety,
- Used the radio more of the time (54.1% vs 41.1%), also suggesting a loss of safety.

The net impact of these behavioural adaptations is difficult to predict. Based on naturalistic driving data, Dingus et al. (2016) estimated the risk associated with following too closely to 13.5 (i.e. those who followed too closely had an accident rate 13.5 times higher than those who did not follow too closely). Closing your eyes was not included, but, analogously, looking away from the road had a risk of 7.1. Operating the radio (turning on/off, tuning) had a risk of 1.9.

Eriksson and Stanton (2018) summarise a number of studies of driver reactions to take over requests. The studies they reviewed reveal a tendency for drivers to react more slowly if the time budget for taking over is long than if the time budget for taking over is short.

Zhang et al. (2019) report a meta-analysis of 129 studies of take-over time. The time budget given was up to 30 seconds, but in most cases between 5 and 10 seconds. A larger time budget was associated with longer response times. A curve fitted to the data indicated an increase in mean response time from about 1 second with the shortest time budget to about 7 seconds with the largest time budget.

Presumably all these trials refer to situations involving a controlled transfer of control from the car to the driver. By controlled is meant that the car knows that it is approaching the end of an operational domain, e.g. it will soon exit a motorway or the motorway ends. In such situations the system may offer a takeover time budget of up to 10 seconds. If a situation is unexpected, such takeover times can be unrealistic. If an animal suddenly jumps out in front of the car, there may not be enough time to react, irrespectice of whether the automation technology reacts or the driver is asked to react. The machine learning capabilities of automated vehicles can remember and retrieve information about regularly occurring and predictable events. It can, for example, remember the curvature of a road and prepare for the next curve. It cannot remember or prepare for an event that has never occurred before and develops suddenly. There will be such events.

Drivers adapt very quickly to automated driving. In a driving simulator trial reported by Omozik et al. (2019), it took only 4 minutes of automated driving before drivers started to glance away from the road for at least 2 seconds per glance. Once into the habit of looking at something other than the road, drivers continued to do so as long as automated driving lasted. Even a short interruption in automated driving did not break the habit of glancing away from the road. Paying attention to roads and traffic is demanding. Not having to do so reduces cognitive load, something drivers learn very quickly.

5.3 Safety margins – contagion models

Connectivity is a technology that is likely to be implemented before vehicles reach the highest level of automation. Connectivity can increase road capacity (or, more precisely, utilise it more effectively) and reduce fuel consumption by packing vehicles densely together in platoons.

One should be very clear about the fact that the advantages of platooning are realised by greatly reducing the safety margins human drivers adopt when following each other. Hence, if human drivers start imitating a platoon, by adopting equally small safety margins, there could be an increase in rear-end collisions, at least during a transition period before all vehicles have connectivity and are obliged to enter a platoon when entering a motorway.

In a transition period, it is possible that one or more lanes on multilane highways could be reserved for platoons, while the other lanes are used by manually driven vehicles. How will driving next to a platoon influence the headways adopted by human drivers? Gouy et al. (2014) performed a driving simulator study intended to shed light on this question. Three conditions were compared:

1. A platoon with short headway (0.3 seconds)
2. A platoon with long headway (1.4 seconds)
3. No platoon

When no platoon was present in the adjacent lane, the mean time headway was 3.3-3.4 seconds. In the short headway platoon condition, drivers reduced their mean time headway to 1.87 seconds. In the long time headway platoon condition, drivers reduced their mean time headway to 1.99 seconds. These are reductions of about 40-45%

compared to the no platoon condition, which indicates considerable behavioural adaptation.

While a time headway or more than 3 seconds is generally regarded as safe, shorter time headways require faster reactions to be safe. Reaction times vary between drivers; however, headways less than 1 second are generally regarded as unsafe. Although a driver may be able to react in less than one second, the remaining time will be too short to come to a safe stop.

A recent study by Rahmati et al. (2019) found that human drivers follow automated cars at a shorter distance than human driven cars.

Many years ago, Evans and Wasielewski (1982) studied the relationship between everyday driving behaviour and accident involvement. One of the characteristics of behaviour they included was time headway (observed from bridges across motorways; number plates were recorded and matched to driver records). For headway, they found the relationship shown in Figure 26.

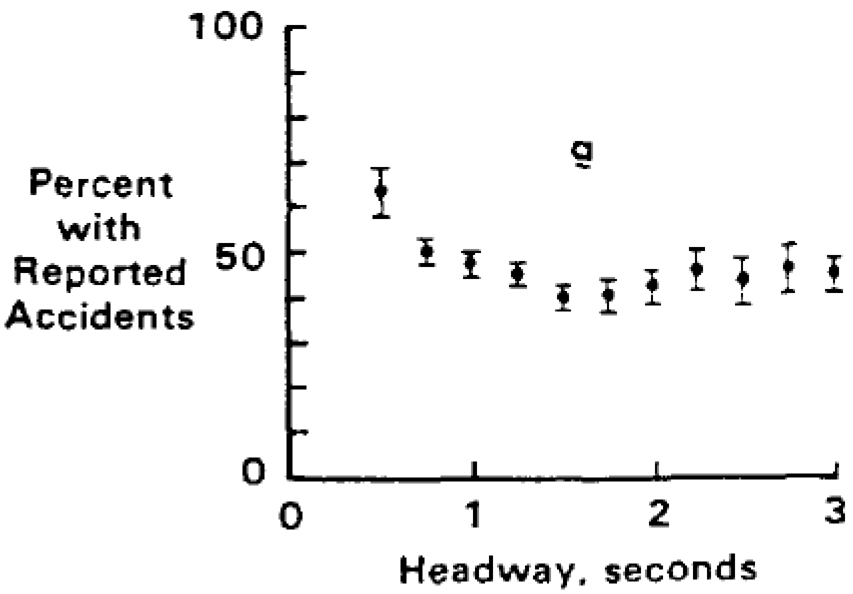

Figure 26: Relationship between time headway and accident involvement. Source: Evans and Wasielewski 1982, Figure 3a.

The relationship has a U-shape, although the bottom of the U is quite shallow. For drivers with a headway of less than 1 second, accident involvement increases. Headways between 1.5 and 3 seconds are not related to accident involvement; indeed, there is a weak tendency for accident involvement to be higher among drivers adopting the largest time headways. These drivers may trust their skills less than other drivers and for that reason choose larger time headways. There is, however, no doubt that adopting very short time headways increases accident involvement.

It should be noted that all accidents were included in this study. Had only rear-end collisions been included, a stronger relationship to time headway would most likely have been found. Yet, the fact that the study is old and that it is the only one of its kind means that it cannot be used to predict the impacts on accidents of behavioural adaptation to platoons. It seems likely that there will be behavioural adaptation, and that this will be detrimental to safety, but quantifying the effect would be speculative.

5.4 Interactions between automated and non-automated road users

One of the biggest challenges in developing automated vehicles is to ensure safe and effective interactions between automated vehicles and non-automated road users. Today, interactions between road users are partly governed by formal rules, partly by informal conventions. Interactions may rely on subtle cues that automated vehicles will be unable to grasp or perform.

Heymann and Degani (2019) have proposed a protocol for the interaction between automated vehicles and pedestrians at pedestrian crossings. The protocol can be explained by reference to Figure 27.

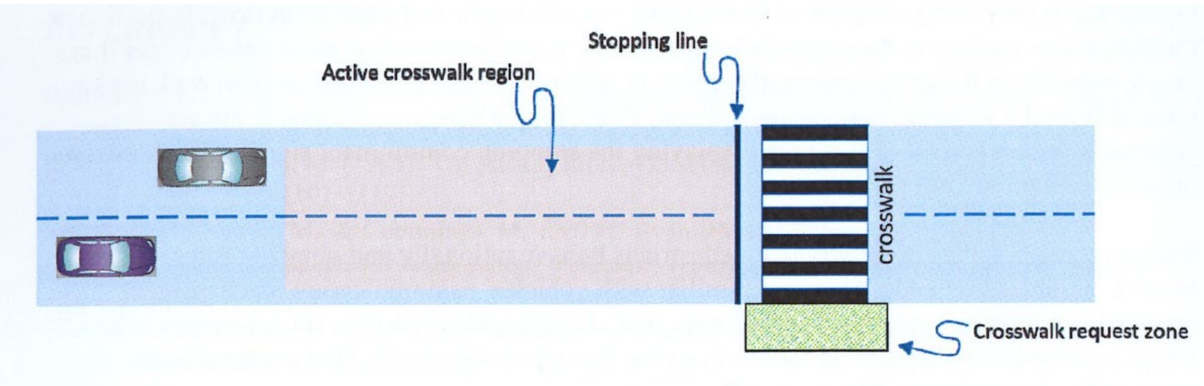

Figure 27: Outline of zones used to define interaction at pedestrian crossings. Source: Heymann and Degani 2019

An active crosswalk region (detectable by vehicle sensors) is delineated ahead of a pedestrian crossing. When a vehicle enters this region, it transmits a crosswalk recognition signal, i.e. a signal that informs pedestrians that the vehicle knows that it is approaching a pedestrian crossing. If the vehicle, when entering the active crosswalk region, detects a pedestrian in the crosswalk request zone or in the crossing itself, it transmits a stopping commitment signal to the pedestrian. This signal is also displayed to other vehicles in the active crosswalk region. The vehicle(s) come(s) to a full stop at the stopping line.

A pedestrian who is in the request zone is obliged to cross the road provided that: (1) There are no vehicles in the active crosswalk region, or (2) All vehicles in the active crosswalk region have stopped or indicated that they will stop. Vehicles start moving again when all pedestrians have crossed completely.

If complied with by everybody, this protocol ensures that automated vehicles will give way to pedestrians, who may cross in safety. Yet, promising as the protocol is, it is easy to imagine situations in which it will not function as intended.

First, not all pedestrians intend to cross the road. Depending on the width of the request zone, a pedestrian who intends to continue walking without crossing, may have to enter the request zone. Vehicles will then be erroneously informed that a pedestrian intends to cross and stop to wait for the pedestrian. Some limit on waiting time must then be imposed before vehicles start moving again.

Second, a pedestrian may enter the request zone when a vehicle approaches the end of the active crosswalk region. Will the vehicle then come to a sudden stop, or will it continue and pass the crosswalk? There will be a dilemma zone, and misunderstandings about who gives way when a vehicle is in the dilemma zone will lead to accidents.

Third, to provide maximum safety the length of the active crosswalk region must be sufficient for large and heavy vehicles to stop. While heavy braking may not be a problem for a driverless truck, it could be a problem for a bus. A bus cannot slow down too fast; otherwise occupants may fall inside it and get injured. A long active crosswalk region could tempt pedestrians to cross before vehicles have reached the stop line. If all pedestrians manage to do so, an ambiguous situation may once again arise. When the vehicle reaches the stop line, all pedestrians will have crossed already and there is no longer any need to stop. How quickly will the system understand this?

Fourth, will the same rules apply to cyclists as to pedestrians? In Norway, a cyclist cannot cycle across a pedestrian crossing, but has to give way to cars. If, however, the cyclist dismounts, he or she is regarded as a pedestrian and then has the right of way. It should come as no surprise that most cyclists do not follow this rule. The stay mounted and cycle across the crossing and cars give way to them. This outcome is in everybody's interest (Bjørnskau 2017). It saves time both for cyclists and the car drivers. But it is an informal solution, which violates the legal regulation.

Fifth, there will be drunk pedestrians (and cyclists) who do not understand the system and will not comply with it. They may, for example, enter the crossing without looking for vehicles or stay put in the request zone so long that cars start moving, and only then start to cross the road. All sorts of aberrant, irrational behaviour can be imagined.

It therefore seems likely that there will be behavioural adaptation to a system as outlined in the protocol by Heymann and Degani (2019). It is likely that such behavioural adaptation will reduce the effectiveness of the system and make it unlikely that it will prevent all crashes involving motor vehicles and pedestrians. However, any numerical estimate of the adaptation effect would be speculative.

5.5 Overloading bridges by platooning

The prospect of increased road capacity due to automated driving leads to the question of effects on infrastructure wear. Shortening of the headway distances leads to a higher number of vehicles on road sections of a given length. The primary effect on infrastructure wear is expected to arise from more concentrated loads of heavy vehicles in combination with automated driving. Especially the formation of truck platoons, which should reduce fuel consumption by using positive aerodynamic effects, will lead to an increased loading of roads and bridges compared to a single vehicle entity. By contrast, increasing the number of passenger cars on sections will have minor effects on infrastructure loading. Secondary effects of automated driving are expected to manifest themselves in changed traffic flows and possibly modified dynamic amplification of axle loads on bridges.

Most impacts of changed loading are expected to concern bridges. The standards for designing are regulated by the *European Committee for Standardization (CEN)*, which covers almost all member countries within the European Union. There is a master code with the main regulations; each member state has his own national annex for covering regional requirements. The load models are regulated in the Eurocode 1991-2 "*Traffic loads on Bridges*". This code defines the characteristic load model for dimensioning new main bridges. The so-called Load Model 1 (LM1 is based on a 1000-year return period of traffic loads (based on a reference in 2003) which is equal to a probability of exceedance of 5% in 50 years for traffic on the main roads in Europe (Calgaro 2008). This enables to fulfill the structural safety regulations of the CEN Standards which should cover the design life of 100 -120 years or more. This restriction is predominantly used for the design of new bridges. For assessing existing bridges, the remaining service life is generally less than the design life and it may be more appropriate to use a shorter return period of traffic events, which is equivalent to a lower level of safety (Leahy 2016). There are different approaches for assessment of existing bridges. Leahy (2016) suggested a return period of 75 years (as it is used in bridge design in the United States).

Another aspect is the change of load patterns due to automated traffic. Stochastic traffic produces random forces, while automated driving or truck platoons lead to synchronised traffic flows with synchronised loading. Synchronization will change the dynamic interaction or can increase the horizontal loading, for example in case of an emergency braking performed by the whole truck platoon. It will be necessary to identify the changed load impacts as a consequence of the truck platoon configuration and to investigate their effects on infrastructure components. For example, bridges with smaller spans (< 30 m) are more sensitive to dynamic interactions, while longer bridges will be more restricted by the maximum number of trucks on the bridge. Steel bridges with low dynamic damping factor behave differently than concrete bridges.

The most **relevant load impacts** are expected on bridges, which will be:

a) Vertical loading
b) Horizontal loading
c) Dynamic interaction

Ad a) It is obvious that increasing the number of trucks/per h or shortening the headway distances leads to a more trucks per section which changes the infrastructure loading. The effects of **vertical loading** are mainly dependent on the vehicle weight, the bridge span and platoon length/headways and must meet the bridge design restrictions.

The change of load can be demonstrated in the following example: A truck platoon consisting of 8 trucks, with average truck load of 39 tons (390 kN) and a single truck length of 14 m, is passing with a speed of 80 km/h over a bridge with a length of 128 m. Using a common mean time headway of 2 seconds (44.4 m) results in a total bridge loading of 2.19 trucks or 85 tons. Reducing the distance between the vehicles to 2 m, the whole platoon will fit with its 8 trucks or 312 tons on the bridge, which is an additional loading of + 267 %.

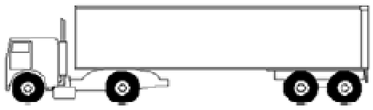

Axle spacing[m]	Axle weight [kN]
3,40	70
6,00	140
1,80	90
	90

Example of an equivalent trailer truck according to EN 1991-2 (Calgaro 2008).

Other restrictions of increased loading could arise from torsional effects due to different loading in individual lanes, or due to local effects corresponding to transverse bearing capacity.

Ad b) The change of **horizontal loading** is mainly expected as the change of braking loads due to abrupt braking of a whole truck platoon on a bridge. According to EN 1991-2 new bridges are designed to a horizonal load which is depending on the total bridge length (Calgaro 2008):

$$Q_{lk} = 360kN + 2{,}7\ L < 900\ kN \quad for\ 1{,}2m < L < 200\ m$$

For bridge with 1,2 m length the horizontal load Q_{lk} due to breaking of LM1 would be 363 kN and rises to a maximum of 900 kN for bridges with 200 m length.
The effects can be illustrated using the above truck platoon with 8 trucks, which performs an emergency break with deceleration of 6 m/sec² (dry asphalt) on a bridge with length L=128 m. Applying Newton's law (force = mass times acceleration), the acting horizontal force of the platoon is:

$$Q_{lk\ p} = m * a_b = \ 8 * 39000\ kg * 6\ ms^{-2} = 1.872\ kN$$

This is far more than the Eurocode load, which would be only 706 kN. It cannot be guaranteed that existing bridges are designed for that level of horizontal forces.
Other effects of horizontal loads are centrifugal loads or accidental loads of single straying trucks. Here also, some changes can be expected due to platooning.

Ad c) **Dynamic effects** are mostly observed when the traffic flow changes from random to more synchronised forms. The changed dynamic amplification and possibly increased resonance loading effects, which are well-known from railway bridges, could now act also on existing road bridges, which were not designed for that.

In general, the dynamic effect of traffic loads is determined by many factors, such as maximum bridge span length, bridge natural frequency, vehicle weight, axle loads, axle configuration, vehicle suspension properties, current vehicle position on the bridge, quality of pavement, and stiffness of structural members. However, a large contribution may be attributed to vibrations of the vehicle induced by the road profile roughness, depending on the velocity and road surface unevenness (Cantieni 1992, Prat 2001).
Changed fatigue consumption of bridges and their components due to synchronised loading could affect the expected bridge lifetime.

The effects of changed loading due to automatisation and truck platooning must consider existing bridges and detailed investigation of different bridge types and truck platoons are needed for assessment of these effects. Bhoopalam et al. (2018) stated that measures to fullfill the infrastructure needs could lead to restrictions on maximum weight or smart division of truck loads, or decoupling of truck platoons at bridges. Other measures could be temporary changes in the truck platoon speed or changes of headway distances, so that the bridge loads would not exceed the original design levels. Another option would be legal limitations on the total number of trucks in a platoon, as demonstrated in the lessons learnt of European truck platooning challenge (Eckhardt et al. 2016).

5.6 Cyber risks

The purpose of this section is to investigate the effects of CATS on cyber security of the transport system. This is done by conducting a security risk assessment as described by the Norwegian Standard for Security Risk Assessment (Standards Norway, NS5832:2014). NS5832:2014 describes how to carry out a risk assessment when the probability of events is difficult or impossible to calculate. Scenarios have been developed to describe some types of events that may occur.

A security risk assessment, as described in NS5832:2014, consists of seven steps:

(1) Asset identification
(2) Determination of security objectives
(3) Threat assessment
(4) Scenario selection
(5) Vulnerability assessment
(6) Assessment of pure (negative) risk
(7) Presentation of the risk situation

Asset identification (1) consists of mapping, assessing and ranking the assets (both material and non-material) that are at risk. Determination of security objectives (2) involves deciding what is the desired or acceptable conditions for the assets during or after an undesired event. The threat assessment (3) should identify and describe possible malicious actors, their intention and capacity for attacking. Scenario selection (4) should result in scenarios that are based on assessments of consequences and threats, and is considered as relevant for further analysis. The vulnerability assessment (5) should point out which vulnerabilities, or security weaknesses, that a threat actor can exploit in the selected scenarios. Finally, the assessment of pure risk (6) should compile the results from the consequences, threat and vulnerability assessments to an assessment of the pure risk for each scenario, and the presentation of the risk situation (7) should summarise the pure risk for all the selected scenarios.

5.6.1 Asset identification

A traditional transport system consists of:

(1) vehicles (cars, buses, bicycles etc.),
(2) the infrastructure the vehicles use (roads, bridges, tunnels, the network of fuel providers etc.), and
(3) transport providers (taxi drivers, public transport companies, logistic companies)
(4) people using the transport system incl. passengers in vehicles, bicyclists, pedestrians and children that play in the streets
(5) commodities that are being transported

CATS will also comprise of connectivity elements:

(6) systems collecting, processing and distributing information about weather, traffic and road works to enable vehicles to optimise travel choices.
(7) platforms for ordering and distributing vehicles wherever and whenever transport is sought.
(8) servers providing (automatic) updates of vehicles.

Finally, CATS will collect vast amount of data:

(9) data about when and where people travel

Which of the above elements are critical for the transport system to be working? (1) No single vehicle is critical for the transport system as such, but connectivity makes it possible to immobilise a larger share of the vehicles, which can seriously hinder transport. (2) The infrastructure the vehicles use is the same as used by non-autonomous vehicles. Some parts of the infrastructure can be bottle necks (bridges, tunnels etc.) and these points could be clogged by immobilised vehicles or hacking of equipment installed at these points (barriers, traffic signals etc.). (3) Most transport providers are not critical for the transport system. An exception might be monopolists in an area, such as Ruter which is responsible for all public transport in the Oslo area. (4) A single person is not critical for the workings of the transport system. Public incidents where individuals are killed, seriously injured or hijacked in an autonomous vehicle might, however, scare the general public from using autonomous vehicles. (5) Most commodities are not of critical value.

(6) Systems that provide information for vehicles to optimise travel choices increase the efficiency of the transport system, but are not critical for the working of the transport system. Downtime or misinformation in such systems, especially if there exists only one such system, can create traffic jams. (7) Platforms for ordering and distributing vehicles are critical in a situation where all vehicles are for public hire instead of private. In a situation with few platforms for ordering and distributing vehicles these can be critical for the transport system to be working. (8) Servers providing (automatic) updates can compromise a large amount of vehicles at the same time so that so many are not working/clogging the streets that the transport system stops working properly.
(9) The data produced is also part of the transport system. Data about individuals travel patterns are sensitive and can provide information about individuals and if it gets astray could reduce trust in the transport system. Data about travel patterns at a higher level will be used to plan the transport system and can be misused by a malicious actor.
A final value connected to (4) and (9) is confidence in the transport system. If potential users do not believe that it is safe to use the transport system, they might choose to not use it after all. Fear of getting killed, injured or being victim of crime might make people not travel.

The most important value is of course the people in the transport system. If everybody using the transport system were killed, the transport system would have no purpose anymore. However, the likelihood of managing to kill all the people using a transport system in one cyberattack is low. And cyberattacks causing a few deaths does not per se threaten the transport system (but might make people avoid using the transport system and thus threaten its existence). Individual elements that can be critical for the functioning of the transport system are (1) servers that provide updates to a large number of vehicles, (2) platforms ordering and distributing transport, (3) physical bottle necks in the transport system and (4) energy refilling infrastructure (including fossile fuels, electricity, hydrogen etc.). In a situation where people have the option of not travelling/using alternative transport (such as non-autonomous vehicles), confidence in the transport system also is a critical value for the transport system.

5.6.2 Security objectives

Before delving into incident-specific security objectives, general security objectives for connected and autonomous cars are discussed.

1. The first objective is an absolute reduction in the number of killed and seriously injured road users compared to the current situation, and a reduction in the risk of being killed or seriously injured per kilometre travelled. This objective is not especially ambitious given the fact that some countries already aim for zero fatalities and seriously injured road users (Vision Zero).
2. The second objective is that cyberattacks should not incur larger costs on society than the efficiency gains achieved by the introduction of connected and autonomous vehicles.
3. The third objective is that breakdown-time, the time the transport system does not work, should be maximum 48 hours per year. Heavy snowfall causes school, public transport and other public functions to close down many places and at least 48 hours of downtime per year is not uncommon.
4. The fourth objective is that breakdowns in the transport system should not cause so many delays in emergency missions that the number of people dying and the economic costs outdo the gains of the first and second objectives.
5. The fifth objective is that the confidence in the transport system should be so large that 95% of the population are willing to travel by connected and autonomous vehicles.
6. The sixth objective is that the confidence in the transport system should be so large that 90% of the population report that they feel secure when travelling with connected and autonomous vehicles.

The general security objectives can be translated into incident-specific security objectives; the desired or acceptable conditions for the assets during or after an undesired event. A difficult question is whether it matters how large the consequences are of a single cyberattack when the consequences of all cyberattacks still meets the general objectives (as described above). For the objectives of number of deaths and seriously injured and economic costs of cyberattacks it does not matter whether the unwanted consequences are caused by one large incident or several smaller incidents. Large incidents might cause more attention and thus fear, but for society it is irrelevant whether all the fatalities occur in June or are evenly distributed throughout the year. For downtime of the transport system the answer is different: Small delays have seldom large consequences, while longer delays can cause shortages of commodities et cetera. The general objectives of confidence in the transport system might be obtainable for most of the time, but perhaps not just after a cyberattack that has caused a lot of damage and/or received a lot of attention.

7. The first incident-specific objective is that no cyberattack should cause more fatalities and seriously injured than defined by the first general security objective per year.

8. The second incident-specific objective is that no cyberattack should incur larger economic costs than defined by the second general security objective per year.
9. The third incident-specific objective is that no cyberattack should cause more than 24 hours of downtime of the transport system.
10. The fourth incident-specific objective is that no cyberattack should delay emergency missions such that people die, fires develop further and conflict situations escalate into violence because of the delay.
11. The fifth incident-specific objective is that 7 days after a cyberattack at least 90% of the population should be willing to use autonomous vehicles again.
12. The sixth incident-specific objective is that 21 days after a cyberattack at least 80% of the population should report that they feel safe as passengers of autonomous vehicles.

5.6.3 Threat assessment

Hacking refers to "activities that seek to compromise digital devices, such as computers, smartphones, tablets, and even entire networks" (Malwarebytes, 2019). Connected and autonomous vehicles are examples of digital devices that can be hacked.
Hacking is typically technical in nature. But hackers can also use psychology to trick the user into clicking on a malicious attachment or providing personal data. These tactics are referred to as "social engineering" (Malwarebytes, 2019). Furthermore, the control systems in autonomous vehicles suffer from limitations that can be exploited by a malicious attacker. For example, the autopilot of Tesla vehicles has been shown to interpret white crossing vehicles as brightly lit sky and act accordingly (Hawkins, 2019). A malicious attacker could change the physical environment to make the control system in the autonomous vehicle act differently. In this section, the term "autopilot manipulation" is used to describe malicious attacks where the "hacker" changes the physical environment to manipulate the autonomous vehicle to act differently.

Cybercriminals can be distinguished by motive. Zamora (2018) lists six main motives for hacking:

(a) Fun/the challenge: When the Black Report (Pogue, 2018) asked hackers about why they hack (several answers possible), 86 percent replied that they liked the challenge of hacking and hacked to learn. Additionally, 35 percent said they did it for the entertainment value or to make mischief (Pogue, 2018).
(b) Financial gain: 21% of the respondents in the Black report survey replied that they hacked for financial gain (Pogue, 2018).
(c) Emotional: Some of the most destructive cybercriminals act out of emotions, such as rage, revenge, "love" or despair. These criminals may cyber stalk, access accounts without authorisation or use Internet of Things (IoT) to commit domestic abuse (Boyd, 2018; Zamora, 2018).
(d) Ego: Strengthening a weak ego is a motivation that can be caused by several psychological vulnerabilities, such as insecurity, financial woes and emotional turmoil (Zamora, 2018).
(e) Political/religious: Six percent of the hackers in the above survey replied that they did it for social or political motives. Such activities can be labeled as hacktivism, cyber terrorism and/or state-supported cybercrime (Zamora, 2018).
(f) Sexual impulses/deviant behavior: This sixth and last category encompass deviant acts motivated by sexual compulsion; rapists, sexual sadists, pedophiles and

serial killers can either use their own skill or hire unscrupulous people to carry out predatory sexual behaviors (Zamora, 2018).

It is also possible to distinguish by hacking actors by capacity: The lone hacker (i), an ad hoc group (ii), an organised group (iii), a state-supported group (iv), and a state (v).

(i) A lone attacker may have more or less capabilities, but does commit the attack alone, which often limits what the attacker can do. An example of a lone attacker would be a rejected lover wanting to take revenge by stealing information from the ex-lover's digital device and then use it to destroy this person's life.
(ii) An ad hoc group is a temporary group formed of people with hacking capabilities to do some mischief. They do not have a formal hierarchy, but have a common goal; such as anger towards an organisation/a company/the authorities or financial gain.
(iii) An organised group collaborate frequently and have designated roles, which enables the group to employ its human resources more efficiently in a cyberattack.
(iv) A state-supported group is an organised group that gets support from a government and thus have even more resources available when attacking. Such groups typically inflict harm that the government wants to cause, but cannot legitimately do themselves.
(v) We say that the attack is conducted by the state if any part of the government's formal organisation conducts the attack.

In the next section, the above actors and motivations are combined to produce some scenarios.

5.6.4 Scenario selection

Five scenarios that describe how malicious actors can proceed to adversely affect the assets have been created for further analysis. These scenarios cover 4 out of 6 motivations for hacking (for the fun/challenge, financial gain, emotions, political/religious) and 4 out of 5 actor capacities (lone actor, ad hoc group, organised group, state).

- Playful teenagers

Three teenagers, who are very interested in computers and autonomous vehicles, decide that they want to play with a full-sized vehicle. One of the teenagers has an uncle who has just bought a fully autonomous vehicle and they decide to hack the vehicle to use it as a remote-controlled car. They have access to the uncle's home and there they find the vehicle's id and other information about the vehicle that makes it easier to gain control over it. The teenagers use a couple of weeks browsing the internet to obtain information on how to hack the vehicle and software/scripts that they can use in the hacking.

One weekend, when they know that the uncle has left the country for a short holiday, they hijack the vehicle. They manage to get it to travel from the uncle's home to the school the teenagers are attending. They discover, however, that they cannot control the speed of the car. They, therefore, search the internet to see if they can find software that make it possible to control the speed. They find such software and manage to install it into the vehicle computer.

They make the car drive into the street again and this time faster than the speed limit and faster than the road environment permits. The vehicle drives past a lawn where some children are playing with a ball. One child throws the ball so that it ends up in the road and another child runs after it. The autonomous vehicle's sensors discover the child, but the vehicle's speed is so high that it is not physically possible for the vehicle to slow down enough before it hits the child and the child is killed immediately.

- Kidnapping for ransom

A group of organised criminals are looking for new ways of earning money. They recruit some members that are adept at hacking and decide to kidnap the ten-year-old child of a billionaire in their geographical proximity. The billionaire has just bought one of the newest autonomous cars available on the market. The hackers suspect that the vehicle has been released to the market prematurely and that there still are bugs in the software that can be exploited in a cyberattack.

The hackers successfully gain access to the in-vehicle system and manage to put surveillance on the vehicle. They are able to follow the vehicle's movements in real time and to listen in on communication and other sounds inside the vehicle. They quickly discover that the child attends dancing lessons every Tuesday evening. The hackers also gain access to the control system and locking system.

The group of organised criminals decides to strike one such Tuesday evening. First, they plant a mobile phone jammer under the vehicle while it is waiting to pick up the child. When the dancing lesson has finished and the child is on the way home, the hackers lock the vehicle and make it impossible for the child to open it. They also hijack the vehicle so that it is driven to a predefined site in the forest. All the time they are listening for sounds that reveal that the child is noticing that something is wrong. At the moment they understand that the child gets worried, they turn on the mobile phone jammer.
When the vehicle has arrived at the predefined site in the forest, the organised criminals open the locks and take the child to a prepared safe place. At the safe place, they make a ransom movie that shows that child is physically unhurt and present their ransom demand. When the criminals have received the money, they give the child back to the parents.

Initially, the parents do not want it made public that they gave in to ransom demands, and keep the hijacking secret. However, after deliberating it with their lawyers, they decide to sue the manufacturers of the autonomous vehicle for not having invested enough in cyber security to prevent such hijackings. The subsequent media attention reduces public confidence in autonomous vehicles and some people avoid using such vehicles for some time.

- Domestic abuse

An IT-expert has been rejected by a lover after a couple of months of dating. The ex-lover had just obtained an autonomous car and all the user information and documentation was lying around in the ex-lover's apartment, easily accessible for the IT-expert. The IT-expert was curious about the technology and had therefore spent some time reading up on all the documentation when visiting the ex-lover.

The break-up is ugly and the IT-expert feels hurt and badly treated by the ex-lover. The IT-expert decides to take revenge by playing some "tricks" on the ex-lover. The IT-expert succeeds with hacking the vehicle, gaining access the control system, heating system and the locking system.

The IT-expert conducts the cyberattack one cold winter morning when the ex-lover is travelling to work. First, the IT-expert locks the vehicle and makes it impossible for the ex-lover to open the vehicle. Then the IT-expert make the vehicle drive a new route where the vehicle can drive at least 70 km per hour for many hours. Finally, the IT-expert turns off the heating system, which makes the temperature in the vehicle very low.

The ex-lover discovers quite fast that something is very wrong with the vehicle and calls up the emergency number. The police try to hunt down the vehicle, but experience some difficulty with intercepting the vehicle because the route is confusing and they therefore do not know where it is travelling next. Finally, after approximately an hour do they manage to intercept and force the vehicle to stop.

The ex-lover is physically unhurt, but very traumatised. Since it is difficult to establish what caused the incident, public confidence in autonomous vehicles is seriously hurt and many people end up avoiding using such vehicles for a period.

- Autopilot manipulation

The introduction of autonomous vehicles has made private vehicle transport much more attractive and thus caused increased car traffic. Even if the new cars cause less or no exhaust emissions, they create particulate matter because of wear and tear when car tires meet the road surface. The increased car traffic has therefore caused increased local pollution and reduced air quality.

Environmental organisations are furious about what they perceive as the politicians' lack of willingness to implement measures that prevent increased car traffic. One of the more radical organisations decides to take action. They aim to reduce confidence in the transport system by exposing the vulnerabilities of the autopilots and, thus, decrease the use of autonomous vehicles.

Many of the local branches decide to do a coordinated campaign. They decide on a date and time (a spring morning around 9am) for the demonstrations. They create physical dummies that are supposed to fool the autopilot into believing that the lane is going another way.

Most of the branches succeed with misleading the vehicles to run off the road. The vehicles are seriously damaged, while the passengers suffer mostly from minor injuries and trauma. One branch, consisting of really furious teenagers, decides to mislead the vehicles into driving in the lane for the opposite direction of travel. This causes one large traffic accident with several people getting seriously injured.

Public confidence in autonomous vehicles is seriously damaged and many people end up avoiding using such vehicles for a period.

- Paralysation of transport system

Hostilities between two countries have been increasing for some time and the two countries are currently on the brink of war. One of the countries has decided to openly demonstrate its cyber capabilities and conduct a cyberattack.

Through public sources, the intelligence services find out the share of different car brands in the target country. They select a car brand with 23% share of the vehicles in the country. After choosing a car brand, they identify the server that the vehicles get over-the-air updates from. They then use time to gain access to the server without being discovered.

After gaining access to the server, the intelligence services add a hostile program to a scheduled update. The program is supposed to make all vehicles stop wherever they are located at 4 pm at a specific date (a week day) after the scheduled update. The intelligence services are successful in infecting all the vehicles of the car brand without being discovered.

On the planned date, the hostile program switches on and all the vehicles stop abruptly. Many other vehicles do not manage to stop in time and this results in multiple traffic accidents all over the country. The immobile vehicles block other traffic, making it impossible to travel by car in many dense city areas. The immobile vehicles also impede emergency vehicles, delaying emergency missions and making it impossible to gain access for the emergency vehicles in some inner city areas.

5.6.5 Vulnerability assessment

Which vulnerabilities, or security weaknesses, can a threat actor exploit in the selected scenarios? Furthermore, will the developers will invest sufficient resources to cyber security to prevent such attacks?

Fully autonomous vehicles are still being under development. Developers therefore have the opportunity to design in security, such as building encryption and cryptographic code signing into a vehicle's system, minimising the attack surface hackers can abuse, and tightly locking down communications with the outside world (Thomson, 2018).

Many of the vehicles already on the roads are connected and thus vulnerable to hacking. The industry has until now not been willing to invest the necessary amount of resources in cyber-security to prevent hacking. One possible explanation is that the car manufacturers have failed to realise that secure software is critical for automotive safety (Consumer Watchdog, 2019), but this explanation was more plausible a few years ago before several well-published hacking experiments (Osborne, 2018). Another possible explanation is that the industry is in a state of Nash Equilibrium in which all of the major car manufacturers are aware of the risk of hacking, but no one is motivated to unilaterally increase the expenditure on cybersecurity. Investing in cybersecurity diverts resources away from development of customer-visible features and thus make them less competitive. Since hardware and software differ between the manufacturers, an actual cyberattack will probably only hit one of the manufacturers. Each manufacturer may therefore prefer to gamble on not being targeted in a cyberattack until external forces, such as regulation, force them to increase expenditure on cybersecurity (Consumer Watchdog, 2019).

A third possible explanation for limited cybersecurity is that the manufacturers want to be able to harvest data to monetise it. To ensure that it is possible to harvest the data, all vehicles need to be accessible to a central computer collecting the data.

- Playful teenagers

In the past, cloning electronic keys has been the most common way of gaining illegal access to a vehicle. But vehicles are getting more connected with technologies such as WIFI and 3G,4G and 5G, giving hackers multiple ways to get in. More than a quarter of current attacks exploit vehicles' cloud servers or mobile apps (Bird, 2019). In Britain, vehicle thefts have increased by 50% in five years partly because of criminals hacking keyless systems (Hull, 2019). In this scenario, the teenagers have access to the vehicle owner's home, which may make it even easier to gain control over the vehicle.

Vehicle manufacturers will have to deal with the risk of vehicle theft. GPS and other types of tracking might reduce thefts motivated by financial gain, but will not have the same preventive effect against perpetrators who only want to "borrow" the vehicle for a limited time. And even if the manufacturers spend enough on protection against unrightful access by strangers, it still might be possible to take control of a vehicle if one has access to the vehicle owner's home. Hence, it seems probable that scenarios where people unrightfully play with other people's autonomous vehicles will occur.

Tuning up an engine to make a vehicle capable of going faster is also quite common, partly because tuning up rarely leads to penalties for the vehicle owner or other people responsible. And if anyone is able to create software that makes the vehicle drive faster than it is programmed to, they may also make the software available for anyone on the internet. Will manufacturers spend enough on cyber security to counter such hacking attempts? Since there are so many that may be motivated to manipulate autonomous vehicles' speed, there is a real probability that a few of them will discover and exploit any vulnerabilities in the programming. It might therefore be impossible to stop it from happening at all. But it might be possible to monitor what hacking software is made available and send updates that counter such software whenever the manufacturers discover its existence. It may also be possible to monitor speeds of the vehicles in real time to detect whether the speed control has been manipulated. Even if it is not possible to stop all tuning up, it should be possible to discover vehicles that have been tuned up and prevent such programming being commonplace.

A related question is whether the manufacturers will have incentives to spend the necessary amount of resources to prevent tuning up. That depends on whether they are held responsible for accidents caused by tuning up. If the manufacturers are able to say that they do not have responsibility for a traffic accident because someone modified the programming of the vehicle, they may not be willing to spend the right amount of resources to prevent tuning up. But if they are held responsible for not preventing speed manipulation, they will probably also be willing to do whatever is required to limit the extent of tuning up.

In summary, it might be difficult to stop all unrightful overtaking of autonomous vehicles, especially by people who have access to the vehicle owner's property. Still, if the manufacturers spend enough resources to prevent manipulation of vehicles' speed, the worst consequences (the killing of the child) can be prevented.

- Kidnapping for ransom

In this scenario, the kidnappers gain private private data about mobility patterns when preparing the crime. An important question is therefore whether the developers will choose to prioritise protection of private data. And will the manufacturers be held responsible if such data is unrightfully obtained?

For manufacturers to be held responsible it must be discovered that private data about mobility has been accessed. In many situations such information leakage can be difficult to discover. And even if data obtained by hacking is used to conduct a crime against a person or property, it might still be a problem of proving that the crime in question was made possible by illegal access to private data. It is therefore uncertain whether the manufacturers will be willing to spend the necessary resources on protecting the vehicle users' private data.

In the kidnapping scenario, the kidnappers install a mobile jammer under the vehicle. Such jammers already exist on the market and are relatively inexpensive.

In addition, the kidnappers gain access to the control system and the locking system. The control system is vulnerable because it is remote-controlled. It is plausible that it will be possible to spoof the real owner's signals for capable hackers. The locking system is also vulnerable because modern vehicles are keyless and that makes them easier to open and lock for malicious actors (see discussion of vehicle theft in the above section about the playful teenagers scenario).

In summary, it seems plausible that organised criminals can exploit new vulnerabilities caused by autonomous vehicles and more easily organise kidnappings without being caught.

- Domestic abuse

In the domestic abuse scenario, the IT-expert gains unrightful access to the control, and locking systems. The challenges of securing the control and locking systems against cyberattacks have been discussed in the above scenarios of playful teenagers and kidnapping for ransom.

In addition, the IT-expert gains control of the heating system for the purpose of increasing the ex-lover's suffering during the drive. Because of the convenience of entering a vehicle that has a comfortable temperature, the heating system will probably also be remotely controlled, and thus vulnerable to hacking. A relevant question is whether developers will prioritise to protect the heating systems from hacking. Since the main purpose of hacking the heating system must be to make travelling less comfortable (and not financial gain et cetera), fewer hackers will be motivated to conduct such an attack and less security is needed. There might, however, exist manual solutions protecting against cyberattacks on the heating system. For example, a manual switch that disconnects the heating systems from all networks.

In summary, it seems plausible that domestic abusers with the right technical knowledge would be able to conduct such an attack as described in this scenario.

- Autopilot manipulation

Artificial intelligence is vulnerable and early versions of autopilot have demonstrated some vulnerabilities. The autopilot of Tesla brands has repeatedly misinterpreted the broad side of large white semi-trailers as the bright sky, leading to two fatal traffic accidents with three years in between (Hawkins, 2019). Since the problem of misinterpretation still exists despite the seriousness of the first accident, the problem is seemingly difficult to solve. It is also reasonable to assume that AIs have many more vulnerabilities that will be discovered when autopilots/autonomous vehicles become more common. Such vulnerabilities can be exploited in protests as described above.

Some researchers have actually succeeded in tricking a Tesla's autopilot into driving into oncoming traffic. They used a can of paint and a brush to trick the Enhanced Autopilot of a Model S 75 into detecting and then following a change in the current lane onto the lane with opposite traffic (Tencent Keen Security Lab, 2019). Developers are undoubtedly doing their best to improve the autopilot and eliminate vulnerabilities. It is, however, uncertain whether it is possible to eliminate all such vulnerabilities.

In summary, the developers should be motivated to prevent autopilot manipulation, but it is still uncertain whether it will be possible to prevent all such attacks.

- Paralysation of transport system

In the scenario, the attacker compromises all vehicles of a specific brand through the over-the-air update. The manufacturer should be motivated to prioritise investing in cyber security of update servers, and it is therefore not plausible that actors with medium-to-low capacity will manage to hack an update server. However, even if the manufacturers prioritise cyber security, it will be difficult to stop all high-capacity attackers. Even high-tech companies, such as Apple, Google and Microsoft, have experienced that their software update infrastructure has been hacked (Schmitt, 2019). Hence, a high-capacity attacker will probably be able to paralyse the transport system.

- Summary

In this analysis, we have included five scenarios, four scenarios where the attacker does only need limited capacity and one scenario only realistic to conduct successfully for high-capacity attackers.

  - The discussions of scenarios "Playful teenagers", "Kidnapping for ransom" and "Domestic abuse" demonstrate that it might be difficult to stop all illegal takeovers of autonomous vehicles, especially by people who have access to the vehicle owner's property.
  - The fourth scenario "Autopilot manipulation" requires very little technical knowledge to conduct successfully and it is uncertain whether it will be possible to remove all possibilities of autopilot manipulation.
  - The last scenario "Paralysation of the transport system" require high-capacity attackers and attempts of such attacks will therefore probably be rare. To prevent all such attacks, however, will probably be impossible.

5.6.6 Assessment of pure risk

The assessments of pure risk are based on the asset identification, threat assessment and vulnerability assessements. For each scenario, we ask:

1. Is it plausible that there are people who are motivated to conduct such cyberattacks?
2. Is it plausible that anyone motivated to conduct such an attack also has capacity to implement it?
3. What is the consequences of the scenario and how does it influence the assets?

- Playful teenagers

Remotely controlled vehicles are a popular toy and it seems plausible that many people would want to use autonomous vehicles as a toy to play with. It is also plausible that people without rightful access to a vehicle will be motivated to try to gain access to it, since car thefts for the purpose of joyriding is already common. With access to the vehicle owner's home, it might also be possible to gain control over the vehicle without much technical expertise.

In this scenario, two assets are damaged: the child and the vehicle. Another possible consequence is that people start to fear autonomous vehicles and therefore do not let their children go out and play.

In summary, it might be difficult to stop all attempts at gaining control over vehicles without rightful access, but it should be possible to limit the consequences; i.e. killings and resulting fear.

- Kidnapping for ransom

Many people may be motivated to obtain private data about mobility patterns, including criminals planning to commit a crime against a person or property. The developers, furthermore, lack sufficient incentives to prioritise protecting against such unrightful access to private data.

Kidnapping for ransom can be an attractive way of obtaining money for unscrupulous criminals if they think it is possible to get away without being caught. If autonomous and connected vehicles make it easier avoid being caught during and after a kidnapping attempt, the prevalence of such kidnappings will probably increase. Since it can be difficult to protect the control and locking systems from hacking, this might indeed happen.

In the kidnapping scenario, only the ten-year-old child is under physical threat. However, both the child and the parents (and maybe other relations involved) experience trauma. After the public is made aware about the kidnapping, people that perceive themselves as attractive targets for such crimes might start avoiding autonomous vehicles and instead choose to travel by older, non-autonomous vehicles.

In summary, autonomous and connected vehicles can enable organized criminals to more easily conduct successful kidnappings for ransoms. In addition to causing trauma among the directly affected, such kidnappings could make large groups of people avoid using them.

- Domestic abuse

There are no reliable prevalence data on domestic abuse but the Crime Survey of England and Wales estimates that 21% of people aged 16 to 59 years in England and Wales have experienced some form of domestic abuse since the age of 16 years (Office for National Statistics, 2018). Use of modern technology when conducting domestic abuse is not a new phenomenon; for years, a subset of abusive partners with technical knowledge have placed spyware on computers or mobile devices, stolen passwords, and generally kept tabs on their other half (Boyd, 2018). Hence, it is very plausible that domestic abusers with the right technical capability may choose to use autonomous vehicles as a tool for terrorising victims.

The physical threat to the ex-lover is limited, but the experience is traumatic. The news coverage of the incident reduces the confidence in autonomous vehicles, resulting in less use of the transport system and more use of older and less environmentally friendly vehicles.

In summary, domestic abuse is very common, but only some domestic abusers have the technical knowledge needed for hacking an autonomous vehicle. Access to the victim's home may, however, help a would-be hacker to obtain the necessary knowledge for conducting a cyberattack. Hence, as well as investing in security to protect the control, locking and heating systems, measures that support vehicle owners in protecting "inside information" about their vehicle might reduce cyberattacks relying on such inside information.

- Autopilot manipulation

Environmental groups engage in a mixture of political methods and activities, everything from confrontation and violence to more conventional styles of political persuasion (Dalton, Recchia and Rohrschneider, 2003). In many cases where more conventional styles of political persuasion fail, such groups may choose to conduct more irregular protests. Hence, it seems plausible that some environmental groups could conduct a campaign like the one described above.

In this scenario, many people suffer from minor injuries and a few people suffer from serious injuries. A couple of the seriously injured passengers never completely recover. Many vehicles are damaged. The incident also causes reduced confidence in autonomous vehicles, resulting in less use of the transport system and more use of older and less environmentally friendly vehicles.

In summary, political organisations sometimes choose to conduct irregular protests. Some of these organisations might choose to conduct protests by changing the physical environment to manipulate the autopilot of autonomous vehicles. Such an attack will only require limited technical knowledge. It is uncertain whether it is possible to create autopilots without vulnerabilities that can exploited in such an attack.

- Paralysation of transport system

Sometimes countries have outstanding issues that they experience as so salient that they prefer to go to open war. Conducting a cyberattack can be easier and less risky than military action. Cyberattacks can be an especially convenient alternative when the countries in question is located far from each other.

The sudden paralysation of the transport system causes many minor traffic accidents with resulting vehicle damage and personal injuries. The consequences of the continued blocking of traffic are more serious. Trade flows are disrupted, causing damage to property and obstructed delivery of time-critical goods. Emergency vehicles are either delayed or not able to arrive at all, causing fires to develop further and people not being rescued in time.

In summary, the scenario on paralysation of the transport system is plausible. States can conduct such cyberattacks as an attractive alternative to military action. Prioritising the protection of update servers from such attacks will make it less probable that hackers succeeds in attacking, but will not stop the most determined high-capacity attackers.

5.6.7 Risk situation

In this analysis, we have included five scenarios; "Playful teenagers", "Kidnapping for ransom", "Domestic abuse", "Autopilot manipulation", and "Paralysation of the transport system". Because of inherent vulnerabilities in connected and autonomous vehicles, it might be difficult to prevent all such hacker attacks as described in the scenarios. The two scenarios "Playful teenagers" and "Autopilot manipulation" can be conducted by people with little technical knowledge and limited resources. The scenarios "Kidnapping for ransom" and "Domestic abuse" require hackers with more technical knowledge. Only high-capacity actors, such as states and state-sponsored groups, will be able to conduct the attacks described in "Paralysation of the transport system".

Since it is not possible to prevent all hacker attacks, it is necessary to implement measures that reduce the consequences of the scenarios. Such measures include 1) kill switches that make it possible to turn off vehicle automation manually and thus override the autopilot (Consumer Watchdog, 2019)("Kidnapping for ransom" and "Domestic abuse"), 2) make the vehicle possible to move by four healthy adults when automation is turned off ("Paralysation of the transport system"), 3) protection of the occupants in traffic accidents ("Autopilot manipulation"), and 4) detection of installed that may cause accidents ("Playful teenagers").

5.6.8 Comparison of risk picture and security objectives

Below we discuss whether we believe the risk picture as described above will fulfill the security objectives as described in section 5.6.2.

1. The first incident-specific objective was that no cyberattack should cause more fatalities and seriously injured road users as defined by the first general security objective per year:
    - The first general objective was an absolute reduction in the number of fatalities and serious injuries compared to the current situation, and a reduction in the risk of being killed or seriously injured per kilometre travelled.
    - It seems improbable that any of the scenarios above should result in more fatalities and seriously injured road users than currently occur annually in traffic.
2. The second incident-specific objective is that no cyberattack should result in larger economic costs as defined by the second general security objective per year.
    - The second objective is that cyberattacks should not incur larger costs on society than the efficiency gains achieved because of the introduction of connected and autonomous vehicles.
    - It seems improbable that any of the scenarios above would involve larger costs on society than the cost savings produced by CATS, thus eliminating these gains.
3. The third incident-specific objective is that no cyberattack should cause more than 24 hours of downtime of the transport system.
    - If connected and autonomous vehicles all have a kill switch and most of the vehicles are movable by four healthy adults, it seems improbable that any cyberattack could cause more than 24 hours of downtime of the transport system.
4. The fourth incident-specific objective is that no cyberattack should delay emergency missions such that people die, fires develop further and conflict situations escalate into violence because of the delay.
    - If connected and autonomous vehicles all have a kill switch and most of the vehicles are movable by four healthy adults, the fourth incident-specific objective might be obtainable.
5. The fifth incident-specific objective is that 7 days after a cyberattack at least 90% of the population should be willing to use autonomous vehicles again.
    - The achievement of this objective depends on the frequency of incidents (as described in the scenarios), the consequences of the incidents and the media coverage of the incidents. The terrorist attacks on September 11, 2001, depressed air travel far longer than 7 days.
6. The sixth incident-specific objective is that 21 days after a cyberattack at least 80% of the population should report that they feel safe as passengers of autonomous vehicles.
    - The achievement of this objective depends on the frequency of incidents (as described in the scenarios), the consequences of the incidents and the media coverage of the incidents.

In summary, objectives 1 and 2 seem realistic and objectives 3 and 4 are obtainable if kill switches are installed and vehicles are easily movable by healthy adults. The achievement of objectives 5 and 6 depends both on the frequency and seriousness of all attacks and the resulting media coverage of these attacks.

6 Library of functional relationships with uncertainties

The objective of this chapter is to summarise and assess the uncertainty of the functional relationships between the market penetration rate of automated vehicles and impacts on mobility, safety and the environment estimated in chapter 4 and 5. All relationships are subject to a ceteris paribus clause, meaning that they apply when everything else is equal. Most of the functional relationships show primary impacts. It is currently possible to describe only a few secondary impacts in terms of functional relationships.

The objective of this chapter is to summarise the functional relationships between the market penetration rate of automated vehicles and the impacts of these vehicles. Three main types of impacts are included: mobility, safety and environment. Most relationships describe primary impacts. They represent a formal interpretation of the results, mainly of traffic simulation studies, presented and discussed in chapter 4. All impacts are uncertain. To indicate uncertainty, the range of findings reported in the literature is used as frame of reference. This approach is, however, only feasible when more than one study dealing with a certain impact has been reported.

6.1 Mobility impacts

All functions have been calibrated so that their value equals 100 at 0% market penetration rate of automated vehicles. Market penetration is stated as a percentage, in steps of 10 (10, 20, …, 90, 100). Maximum impact is reached at 100% market penetration. In all tables, the following notation is used:

e = the exponential function, i.e. the base of natural logarithms (2.71828) raised to the power of a term in parenthesis
p% = percent market penetration of automated vehicles; stated as a whole number (10, 20, etc.)
p%^{2} = the squared value of market penetration of automated vehicles
Max (100%) = the value of the impact (dependent variable) at 100% market penetration of automated vehicles

As an example, the first impact listed in Table 10 is change in lane capacity on motorways. This impact is modelled by the following function:

Change in lane capacity = $0.0135 \cdot x^2 - 0.1505 \cdot x + 100$

This function results in a value of 220 for lane capacity at 100% market penetration, i.e. an increase of 120 percent. Uncertainty is stated in the form of lower and upper bounds for impacts. In general, lower limit refers to the smallest impact, upper limit to the largest impact.

Table 10 Functional relationships describing impacts of CATS on mobility.

Impact	Best estimate	Lower limit	Upper limit
	Functions based on studies quoted in chapter 4		
Road (lane) capacity	$0.0135 \cdot p\%^2 - 0.1505 \cdot p\% + 100$ Max = 220	$100 \cdot e(p\% \cdot 0.003365)$ Max = 140	$0.028 \cdot p\%^2 - 0.1505 \cdot p\% + 100$ Max = 375
Junction capacity (urban road capacity)	$100 \cdot e(p\% \cdot 0.003)$ Max = 135.0	$100 + 0.2564 \cdot p\%$ Max = 125.6	$100 \cdot e(p\% \cdot 0.0036)$ Max = 143.3
Delays on motorways	$0.0011 \cdot p\%^2 - 0.4164 \cdot p\% + 100$ Max = 69.4	$0.0011 \cdot p\%^2 - 0.3 \cdot p\% + 100$ Max = 80.0	$0.0011 \cdot p\%^2 - 0.51 \cdot p\% + 100$ Max = 60.0
Delays in roundabouts	$0.0015 \cdot p\%^2 - 0.5 \cdot p\% + 100$ Max = 65.0	$0.0009 \cdot p\%^2 - 0.121 \cdot p\% + 100$ Max = 78.9	$0.0039 \cdot p\%^2 - 0.8662 \cdot p\% + 100$ Max = 52.4
Delays in signalised junctions	$0.001 \cdot p\%^2 - 0.2 \cdot p\% + 100$ Max = 90.0	$0.0005 \cdot p\%^2 - 0.1035 \cdot p\% + 100$ Max = 94.7	$0.0021 \cdot p\%^2 - 0.3592 \cdot p\% + 100$ Max = 85.1
Travel time on motorways	$100 \cdot e(p\% \cdot -0.001)$ Max = 90.5	$100 \cdot e(p\% \cdot -0.00062)$ Max = 94.0	$100 \cdot e(p\% \cdot -0.0023)$ Max = 79.5
	Functions derived from the quoted functions		
Delays in junctions (mean of signalised and roundabouts)	$0.0012 \cdot p\%^2 - 0.3444 \cdot p\% + 100$ Max = 77.6	$-0.00028 \cdot p\%^2 - 0.108 \cdot p\% + 100$ Max = 86.4	$0.0028 \cdot p\%^2 - 0.594 \cdot p\% + 100$ Max = 68.6
Travel time in cities	$100 \cdot e(p\% \cdot -0.00062)$ Max = 94.0	$100 \cdot e(p\% \cdot -0.0003)$ Max = 97.0	$100 \cdot e(p\% \cdot -0.00128)$ Max = 88.0

The upper and lower limit functions have been determined so that they encompass all or most of the data points in primary studies. Figure 28 gives an example of this. It shows the best estimated function and upper and lower limit functions for lane capacity on motorways. It is seen that the upper and lower function diverge more and more as 100% market penetration is approached. This means that uncertainty increases as market penetration approaches 100 percent. This is regarded as reasonable.

It is seen that most data points are located inside the lower and upper limit functions. Using more complex functional forms for the upper and lower limit functions might have brought the data points located outside the functions inside these functions, but the fitted upper and lower limit functions were judged as sufficiently accurate to justify parsimony and simplicity rather than ad hoc complexity designed to bring the few data points that are outside the upper and lower limit curves inside these curves.

In addition to the functional relationships estimated on the basis of published studies, Table 10 contains a couple of functions derived from the functions quoted in Chapter 4. The first of these functions refers to delays in junctions. It is a simple average of the functions for signalised junctions and roundabouts. Both types of junctions are found in cities and their exact mix varies from city to city. Without specific data, a simple average was chosen.

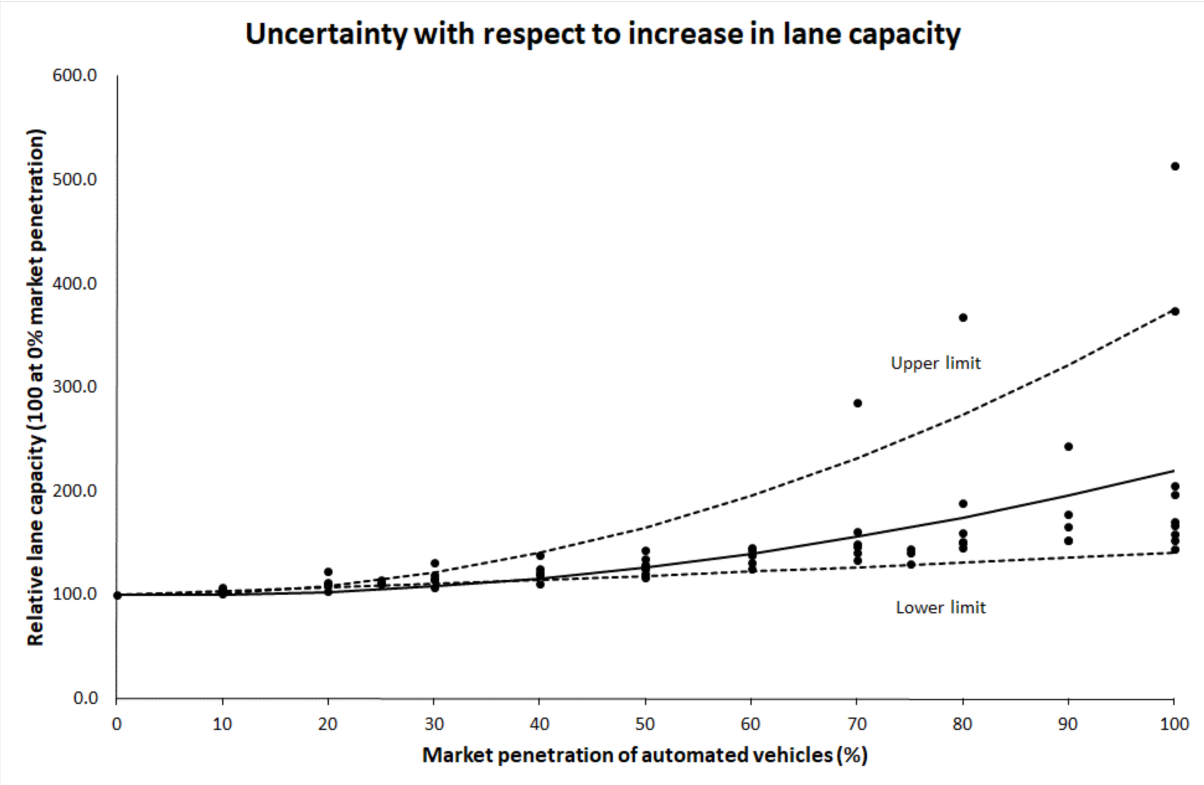

Figure 28: Functions indicating increase in lane capacity and uncertainty of the increase

The other function refers to travel time saving in cities. For motorways, it can be observed that delays, according to the best estimate, are reduced by a little more than 30% at 100% market penetration of automated vehicles. However, travel time on motorways is estimated to be reduced by only 9.5% at 100% market penetration of automated vehicles. The difference between these estimates can be explained in terms of congestion. Delays are caused by congestion, but not all of traffic is congested. When traffic flows freely, there are no delays. Since only about 25-30% of traffic is congested, the mean gain in travel time is considerably smaller than the reduction of congestion delays.

It is reasonable to assume that the delays in junctions are a characteristic of urban traffic, as the density of junctions is much higher in cities than in rural areas. The best estimate of the mean reduction of delay in junction, according to the average function fitted as an average of signalised junctions and roundabouts, is 22.4% reduction at 100% market penetration of automated vehicles. By analogy to the results for motorways, it will be assumed that travel time gain in urban areas will be about 25-30% of the reduction in delays, i.e. about 6% at 100% market penetration of automated vehicles. Lower limit is 3% saving in travel time; upper limit is 12% saving in travel time. See the functions in the bottom row of Table 10.

In general, there is considerable uncertainty. As an example, the upper limit value for increase in lane capacity at 100% market penetration is an increase of 275%, whereas the lower limit is an increase of only 40%.

6.2 Safety impacts

Table 11 presents functions describing the road safety impacts of connected and automated vehicles.

Table 11: Functional relationships describing impacts of CATS on safety.

Impact	Best estimate	Lower limit	Upper limit
	Functions based on studies quoted in chapter 4		
Rearend and lane change collisions (motorways)	$0.0035 \cdot p\%^2 - 1.2868 \cdot p\% + 100$ Max = 6.3	$0.005 \cdot p\%^2 - 1.17 \cdot p\% + 100$ Max = 33.0	$0.0035 \cdot p\%^2 - 1.34 \cdot p\% + 100$ Max = 1.0
Accidents in signalised junctions	$100 \cdot e(p\% \cdot -0.007)$ Max = 49.7	$100 \cdot e(p\% \cdot -0.0011)$ Max = 89.6	$100 \cdot e(p\% \cdot -0.02)$ Max = 13.5
Accidents in priority junctions	$-0.001 \cdot p\%^2 - 0.12 \cdot p\% + 100$ Max = 78.0	$-0.0054 \cdot p\%^2 + 0.412 \cdot p\% + 100$ Max = 87.2	$-0.0035 \cdot p\%^2 + 0.042 \cdot p\% + 100$ Max = 69.2
Accidents in roundabouts (single study only)	$-0.009 \cdot p\%^2 + 0.3022 \cdot p\% + 100$ Max = 40.2	$-0.0088 \cdot p\%^2 + 0.32 \cdot p\% + 100$ Max = 44.0	$-0.0092 \cdot p\%^2 + 0.29 \cdot p\% + 100$ Max = 37.0
Cyclist and pedestrian accidents	$0.0059 \cdot p\%^2 - 1.1562 \cdot p\% + 100$ Max = 43.4	$0.0059 \cdot p\%^2 - 1.22 \cdot p\% + 100$ Max = 37.0	$0.0059 \cdot p\%^2 - 1.12 \cdot p\% + 100$ Max = 47.0
	Functions derived from the quoted functions		
Urban junctions (mean of signals, priority, roundabouts)	$-0.0028 \cdot p\%^2 - 0.1642 \cdot p\% + 100$ Max = 55.6	$-0.0048 \cdot p\%^2 + 0.1974 \cdot p\% + 100$ Max = 71.7	$-0.0041 \cdot p\%^2 - 0.4213 \cdot p\% + 100$ Max = 43.9

The functions for roundabouts are based on the study by Morando et al. (2018). The study provided confidence intervals. These were applied to indicate uncertainty. It should be noted that this resulted in a narrow range of outcomes, which most certainly underestimates uncertainty, given the fact that when multiple studies are available, their results tend to differ considerably.

A mean function for junctions was estimated based on the functions for signalised junctions, priority junctions and roundabouts. Each type of junction had a weight of 1/3.

6.3 Fuel consumption

Table 12 shows the estimated function for fuel consumption and an estimate of its uncertainty.

Presumably, the studies of fuel consumption assume that vehicles are still operated by internal combustion engines running on fossil fuel. As an approximation, it may be assumed that emissions of local pollution per vehicle are proportional to fuel consumption. CO_2 emissions are also proportional to fuel consumption.

Table 12: Functional relationships describing impacts of CATS on fuel consumption.

Impact	Best estimate	Lower limit	Upper limit
	Functions based on studies quoted in chapter 4		
Fuel consumption	$0.0009 \cdot p\%^2 - 0.3948 \cdot p\% + 100$ Max = 69.5	$0.0009 \cdot p\%^2 - 0.75 \cdot p\% + 100$ Max = 34.0	$0.004 \cdot p\%^2 - 0.68 \cdot p\% + 100$ Max = 72.0

If vehicles become electric, emissions become zero. It is important to note that this does not mean that a zero emission vehicle is a zero pollution vehicle. An electric vehicle will produce local air pollution by producing microparticles that are suspended in air and can be inhaled. These particles are produced by the friction between tyres and road surface as well as by the wear of the mechanical parts of the braking system.

In a study of the health impacts of local transport in Warsaw, Poland (Tainio 2015), microparticles had by far the largest health impact of any component of air pollution, producing health losses that were considerably larger than those associated with pollutants that have traditionally been attributed to fossil fuels, like NO_x or SO_2.

6.4 Other impacts and feedback impacts

It is not known when or how fast connected and automated vehicles will penetrate the market. However, given the fact that a passenger car today has a service life of some 15-20 years (the mean age at scrapping varies between countries; in countries where cars are comparatively expensive, like Denmark or Norway, people keep them longer than in countries where cars are cheaper), it is not unreasonable to assume that full market penetration may take some 15-20 years.

Inspired by the logistic functions developed by Rosen and Sander (2009) to model the probability of death as a function of impact speed in pedestrian accidents, similar logistic functions have been developed for the market penetration of automated cars. The functions are shown in Table 13 and Figure 29.

Table 13: Functional relationships describing market penetration of automated vehicles.

Impact	Best estimate	Lower limit	Upper limit
	Market penetration rate		
Market penetration by year (years = 0, 1, 2, etc.)	$100/(1 + e(6.9 - 0.9 \cdot year))$ 100 % in 17 years	$100/(1 + e(8.5 - 0.6 \cdot year))$ 100 % in 27 years	$100/(1 + e(5.3 - 1.2 \cdot year))$ 100 % in 11 years

Full market penetration takes 17 years according to the best estimate, 27 years as a pessimistic alternative and 11 years as an optimistic alternative.

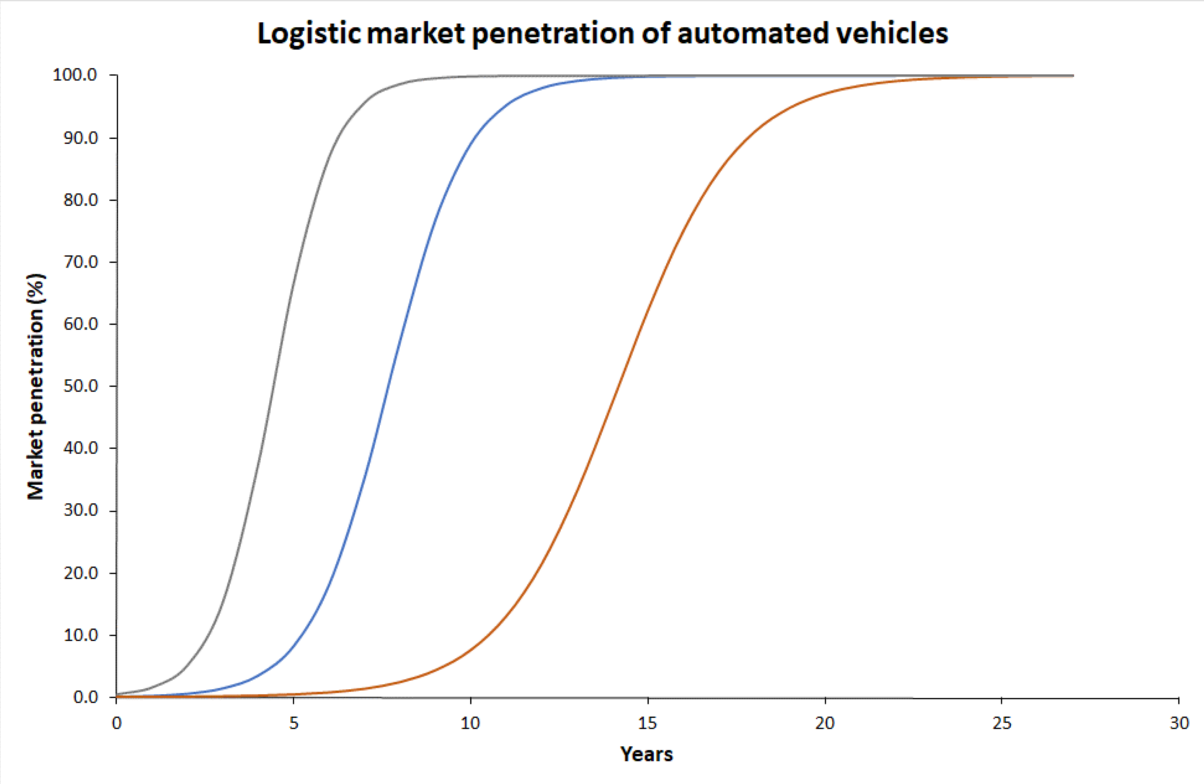

Figure 29: Market penetration curves for automated vehicles

The discussion of feedback impacts in Chapter 5 concluded that travel demand induced by increased road capacity is sufficiently well known to be estimated. This impact is also important, as all other impacts, like savings in travel time or accidents, depend strongly on changes in traffic volume. It was concluded that the best estimate of induced travel demand is 75% of the increase in road capacity, with a lower bound of 50% and an upper bound of 100%. Table 10 presented dose-response curves for lane capacity on motorways and junction capacity, with associated upper and lower limits. Based on these functions, Table 14 presents functions for induced travel demand. Given that there is uncertainty both about the increase in road capacity and the amount of induced travel, the functions have been defined as follows:

1. The best estimate is a function in which travel demand increases by 75% of the best estimate of the increase in road capacity.
2. The lower limit is a function in which travel demand increases by 50% of the lower limit of the increase in road capacity (there is logic in this: a smaller increase in road capacity may generate less new traffic than a large increase in road capacity).
3. The upper limit is a function in which the increase in travel demand equals the upper limit of the increase in road capacity.

It is seen that the range of induced travel demand on motorways is between 20% and 375%. In cities, the corresponding range is between 13% and 43%.

Table 14: Functional relationships describing induced travel demand.

Impact	Best estimate	Lower limit	Upper limit
	Functions based on studies quoted in chapter 5		
Travel demand on motorways	$0.0105 \cdot p\%^2 - 0.1505 \cdot p\% + 100$ Max = 190	$100 \cdot e(p\% \cdot 0.00182)$ Max = 120	$0.028 \cdot p\%^2 - 0.1505 \cdot p\% + 100$ Max = 375
Travel demand in cities (junctions)	$100 \cdot e(p\% \cdot 0.0023)$ Max = 126.2	$100 \cdot e(p\% \cdot 0.00121$ Max = 112.9	$100 \cdot e(p\% \cdot 0.0036)$ Max = 143.3

A final potential impact to be discussed is behavioural adaptation to automated vehicles among drivers of non-automated vehicles. Only one driving simulator study has found this: drivers of manual vehicles copied the shorter headways of platooned automated vehicles. The likelihood and magnitude of this copying behavior will probably depend on the market penetration of automated vehicles. At low penetration, manual vehicles will dominate traffic and there will be few platoons of automated vehicles benefitting from short headways. As automated vehicles become more common, platoons characterised by short headways will be observed more often. Finally, as automated vehicles dominate, copying their shorter headways may decline again, as the last drivers to adopt automated vehicles may be those who are least interested in the benefits offered by the technology and may prefer a defensive driving style not giving high priority to short headways.

A function consistent with these hypotheses will be bell-shaped, rising to a maximum and then declining again. Adding such a function to the original function results in the new functions shown in Table 15. It is recognised that the function representing behavioural adaptation is somewhat speculative.

Table 15: Functional relationships describing changes in conflicts on motorways – allowing for behavioural adaptation.

Impact	Best estimate	Lower limit	Upper limit
	Changes allowing for behavioural adaptation		
Rearend and lane change conflicts	$-0.0001p\%^3 + 0.0129p\%^2 - 1.227p\% + 100$ Max = 6.3	$-0.0001p\%^3 + 0.0135p\%^2 - 1.02p\% + 100$ Max = 33.0	$-0.0001p\%^3 + 0.0135p\%^2 - 1.34p\% + 100$ Max = 1.0

7 Methods for predicting impacts

The objective of this chapter is to show how to use different methods for predicting impacts of connected and automated vehicles. The application of each method consists of several steps. The first step is to define a use case or a type of traffic environment for which impacts are to be predicted. The second step is to decide which impacts are to be included in analysis. The third step is to estimate primary impacts. To this end, the functional relationships described in Chapter 6 can be applied. The fourth step is to estimate rebound effects, including behavioural adaptation. The fifth step is to estimate net impacts when primary impacts have been corrected for rebound effects. The sixth step is to assess uncertainty and develop best and worst case scenarios. Examples of how to use the methods are given.

7.1 Defining use cases or systems for impact analysis

Table 16 summarises the use cases that have been proposed in Levitate. The body of the Table lists examples of technologies associated with each scenario. Thus, for automated urban transport, fully automated vehicles that are able to operate in an urban traffic environment are needed. The vehicles will have technology permitting connectivity and can be integrated into mobility as a service.

The passenger car use case will comprise automation in all traffic environments. The studies reviewed for this use case deal with automation levels 4 or 5, or do not specify the automation level.

The freight transport use case also comprises all traffic environments. Estimates of some impacts are available for some of the technologies listed. For other technologies, the review presented in Chapter 4 has not identified any estimates of impacts. It should be stressed, however, that the review presented in Chapter 4 looked for studies permitting the estimation of dose-response curves for as many potential impacts of connected and automated vehicles as possible, and not for studies of the impacts of specific technologies. Additional literature surveys made as part of work packages 5, 6 and 7 will look for studies of impacts of the specific technologies listed in Table .

For all use cases, a distinction should be made between three types of traffic environment:

1. An urban traffic environment
2. A rural traffic environment
3. Motorways

The potential impacts of vehicle automation vary with respect to the type of traffic environment. It is therefore necessary to develop separate estimates of potential impacts for each traffic environment.

Table 16: Summary of use cases proposed in Levitate.

Use-case	Sub-use cases
Passenger cars	Road use pricing: • Empty km pricing • Static toll on non-automated vehicles • Static toll on all vehicles • Dynamic toll on non-automated vehicles • Dynamic toll on all vehicles
	Automated ride sharing
	Reduction of parking space: • Replace with public space • Replace with driving lanes • Replace with short-term parking
	Provision of economic incentives for buying Avs
	Provision of dedicated lanes for AVs on urban highways
	Highway platooning
Use-case	**Sub-use cases**
Urban transport	Automated shuttles: • Point to point shuttles • Anywhere to anywhere shuttles • Last mile shuttles
Use-case	**Sub-use cases**
Freight transport	Automated urban delivery
	Local freight consolidation
	Hub to hub automated transfer
	Highway platooning

7.2 Defining relevant impacts

A total of 33 potential impacts of connected and automated vehicles have been identified. Not all of these impacts will be relevant for all use cases. To identify the most relevant impacts, the policy objectives set by stakeholders, such as city governments, public transport operators, car manufacturers or freight transport operators will be used as a frame of reference. All stakeholders primarily want to know whether they have realised their policy goals or not. Any policy may, obviously, have unintended side-effects in addition to its declared objectives. If there is knowledge permitting such additional impacts to be identified, they should be included in an impact assessment.

A typical example of un-intended impacts is rebound impacts, or behavioural adaptation to technological changes. If, as is widely assumed in the literature, automated transport systems will reduce the generalised costs of travel, it is highly likely that this will induce more travel. The reduced cost of travel may not just cause more trips to be made, but also longer trips and, in the long term, urban sprawl. These impacts would normally be regarded as undesirable by, for example, a city government. Yet, including them is important, as it is to some extent possible to take action to counteract the rebound impacts. One may, as an example, introduce road pricing as a form of tax to prevent transport from becoming cheaper.

Estimates of impacts will therefore be developed both with and without inclusion of any effects of behavioural adaptation, in order to see how these impacts influence the total benefits of connected and automated vehicles.

In this chapter, we will see how far one can get by applying the dose-response curves presented in Chapter 4, while recognising that the curves will not include all relevant impacts. Impacts not covered by the dose-response curves represent research needs, some of which may be covered by additional research made as part of Levitate.

To illustrate the use of the dose-response curves to estimates impacts, it will be assumed that all cars are automated. It will be assumed that mopeds and motorcycles remain operated by a human driver. Cyclists and pedestrians will also remain non-automated (although, in principle, one could make it a requirement for cyclists and pedestrians to wear a communication device capable of communicating with automated vehicles; it will be assumed that no such requirement is made).

7.3 Case illustration: Automation of all cars

The replacement of human operated cars by fully automated cars at the highest level of automation represent the broadest scope of a use case or scenario for the introduction of connected and automated vehicles and may therefore be interpreted as an attempt to estimate the maximum potential impacts of connected and automated vehicles. As explained above, impacts will be estimated separately for urban, rural and motorway environments.

7.3.1 Identifying accidents involving road users who will not be automated

Mopeds, motorcycles, cyclists and pedestrians will not be automated. Although automation technology may prevent many accidents involving, for example, cars and cyclists, it will not prevent accidents involving only non-automated road users. Based on Norwegian accident statistics (each country would need to produce its own statistics), it can be estimated that injuries involving non-automated road users represent 7% of all injuries in urban areas, 4% in rural areas and 1% on motorways. These injuries will not be affected by automation technology and it is assumed that their number remains unchanged. Automation technology will therefore potentially influence 93% of injured road users in urban areas, 96% in rural areas and 99% on motorways.

7.3.2 Primary and net impacts on road capacity

In the urban traffic environment, it will be assumed that road capacity depends on junction capacity. Estimates of potential changes in junction capacity are only available from a single study (Atkins 2016) and only for roundabouts and signalised junctions. It will be assumed that the function for roundabouts can also be applied for priority junctions. In chapter 6, a functional relationship for all junctions was developed as a mean of the relationships for signalised junctions and roundabouts. This functional relationship has been applied below in order to estimate the change in road capacity in urban areas associated with vehicle automation.

The best estimates of the primary impact and net impact after adjusting for the rebound impact on junction capacity under these assumptions is:

	Relative change in road capacity	
Market penetration of automated vehicles (X)	Primary impact	Net impact
0	100.0	100.0
10	103.0	100.7
20	106.2	101.3
30	109.4	102.0
40	112.7	102.7
50	116.2	103.4
60	119.7	104.1
70	123.4	104.8
80	127.1	105.5
90	131.0	106.2
100	135.0	106.9

Current road capacity is given the value of 100. The numbers represent percentage changes. Without any growth in traffic volume, a maximum increase in road capacity of 35% at 100% market penetration is predicted. It has been assumed, see the discussion in section 5.1, that 75% of added capacity will be consumed by induced travel. Increase in travel at 100% market penetration of automated vehicles will be 75% of 35%, or 26.2%. It is assumed that impacts combine multiplicatively. Thus, the net impact at 100% market penetration will be an increase of just 6.9% in road capacity (135.0/126.2 = 106.9). It is thus seen that the assumptions made about behavioural adaptation to increased road capacity, in the form of an increased amount of travel, has a large influence on the results.

7.3.3 Primary and net impacts on travel time

A functional relationship describing changes in travel time, all else equal was estimated in Chapter 6. This functional relationship applied to mean travel time both in congested and uncongested conditions, i.e. mean travel time for a 24-hour period. Estimates are once again developed under the assumption of no change in travel demand and under the assumption that increased travel will reduce the increase in road capacity by 75%. Even after increase in travel demand, a net increase in road capacity remains. The studies reviewed in Chapter 6, suggest that the saving in mean travel time is a function of the increase in road capacity as follows:

Saving (%) in mean travel time = $0.0006X^2 - 0.1786X - 0.089$

Applying this function (omitting the constant term) to the net increase in road capacity listed above, the results are:

	Relative change in travel time	
Market penetration of automated vehicles (X)	Primary impact	Net impact
0	100.0	100.0
10	99.4	99.3
20	98.8	98.6
30	98.2	97.8
40	97.6	97.1
50	96.9	96.3
60	96.3	95.6
70	95.8	95.0
80	95.2	94.2
90	94.6	93.5
100	94.0	92.8

Current travel time is given the value 100. The other entries show travel time relative to this value. Given the net increase in road capacity, a net reduction in travel time of about 7% is estimated. In this case, the net impact is a little larger than the primary impact.

7.3.4 Primary and net impacts on fuel consumption and emissions

The primary impacts on fuel consumption and emissions were estimated by applying the function developed in Chapter 6.

	Relative change in fuel consumption and emissions	
Market penetration of automated vehicles (X)	Primary impact	Net impact
0	100.0	100.0
10	96.1	98.4
20	92.5	96.9
30	89.0	95.4
40	85.6	94.0
50	82.5	92.7
60	79.6	91.5
70	76.8	90.4
80	74.2	89.4
90	71.8	88.6
100	69.5	87.7

Net impacts were estimated by multiplying primary impacts with the expected growth in traffic volume. It is seen that net impacts are smaller than primary impacts.

7.3.5 Primary and net impacts on accidents

Estimating impacts on safety is more complicated than estimating effects on capacity and travel time. There are three reasons for this:

1. There are different dose-response curves for different types of accident or different groups of road users.
2. There is a category of accidents that will not be affected by vehicle automation, i.e. accidents involving vehicles and road users that will not be automated.
3. To estimate net impact, one should adjust for the expected increase in traffic volume.

A further source of complexity, is the incomplete reporting of injuries to cyclists and pedestrians in official road accident statistics. However, since official road accident statistics is by far the most commonly used source of data in road safety impact assessments, it will be used in the example developed below. In Chapter 6, a mean impact function for accidents in junctions was developed. This function will be applied to accidents in which occupants of motor vehicles are injured. Based on accident statistics for the city of Oslo – this must be customised to statistics for each city – the function for junction accidents has been given a weight of 0.6 and the function for cyclist and pedestrian accidents a weight of 0.4. The weighted mean curve applies to all accidents in urban areas. 7 percentage points were added to each data point of the weighted mean curve (except the data point for 0% market penetration) to account for the fact that 7 percent of the current number of injuries are sustained by non-automated road users and will not be influenced by automation. These injuries are assumed to remain unchanged as the market penetration of automated vehicles increases. Finally, the dose-response curve was adjusted for expected increase in traffic volume. This was done by multiplying each data point by the expected growth in traffic. This assumes that the number of accidents is proportional to traffic volume. This is most likely not correct, but it is the simplest assumption that can be made.

The results are shown below:

Market penetration (%)	Mean curve with 7% added to each entry	Mean curve adjusted for traffic increase
0	100.0	100.0
10	101.5	103.8
20	96.1	100.6
30	90.8	97.4
40	85.6	94.0
50	80.7	90.6
60	75.8	87.2
70	71.1	83.7
80	66.5	80.1
90	62.0	76.5
100	57.7	72.8

It is seen that adjusting for an expected increase in traffic has a major impact on the size of the estimated changes in the number of accidents. At 100 % market penetration of CATS, the estimated accident reduction is about 42% with constant traffic volume but only about 27% if traffic increases in line with historic demand elasticities for road capacity.

7.3.6 Summary of estimated impacts for urban areas

Table 17 summarises estimated net impacts of connected and automated vehicles in urban areas.

Table 17: Summary of estimated net impacts of connected and automated vehicles in urban areas. Relative changes from a baseline 0% market penetration.

Market penetration of automated vehicles (%)	Traffic volume	Road capacity	Travel time	Fuel consumption	Number of accidents
0	100.0	100.0	100.0	100.0	100.0
10	102.4	100.7	99.3	98.4	103.8
20	104.8	101.3	98.6	96.9	100.6
30	107.2	102.0	97.8	95.4	97.4
40	109.8	102.7	97.1	94.0	94.0
50	112.4	103.4	96.3	92.7	90.6
60	115.0	104.1	95.6	91.5	87.2
70	117.7	104.8	95.0	90.4	83.7
80	120.5	105.5	94.2	89.4	80.1
90	123.3	106.2	93.5	88.6	76.5
100	126.2	106.9	92.8	87.7	72.8

The impacts it has been possible to estimate by applying the dose-response curves are limited to those occurring within the traffic system. The estimated impacts can, as an approximation, be assumed to be invariant with respect to whether individual mobility will remain dominant or a shift to shared mobility will occur. As noted in Chapter 4, if there is a shift to shared mobility, there will be fewer cars than today, but each car will be operated for a longer period, implying that the number of cars moving on roads may not change very much from today.

The impacts presented in Table 17 are net impacts. The difference between primary impacts and net impacts is that estimates of primary impacts assume unchanged traffic volume, whereas estimates of net impacts include the growth in traffic volume. Since automated vehicles are widely believed to lower the costs of travel and transport, and since additions to road capacity tend to be filled up by induced traffic, it has been assumed that there will be traffic growth. According to the conventional approach to cost-benefit analysis, a growth in traffic is a benefit; it increases the consumer surplus of travel (more travel is served at lower cost). It does, however reduce the benefits in terms of total fuel consumption and total number of accidents compared to a situation with no traffic growth. There is less reduction in fuel consumption and less reduction in the number of accidents if traffic increases than if it does not. It is, a priori, not clear whether there will be any difference in total benefits between estimates based on primary impacts and estimates based on net impacts. From a theoretical point of view, basing cost-benefit analyses on net impacts would be the most correct approach.

Since the generalised costs of travel are normally taken to include the valuation of travel time, it might be double counting if the gain in travel time listed in Table 17 was to be added as an additional benefit. The costs of accidents are, however, to some extent external, i.e. not included in the generalised costs of travel and might therefore be added as an additional benefit.

The estimated net impacts at 100% market penetration of connected and automated vehicles are quite modest. According to the estimates, there will be a modest net increase in road capacity and a modest reduction of travel time. There will be a small reduction in fuel consumption and a somewhat larger reduction in the number of accidents. The estimated reduction in the number of accidents is not larger than what one might accomplish by means of a combined use of traditional road safety measures, such as traffic calming, reducing speed limits and converting junctions to roundabouts.

All estimates are uncertain. Before discussing uncertainty, an example of estimated impacts of connected and automated vehicles in a motorway environment will be given.

7.3.7 Summary of estimated impacts for motorways

Table 18 summarises estimated impacts of vehicle automation for motorway operations. The increase in road capacity, as shown in Figure 7, was first estimated. It was assumed that 75% of the increase in road capacity will be consumed by increased traffic. For travel time savings, the function shown in Figure 12 was used. It was assumed that traffic growth would eliminate 75% of the saving in travel time. For safety impacts, the functional relationship shown in Table 18 was applied. It was adjusted to account for the fact that 1% of accidents on motorways will involve non-automated road users or modes. It was also adjusted for expected growth in traffic.

It is seen that traffic is expected to grow by about 90% at 100 % market penetration of automated vehicles. There will nevertheless be a net increase in road capacity of about 16%. Travel time savings, after adjusting for traffic growth, will be about 15% at 100% market penetration. Fuel consumption is expected to increase a little; a result consistent with Chen et al (2019). A large reduction of the number of accidents is predicted. This is plausible, as most accidents on motorway are more easily prevented by automation technology than accidents in more complex traffic environments. On motorways, there are no pedestrians or cyclists, no slow-moving motor vehicles, no sharp curves, no at-grade junctions, and, in general a high standard of maintenance for traffic control devices and the road surface. Sensor technology is well capable of detecting nearby vehicles when lane changes take place and well capable of intelligent cruise control in vehicle following situations. Estimated effects on accidents have been adjusted for traffic growth and, admittedly somewhat speculatively, for behavioural adaptation among non-automated vehicles in terms of shorter headways. Effects on accidents are highly non-linear and large reductions are only seen when the market share of automated vehicles reaches 70% or more.

Effects on rural roads are likely to be somewhere in-between those estimated for urban roads and those estimated for motorways.

Table 18: Summary of estimated net impacts of connected and automated vehicles on motorways. Relative changes from a baseline 0% market penetration.

Market penetration of automated vehicles (%)	Traffic volume	Road capacity	Travel time	Fuel consumption	Number of accidents
0	100.0	100.0	100.0	100.0	100.0
10	99.5	100.3	98.7	95.7	89.5
20	101.2	101.2	97.3	93.6	81.8
30	104.9	102.6	95.8	93.4	76.7
40	110.8	104.3	94.2	94.9	73.2
50	118.7	106.3	92.7	98.0	70.3
60	128.8	108.4	91.2	102.4	67.0
70	140.9	110.4	89.6	108.2	61.6
80	155.2	112.4	88.1	115.1	52.5
90	171.5	114.2	86.7	123.1	37.3
100	190.0	115.8	85.3	132.1	13.0

7.4 Sensitivity analysis: worst and best outcomes

All impacts of connected and automated vehicles are uncertain. To show uncertainty, estimates of impacts have been developed based on the lower and upper limits of outcomes as specified in the functional relationships in Chapter 6. In general, the lower limit of an impact shows the smallest impact and the upper limit of an impact shows the largest impact. The worst case is that all impacts take on their smallest values, the best case is that they take on the largest values. However, as impacts go in opposite directions, to show the range of possible outcomes, it is necessary to consider the direction of impacts. The worst outcome will therefore be defined as follows:

1. There is full compensation for the maximum increase in road capacity, i.e. no net increase in road capacity; any increase is fully filled up by increased traffic volume, combined with:
2. There are no changes in travel time, as the ratio of volume to capacity is unchanged, i.e. there is no change in congestion (Kucharski and Drabicki 2017),
3. The lowest estimate for reduction of fuel consumption,
4. The lowest estimate for reduction of the number of accidents.

Estimated impacts for the worst and best outcome will be illustrated for the urban case. Table 19 shows the estimated worst outcome for urban areas. There is considerable traffic growth; no net increase in road capacity and hence no change in travel time (congestion will remain at the same level as before). There is a temporary reduction in fuel consumption; however, traffic growth ultimately eliminates this reduction. A small reduction of the number of accidents has been estimated when the market penetration of automated vehicles reaches 80% or more.

Table 19: Worst case outcome for urban areas. Relative changes from a baseline 0 % market penetration.

Market penetration of automated vehicles (%)	Traffic volume	Road capacity	Travel time	Fuel consumption	Number of accidents
0	100.0	100.0	100.0	100.0	100.0
10	103.7	100.0	100.0	97.0	107.0
20	107.5	100.0	100.0	94.6	106.8
30	111.4	100.0	100.0	92.7	106.3
40	115.5	100.0	100.0	91.5	105.5
50	119.7	100.0	100.0	91.0	104.4
60	124.1	100.0	100.0	91.3	103.0
70	128.7	100.0	100.0	92.6	101.1
80	133.4	100.0	100.0	95.0	98.8
90	138.3	100.0	100.0	98.4	96.1
100	143.3	100.0	100.0	103.2	92.9

The best outcomes for urban areas are shown in Table 20. As it is the growth in traffic that reduces other impacts, the best outcome is associated with the lowest growth in traffic.

Table 20: Best case outcome for urban areas. Relative changes from a baseline 0 % market penetration.

Market penetration of automated vehicles (%)	Traffic volume	Road capacity	Travel time	Fuel consumption	Number of accidents
0	100.0	100.0	100.0	100.0	100.0
10	101.2	101.3	98.5	94.8	101.4
20	102.4	102.6	97.0	89.4	95.9
30	103.7	103.9	95.6	84.0	90.6
40	105.0	105.0	94.1	78.4	85.4
50	106.2	106.2	92.7	72.8	80.5
60	107.5	107.3	91.3	67.0	75.7
70	108.8	108.4	90.0	61.1	71.2
80	110.2	109.4	88.6	55.1	66.8
90	111.5	110.4	87.3	49.1	62.7
100	112.9	111.3	86.0	42.9	58.8

It is seen that both travel time, fuel consumption and the number of accidents is reduced. Figure 30 illustrates the range of outcomes for the number of accidents.

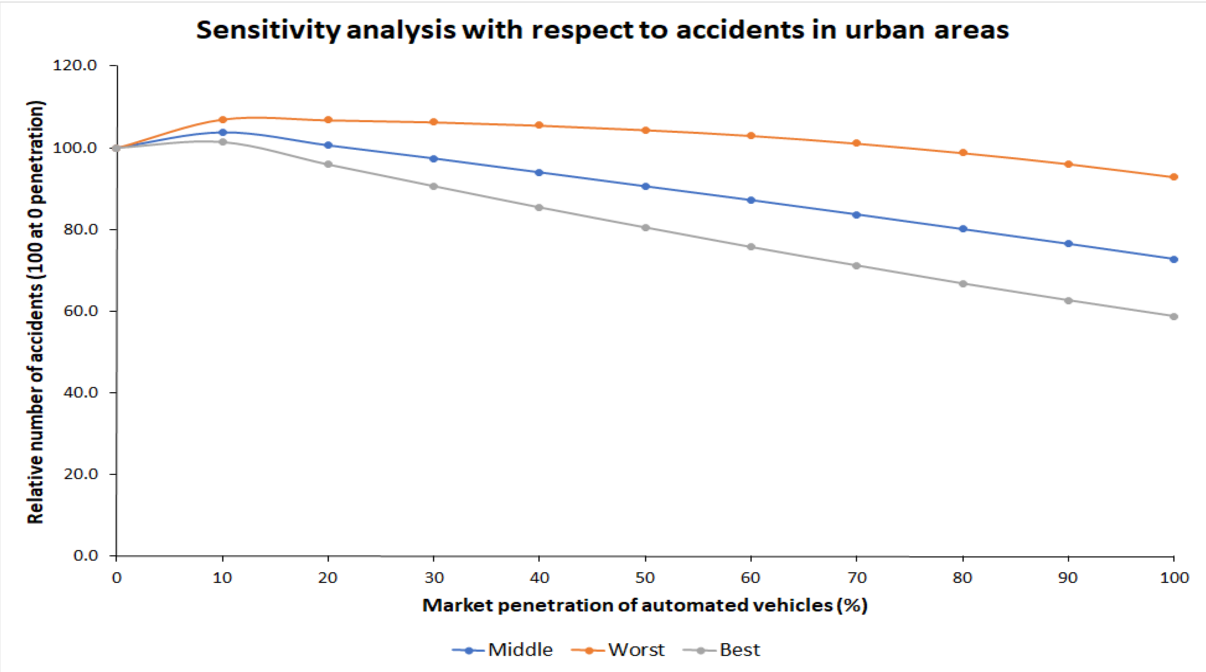

Figure 30: Sensitivity analysis with respect to accidents in urban areas

8 Discussion

This chapter discusses how far the dose-response approach can be taken when trying to estimate impacts of connected and automated vehicles. It also discusses whether some potential impacts can be better modelled as stepwise functions. Finally, it discusses whether all potential impacts that have been identified need to be predicted or whether predictions are inconsequential, as some potential impacts are unlikely to influence the adoption of connected and automated vehicles.

8.1 The dose-response approach – advantages and limitations

As noted in Chapter 3, the transition to a transport system based on connected and automated vehicles is likely to take place gradually over an extended period of time. Impacts are therefore likely to develop gradually as the market penetration of connected and automated vehicles increases. This suggests that it is fruitful to describe impacts in terms of continuous functions, having the level of automation or the degree of market penetration of vehicles at a given level of automation as the principal independent variable.

As far as the level of automation is concerned, it is impossible to predict when vehicles at SAE levels 3, 4 or 5 will be introduced, and how the mix of vehicles at different levels of automation will be in the fleet of vehicles. Any number of mixes can be imagined, such as 40% non-automated, 20% level 3, 20% level 4 and 20% level 5, or any other mixture of percentages one can make up. To make the task of predicting impacts analytically tractable, it has been assumed that all potential impacts are unidirectional and can be thought of as continuous functions. One can imagine more complex relationships, in which the experiences gained when a certain technology has reached a certain level of market penetration lead to changes in technology that changes the impacts it will subsequently have. This form of "learning-by-doing" is likely to take place, but it is not possible to predict its impacts with any confidence. One can imagine that impacts will accumulate slowly at first and then accelerate as a result of learning-by-doing. However, many functional forms are consistent with such a relationship.

It is most consistent with current knowledge, and represents an effective way of summarising it, to model the impacts of connected and automated vehicles as continuous dose-response curves with the level of market penetration representing the dose and a given impact representing the response. Indeed, most of the traffic simulation studies found in the literature either contain dose-response curves or a set of data points to which a dose-response curve can be fitted.

The principal advantages of modelling impacts of connected and automated vehicles by means of dose-response functions are:

1. The functions are consistent with a gradual introduction of connected and automated vehicles.
2. The functions bypass the need to make predictions about when, and how long, the introduction of connected and automated vehicles will take. These aspects are implicit in the market penetration rate of connected and automated vehicles.
3. The functions are flexible with respect to shape: they can be linear or non-linear, steep or flat.
4. It is possible to account for interactions between impacts, or rebound impacts, by multiplying dose-response curves with each other.
5. It is possible to produce mixed dose-response curves that are weighted or non-weighted combinations of several curves.
6. If a dose-response curve is based on many data points, the dispersion of the data points can be interpreted as an indicator of the uncertainty surrounding the curve.
7. It is possible to develop dose response curves for at least the principal impacts of connected and automated vehicles as observed in traffic, such as road capacity, traffic volume, travel time, travel delay, fuel consumption and the number of accidents.

There are, on the other hand, several limitations in modelling the impacts of connected and automated vehicles as dose-response functions:

1. It is, at the present state of knowledge, only possible to include easily observable impacts in traffic. There are no dose-response functions for some potentially important impacts, like a transition to shared mobility or to electric vehicles.
2. Some potential impacts of connected and automated vehicles may not develop gradually in the manner of a continuous function, but rather take the form of one or a few major steps.
3. Little is known about the uncertainty associated with some dose-response curves. Curves based on few data points may give a spurious impression of precision, as functions tend to fit better when there are few data points than when there are many.
4. Some dose-response curves, in particular those for road safety, are based on surrogate measures, mainly traffic conflicts. In this deliverable, no attempt has been to establish conversion rates between the number of conflicts and the number of accidents.
5. The dose-response curves presented in this deliverable mostly refer to potential impacts of vehicles at the highest level of automation. The curves therefore represent maximum potential impacts.

These limitations mean that a policy support tool designed to help policymakers estimate impacts of connected and automated vehicles cannot be based exclusively on dose-response functions.

8.2 Non-continuous impacts

Some potential impacts of connected and automated vehicles may not take the form of continuous functions.

There will be a stepwise change in vehicle ownership costs. Connected and automated vehicles are expected to be more expensive to buy than current vehicles, but cheaper to operate. The additional ownership costs are expected to come down over time, as economies of scale are realised in the production of the vehicles. These changes are very difficult to predict and may not necessarily resemble any continuous function.

Trust in technology is also likely to be discontinuous. Building trust takes times and a single mishap may generate distrust that may take long to overcome. It is likely that people will be less forgiving with respect to malfunctions in automation technology than with respect to errors made by a human driver.

Access to travel for people who are currently not able to drive a car is also likely to become possible only at the highest level of automation. Public transport is, at least to some extent, accessible today for children, the blind or people who for other reasons are not allowed to drive a motor vehicle. But public transport today has human drivers, and there will usually be other passengers who may assist a disabled person should the need for it arise. It is only when automation technology has become very reliable that it can become fully accessible to those who are currently excluded from driving motor vehicles. Thus, gain in accessibility is likely to be a stepwise change, not a gradual one.

8.3 Other potential impacts

Some of the potential impacts of connected and automated vehicles may not need to be predicted, because these potential impacts are likely to be regarded either as irrelevant (i.e. their occurrence will not influence the introduction of CATS) or too small to be worth the effort of trying to predict them. Other impacts are likely to be self-corrected through the market mechanism.

A case belonging to the latter category is travel comfort. People are unlikely to accept automated vehicles if travel comfort in them is substantially below current travel comfort in a car or a public transport vehicle. An automated car that accelerates in jerks, brakes very hard, and has no well-functioning ventilation system will be regarded as a step backwards and will not sell in the market. The public will demand at least the same level of comfort as in current vehicles. Motion sickness could be an issue in automated vehicles, which it mostly is not in cars today. Car passengers experience changes in speed and direction as more violent than car drivers do. To compensate for this, and reduce the risk of motion sickness, automated vehicles would have to adopt a very soft and smooth style of driving, avoiding any sudden changes in speed and/or direction. Car manufacturers will simply have to adapt to these expectations.

Three other potential impacts whose relevance is at least open to discussion are changes in geographic accessibility, employment and public finances. The analysis by Meyer et al. (2017) show minor changes in geographic accessibility as a result of automated vehicles. In a conservative scenario, mean accessibility, weighted by population size increased by 1.4%. In an optimistic scenario, the increase was 10%. These small effects are not surprising. Automated vehicles will, at least in the short term, not change the road system or the location of businesses and residential areas. They may stimulate urban sprawl, which, if anything, reduces accessibility. The small effects found by Meyer et al. (2017), and the apparently large uncertainty associated with them, suggests that putting more effort into modelling these effects should have low priority.

Similar points of view are relevant with respect to effects on employment. The International Transport Forum (2017) discusses potential effects on employment. It is clear that full automation of transport means that drivers will lose their jobs. However, it is not likely that this loss of jobs will be seen as a decisive argument against introducing automated vehicles. There are many historical examples of how the introduction of new technology has caused some types of jobs to disappear and new types of jobs to arise. This is the very nature of economic development. It is therefore suggested not to include changes in employment as a potential impact of connected and automated vehicles.

Changes in public finances as a result of the introduction of connected and automated vehicles has also been mentioned in some studies. An example is loss of income from parking fees if shared mobility replaces individual mobility. Again, however, public finances are very flexible. A city ought to be able to replace the loss of income from parking fees by creating another source of revenue. Road pricing is one possibility. Road pricing may be used to prevent transport from becoming cheaper. Although cheaper transport, in particular cheaper commercial transport where the costs of having a driver are saved, has been put forward as one the advantages of connected and automated transport, it is highly debatable if this is indeed an advantage. Cheaper transport is likely to mean more transport, and more transport will generate rebound effects that diminish many of the potential benefits of connected and automated vehicles. An increased traffic volume will reduce savings in travel time, may generate more emissions, may diminish gains in road safety and may ultimately require an expansion of road capacity.

9 Conclusions

This chapter summarises the main conclusions of the study presented in this deliverable with respect to: (1) The methods that can be applied to estimate and predict potential impacts of connected and automated vehicles; (2) The choice of how best to represent the impacts of connected and automated vehicles; (3) The need for more knowledge about the potential impacts of connected and automated vehicles; and (4) The choice of impacts to include in an assessment of the potential impacts of connected and automated vehicles.

The main conclusions of the study presented in this deliverable will be presented under four major headings.

9.1 Methods for estimating and predicting impacts

A survey was made of methods that can be applied for the purpose of estimating and predicting potential impacts of connected and automated vehicles. A distinction was made between retrospective methods and future-oriented methods. Both classes of methods are relevant for predicting impacts of connected and automated vehicles. The main conclusions are:

1. Various forms of traffic simulations are the most commonly applied method for predicting the impacts of connected and automated vehicles.
2. Traffic simulations have been used to study several potential impacts of connected and automated vehicles, including impacts on:
    a. Road capacity, including junction capacity
    b. Traffic volume
    c. Travel time, including congestion delays
    d. Fuel consumption
    e. Road accidents
3. The results of most simulation studies show potential impacts of connected and automated vehicles as a function of their market penetration rate.
4. Most simulation studies refer to an urban traffic environment or to motorways. There are few simulation studies dealing with impacts on rural all-purpose roads.
5. The simulation studies all rely on a ceteris paribus condition, meaning they assume that all else remains equal. This means that they do not account for potential rebound effects of connected and automated vehicles.
6. One rebound effect which is very likely to occur as a result of increased road capacity is induced traffic. The effects of increased road capacity on travel demand have been extensively studied and summarised by means of meta-analyses.
7. Cost analyses and studies of willingness-to-pay are relevant to provide a basis for cost-benefit analyses of connected and automated vehicles. A few such studies have been reported.

9.2 How to describe and measure impacts

There are many ways of describing and measuring the potential impacts of connected and automated vehicles. The following approaches to quantifying impacts are the most common:

1. Impacts are modelled as a dose-response curve having the rate of market penetration of automated vehicles as the dose and the size of an impact as the response.
2. The dose-response model is flexible and applicable to many potential impacts of connected and automated vehicles.
3. Not all impacts of connected and automated vehicles will take the form of a continuous function. Some impacts are likely to resemble stepwise changes.

In this deliverable, dose-response curves have been estimated for the following impacts:

1. Impacts on capacity and mobility
    a. Lane capacity
    b. Junction capacity
    c. Delays on motorways
    d. Delays in roundabouts
    e. Delays in signalised junctions
    f. Travel time on motorways
    g. Delays in junctions (average of roundabouts and signalised)
    h. Travel time in cities (derived from changes in capacity and delays)
2. Impacts on safety
    a. Rear-end and lane change collisions (motorways)
    b. Accidents in signalised junctions
    c. Accidents in roundabouts
    d. Accidents in priority controlled junctions
    e. Cyclist and pedestrian accidents
    f. Accidents in urban junctions (average of signals, roundabouts, priority)
3. Impacts on fuel consumption (and emissions)
    a. Fuel consumption (and emissions, if these are assumed to be proportional to fuel consumption)
4. Other impacts and feedback impacts
    a. Diffusion curves for market penetration of automated vehicles
    b. (Induced) travel demand on motorways
    c. (Induced) travel demand in cities
    d. Behavioural adaptation to platoons by manually driven vehicles

In addition to the dose-response functions, the following impacts were quantified:

4. Willingness-to-pay for automation technology
5. Demand function for automated vehicles
6. Valuation of travel time in automated vehicles

Connected and automated vehicles are likely to increase road capacity, reduce travel time, reduce emissions and reduce the number of accidents. A rebound effect in terms of increased traffic is likely to occur.

9.3 Need for more knowledge

Some potential impacts of connected and automated vehicles are, at the current state of knowledge, more difficult to predict than others. Impacts that are highly uncertain include:

1. Whether there will be a widespread transition from individual to shared mobility. There is no consensus on whether individual use of motor vehicles will continue at present levels or be replaced by various forms of shared mobility.
2. It is not clear what type of propulsion energy connected and automated vehicles will use. Some researchers expect the introduction of connected and automated vehicles to be associated with a transition to electric propulsion.
3. Connected and automated vehicles are vulnerable to cyber attacks. However, the risk of such attacks cannot be quantified. Potential scenarios can be described.
4. The costs of connected and automated vehicles are highly uncertain. It is not clear that connected and automated vehicles will be as affordable as current motor vehicles. The costs of automation technology may influence the level of inequality in access to transport.
5. Behavioural adaptation to connected and automated vehicles, in particular during the transition period before full market penetration. While some studies suggest various forms of behavioural adaptation, predicting its form and impacts is impossible or speculative.
6. Changes in employment are difficult to predict. While full automation will eliminate the need for drivers, other potential changes in employment are less known.

9.4 Impact inclusion criteria

Some potential impacts of connected and automated vehicles are very poorly known or impossible to predict. Some impacts are likely to be small. Some impacts can be neutralised by means of compensatory policy instruments. And finally, some impacts are likely to be regarded as irrelevant for policy making, although they are predictable, and per se possibly even regarded as undesirable. These impacts include:

1. Changes in travel comfort, as related to physical components such as acceleration, braking and ventilation. Vehicles performing poorly with respect to comfort will not sell and manufacturers will be forced to adapt to feedback from the market.
2. Changes in employment. Fully automated vehicles do not need a driver. However, the loss of driver jobs is unlikely to be viewed as an important, or even relevant, objection to the introduction of automated vehicles.
3. Changes in geographic accessibility. Accessibility is mainly a function of topography and the likely impacts on it of vehicle automation are small.
4. Changes in public finances. Loss of any source of income can almost always be compensated for by the creation of another source of income.

References

Ahanchian, M., Gregg, J. S., Tattini, J., Karlsson, K. B. 2019. Analyzing effects of transport policies on travelers' rational behaviour for modal shift in Denmark. Case Studies on Transport Policy, 7, 849-861.

Anderson, J. M., Kalra, N., Stanley, K. D., Sorensen, P., Samaras, C., Oluwatola, O. A. 2016. Autnomous vehicle technology. A guide for policymakers. Santa Monica, Rand Corporation.

Antin, J. F., Guo, F., Fang, Y., Dingus, T. A., Perez, M. A., Hankey, J. M. 2017. A validation of the low mileage bias using naturalistic driving study data. Journal of Safety Research, 63, 115-120.

Atkins, 2016. Research on the impacts of connected and autonomous vehicles on traffic flow. Stage 2: Traffic modelling and analysis. Technical Report.

Bansal, P., Kockelman, K. M. 2017. Forecasting Americans' long-term adoption of connected and autonomous vehicle technologies. Transportation Research Part A, 49-63.

Bansal, P., Kockelman K. M. 2018. Are we ready to embrace connected and self-driving vehicles? A case study of Texans. Transportation, 45, 641-675.

Bansal, P., Kockelman, K. M., Singh, A. 2016. Assessing public opinions of and interest in new vehicle technologies: An Austin perspective. Transportation Research Part C, 67, 1-14.

Bhoopalam, A.K., Agatz, N., Zuidwijk, R., 2018. Planning of truck platoons: a literature review and directions for future research. Transportation Research Part B, 107, 212-228.

Bird, J. 2019. Car hacking threatens vision of connected mobility. *Financial Times*. Retrieved from https://www.ft.com/content/163f08c6-6ce3-11e9-9ff9-8c855179f1c4

Bjørnskau, T. 2017. The Zebra Crossing Game – Using game theory to explain a discrepancy between road user behaviour and traffic rules. Safety Science, 92, 298-301.

Boggs, A. M., Arvin, R., Khattak, A. J. 2020. Exploring the who, what, when, where, and why of automated vehicle disengagements. Accident Analysis and Prevention, 136, 105406.

Booz Allen Hamilton et al. 2018. Dedicating lanes for priority or exclusive use by connected and automated vehicles. NCHRP research report 891. Washington DC, Transportation Research Board.

Boyd, C. 2018. IoT domestic abuse: What can we do to stop it? Retrieved from https://blog.malwarebytes.com/cybercrime/2018/07/iot-domestic-abuse-can-stop/

Buttice, C. 2019. Hacking Autonomous Vehicles: Is This Why We Don't Have Self-Driving Cars Yet? . *Techopedia*. Retrieved from https://www.techopedia.com/hacking-autonomous-vehicles-is-this-why-we-dont-have-self-driving-cars-yet/2/33650

Bösch, P. M., Becker, F., Becker, H., Axhausen, K. W. 2018. Cost-based analysis of autonomous mobility services. Transport Policy, 64, 76-91.

Calgaro, J.A. Traffic Loads Traffic Loads on Bridges - EN 1991-2. Information workshop Eurocode Background and applications, Brussels, 18-20 February 2008.

Campbell, P. 2018. Hackers have self-driving cars in their headlights. *Financial Tiems*. Retrieved from https://www.ft.com/content/6000981a-1e03-11e8-aaca-4574d7dabfb6

Cantieni R. Dynamic behavior of highway bridges under the passage of heavy vehicles (EMPA Report No. 220), 1992. p. 220.

Cervero, R. 2003. Road expansion, urban growth, and induced travel: a path analysis. Journal of the American Planning Association, 69, 145-163.

Chaturvedi, D. K. 2009. Modeling and simulation of systems using MATLAB and simulink. London, Taylor and Francis.

Chen, Y., Gonder, J., Young, S., Wood, E. 2019. Quantifying autonomous vehicles national fuel consumption impacts: A data-rich approach. Transportation Research Part A, 122, 134-145.

Cohn, J., Ezike, R., Martin, J., Donkor, K., Ridgway, M., Balding, M. 2019. Examining the equity impacts of autonomous vehicles: a travel demand model approach. Transportation Research Record, 2673, 23-35.

Combs, T. S., Sandt, L. S., Clamann, M. P., McDonald, N. C. 2018. Automated vehicles and pedestrian safety: exploring the promise and limits of pedestrian detection. American Journal of Preventive Medicine. doi.org/10.1016/j.amepre.2018.06.024.

Consumer Watchdog. 2019. Kill Switch. Why connected cars can be killing machines and how to turn them off. Retrieved from Santa Monica, CA: https://www.consumerwatchdog.org/report/kill-switch-why-connected-cars-can-be-killing-machines-and-how-turn-them

Cunningham, M. L., Regan, M. A., Ledger, S. A., Bennett, J. M. 2019. To buy or not to buy? Predicting willingness to pay for automated vehicles based on public opinion. Transportation Research Part F, 65, 418-438.

Dalton, R. J., Recchia, S., Rohrschneider, R. 2003. The Environmental Movement and the Modes of Political Action. *Comparative Political Studies, 36*(7), 743-771.

Daziano, R. A., Sarrias, M., Leard, B. 2017. Are consumers willing to pay to let cars drive for them? Analyzing response to autonomous vehicles. Transportation Research Part C., 78, 150-164.

Dingus, T. A., Guo, F., Lee, S., Antin, J. F., Perez, M., Buchanan-King, M., Hankey, J. 2016. Driver crash risk factors and prevalence evaluation using naturalistic driving data. Proceedings of the National Academy of Sciences (doi/10.1073/pnas.1513271113).

Doffman, Z. 2019. State-Sponsored Cyberattacks 'Challenge The Very Concept Of War'. *Forbes.*

Ebbesen, E. B., Haney, M. 1973. Flirting with death: Variables affecting risk-taking at intersections. Journal of Applied Social Psychology, 3, 303-324.

Eckhardt, J., Aarts, L. , van Vliet, A. , Alkim, T. , 2016. European truck platooning challenge 2016. Lessons learnt.

Elvik, R. 2009. An exploratory analysis of models for estimating the combined effects of road safety measures. Accident Analysis and Prevention, 41, 876-880.

Elvik, R. 2010. An exploratory study of mechanisms by which exposure influences accident occurrence. Transportation Research Record, 2148, 76-82.

Elvik, R. 2012. Speed limits, enforcement, and public health consequences. Annual Review of Public Health, 33, 225-238.

Elvik, R. 2015. Some implications of an event-based definition of exposure to the risk of road accident. Accident Analysis and Prevention, 76, 15-24.

Elvik, R., Amundsen, A. H. 2000. Improving Road Safety in Sweden. Report 490. Oslo, Institute of Transport Economics.

Elvik, R., Amundsen, A. H. 2004. Effects on road safety of new urban arterial roads. Accident Analysis and Prevention, 36, 115-123.

Elvik, R., Erke, A., Christensen, P. 2009. Elementary units of exposure. Transportation Research Record, 2103, 25-31.

Elvik, R., Ramjerdi, F. 2014. A comparative analysis of the effects of economic policy instruments in promoting environmentally sustainable transport. Transport Policy, 33, 89-95.

Eriksson, A., Stanton, N. A. 2018. Driver reactions to automated vehicles. Boca Raton, CRC Press.

Evans, L., Wasielewski, P. 1982. Do accident involved drivers exhibit riskier everyday driving behavior? Accident Analysis and Prevention, 14, 57-64.

Fagnant, D. J., Kockelman, K. 2015. Preparing a nation for autonomous vehicles: opportunities, barriers and policy recommendations. Transportation Research Part A, 77, 167-181.

Favaro, F. M., Nader, N. Eurich, S. O., Tripp M., Varadaraju, N. 2017. Examining accident reports involving autonomous vehicles in California. Plos One (doi.org/10.1371/journal.pone.0184952).

Fernandes, P., Nunes, U. 2012. Platooning with IVC-enabled autonomous vehicles: Strategies to mitigate communication delays, improve safety and traffic flow. IEEE Transactions on Intelligent Transport Systems, 13, 91-106.

Flügel, S., Halse, A. H., Hulleberg, N., Jordbakke, G. N. 2019. Estimating the effect of vehicle automation on cars drivers' and car passengers' valuation of travel time savings. Working paper 51396. Oslo, Institute of Transport Economics.

Fridstrøm, L., Østli. V. 2016. Kjøretøyparkens utvikling og klimagassutslipp. Framskrivinger med modellen BIG. Rapport 1518. Oslo, Transportøkonomisk institutt.

Gondaliya, P., Manisha, C. 2015. Modeling and system identification of liquid level system. International Journal of Science, Engineering and Technology, 4, 1716-1719.

Goodin, D. 2018. A $225 GPS spoofer can send sat-nav-guided vehicles into oncoming traffic. *Ars Technica*. Retrieved from https://arstechnica.com/information-technology/2018/07/a-225-gps-spoofer-can-send-autonomous-vehicles-into-oncoming-traffic/

Gouy, M., Wiedemann, K., Stevens, A., Brunett, G., Reed, N. 2014. Driving next to automated vehicle platoons: How do short time headways influence non-platoon drivers' longitudinal control? Transportation Research Part F, 264-273.

Grush, B., Niles, J. 2018. The end of driving. Transportation systems and public policy planning for autonomous vehicles. Oxford, Elsevier Science.

Hawkins, A. J. 2019. Tesla didn't fix an Autopilot problem for three years, and now another person is dead. *The Verge*. Retrieved from https://www.theverge.com/2019/5/17/18629214/tesla-autopilot-crash-death-josh-brown-jeremy-banner

Heymann, M., Degani. A. 2019. Autonomous vehicle interactions with other road users: conflicts and resolutions. Proceedings of the tenth international symposium on human factors in driver assessment, training and vehicle design.

Herrmann, A., Brenner, W., Stadler, R. 2018. Autonomous driving. How the driverless revolution will change the world. Bingley, Emerald publishing.

Hjorthol, R., Engebretsen, Ø, Uteng, T. P. 2014. Den nasjonale reisevaneundersøkelsen 2013/14 – nøkkelrapport. Rapport 1383. Oslo, Transportøkonomisk institutt.

Hull, R. 2019. A car is stolen EVERY 5 MINUTES in Britain: Vehicle thefts up 50% in 5 years thanks to criminal gangs hacking keyless systems and reduced policing. *This is Money*. Retrieved from https://www.thisismoney.co.uk/money/cars/article-6631925/Vehicle-thefts-risen-50-5-years-Home-Office-figures-show.html

Hymel, K. 2019. If you build it, they will drive: Measuring induced demand for vehicle travel in urban areas. Transport Policy, 76, 57-66.

Høye, A., Hesjevoll, I. S., Vaa, T. 2015. Førerstøttesystemer. Status og potensial for fremtiden. Rapport 1450. Oslo, Transportøkonomisk institutt.

International Transport Forum (ITF). 2017. Managing the transition to driverless road freight transport. Case-Specific Policy Analysis. Paris, OECD Publishing.

Jamson, A. H., Merat, N., Carsten, O., Lai, F. 2013. Behavioural changes in drivers experiencing highly-automated vehicle control in varying traffic conditions. Transportation Research Part C, 30, 116-125.

Jeong, E., Oh, C., Lee, S. 2017. Is vehicle automation enough to prevent crashes? Role of traffic operations in automated driving environments for traffic safety, Accident Analysis & Prevention, 104, 115–124. doi: 10.1016/j.aap.2017.05.002.

Kalra, N, Groves, D. G. 2017. The enemy of good. Estimating the cost of waiting for nearly perfect automated vehicles. Santa Monica, The Rand Corporation.

Kalra, N., Paddock, S. M. 2016. Driving to safety: How many miles of driving would it take to demonstrate autonomous vehicle reliability? Transportation Research Part A, 94, 182-193.

Kockelman, K., Bansal, P., et al. 2016. Implications of connected and automated vehicles on the safety and operations of roadway networks: a final report. CTR technical report 0-6849-1. The University of Texas at Austin.

Kockelman, K., Boyles, S. 2018. Smart transport for cities and nations. The rise of self-driving and connected vehicles. The University of Texas at Austin.

Kolarova, V., Steck, F., Cyganski, R., Trommer, S. 2018. Estimation of the value of time for automated driving using revealed and stated preference methods. Transportation Research Procedia, 31, 35-46.

Kovaceva, J., et al. (14 authors). 2019. Assessment of the PROSPECT safety systens including socio-economic evaluation. Deliverable D2.3 of the PROSPECT project (proactive safety for pedestrians and cyclists). Gothenburg, Chalmers University of Technology.

Kyriakidis, M., Happee, R., de Winter, J. C. F. 2015. Public opinion on automated driving: Results of an international questionnaire among 5000 respondents. Transportation Research Part F, 32, 127-140.

Kucharski, R., Drabicki, A. 2017. Estimating macoscopic volume delay functions with the traffic density derived from measured speeds and flows. Journal of Advanced Transportation, article ID 4629792.

Leahy C., OBrien E., O'Connor A. 2016.The effect of traffic growth on characteristic bridge load effects. Transportation Research Procedia 14, 3990 – 3999.

Leslie, A., Kiefer, R. J., Meitzner, M. R., Flannagan, C. A. 2019. Analysis of the field effectiveness of General Motors production active safety and advanced headlighting systems. UMTRI-2019-6. Ann Arbor, University of Michigan Transportation Research Institute.

Le Vine, S., Zolfaghari, A., Polak, J. 2015. Autonomous cars: The tension between occupant experience and intersection capacity. Transportation Research Part C, 52, 1-14.

Li, Z. et al. 2013. Modeling Reservation-Based Autonomous Intersection Control in VISSIM Modeling Reservation-Based Autonomous Intersection Control in VISSIM, Transportation Research Record: Journal of the Transportation Research Board, 2381, 81–90. doi: 10.3141/2381-10.

Li, Y., Wang, H., Xing, L., Liu, S., Wei, X. 2017. Evaluation of the impacts of cooperative adaptive cruise control on reducing rear-end collision risks on freeways. Accident Analysis and Prevention, 98, 87-95.

Li, Y., Tu, Y., Fan, Q., Dong, C., Wang, W. 2018. Influence of cyber-attacks on longitudinal safety of connected and automated vehicles. Accident Analysis and Prevention, 121, 148-156.

Linvill, D., Warren, P. 2019. That Uplifting Tweet You Just Shared? A Russian Troll Sent It. *Rolling Stone*. Retrieved from https://www.rollingstone.com/politics/politics-features/russia-troll-2020-election-interference-twitter-916482/

Liu, P., Guo, Q., Ren, F., Wang, L., Xu, Z. 2019. Willingness to pay for self-driving vehicles: influences of demographic and psychological factors. Transportation Research Part C, 100, 306-317.

Liu, P. , Yang, R., Xu, Z. 2019. Public acceptance of fully automated driving: Effects of social trust and risk/benefit perceptions. Risk Analysis, 39, 326-341.

Liu, X., Shen, D., Lai, L., Le Vine, S. 2020. Optimizing the safety-efficiency balancing of automated vehicle car-following. Accident Analysis and Prevention, 136, 105435.

Malwarebytes. 2019. What is hacking? Retrieved from https://www.malwarebytes.com/hacker/

Meadows, D.H., Meadows, D. L., Randers, J., Behrens W. W. 1972. The limits to growth. New York, Signet books.

Meyer, J., Becker H., Bösch, P. M., Axhausen, K. W. 2017. Autonomous vehicles: the next jump in accessibility? Research in Transportation Economics, 62, 80-91.

Milakis, D., van Arem, B, van Wee, B. 2017. Policy and society related implications of automated driving: A review of literature and directions for future research. Journal of Intelligent Transportation Systems, 21, 324-348.

Milakis, D., Snelder, M., van Arem, B., van Wee, B., Correia, G. C. 2017. Development and transport implications of automated vehicles in the Netherlands: scenarios for 2030 and 2050. European Journal of Transport Infrastructure Research, 17, 63-85.

Moore, R. 2014. Cybercrime: Investigating high-technology computer crime. London, Routledge.

Morando, M. M., Tian, Q., Truong, L. T., Vu, H. V. 2018. Studying the safety impact of autonomous vehicles using simulation-based surrigate safety measures. Journal of Advanced Transportation. doi.org/10.1155/2018/6135183.

Nieuwenhuijsen, J., Correia, G. H., Milakis, D., van Arem, B., van Daalen, E. 2018. Towards a quantitative method to analyze the long-term innovation diffusion of automated vehicles technology using system dynamics. Transportation Research Part C, 86, 300-327.

Noland, R. B., Lem, L. L., 2002. A review of the evidence for induced travel and changes in transportation and environmental policy in the US and the UK. Transportation Research, Series D, 7, 1-26.

Noy, I. Y., Shinar, D., Horrey, W. J. 2018. Automated driving: Safety blind spots. Safety Science, 102, 68-78.

Office for National Statistics. 2018. *Domestic abuse: findings from the Crime Survey for England and Wales: year ending March 2018*. Retrieved from London: https://www.ons.gov.uk/peoplepopulationandcommunity/crimeandjustice/articles/domesticabusefindingsfromthecrimesurveyforenglandandwales/yearendingmarch2018

Office of the Chairman of the Joint Chiefs of Staff. 2019. *DOD Dictionary of Military and Accociated Terms (July 2019)*. Washington DC, United States Joint Chiefs of Staff.

Olia, A., Abdelgawad, H., Abdulhai, B., Razavi, S. N. 2016. Assessing potential impacts of connected vehicles: Mobility, environmental, and safety perspectives. Journal of Intelligent Transportation Systems, 20, 229-243.

Omozik, K., Yang, Y., Kuntermann, I., Hergeth, S., Bengler, K. 2019. How long does it take to relax? Observation of driver behaviour during real-world conditionally automated driving. Proceedings of the tenth international symposium on human factors in driver assessment, training and vehicle design.

Osborne, C. 2018. The most interesting Internet-connected vehicle hacks on record. *ZDNet*. Retrieved from https://www.zdnet.com/article/these-are-the-most-interesting-ways-to-hack-internet-connected-vehicles/

Papadoulis, A., Quddus, M., Imprialou, M. 2019. Evaluating the safety impact of connected and autonomous vehicles on motorways. Accident Analysis and Prevention, 124, 12-22.

Pfaffenbichler, P., Emberger, G., Shepherd, S. 2010. A system dynamics approach to land use transport interaction modelling: the strategic model MARS and its application. System Dynamics Review, 3, 262-282.

Prat M. 2001. Traffic load models for bridge design: recent developments and research. Prog Struct Eng Mater, 3, 326–334.

Pogue, C. 2018. The 2018 Black Report. Decoding the minds of hackers. Retrieved from https://www.nuix.com/black-report/black-report-2018

Rahman, M. S., Abdel-Aty, M. 2018. Longitudinal safety evaluation of connected vehicles' platooning on expressways, Accident Analysis and Prevention, 117, 381–391.

Rahman , M. S., Abdel-Aty, M., Lee, J., Rahman, M. H. 2019. Safety benefits of arterials' crash risk under connected and automated vehicles. Transportation Research Part C, 100, 354-371.

Rahman, M. S., Abdel-Aty, M., Wang, L., Lee, J. 2018. Understanding the highway safety benefits of different approaches of connected vehicles in reduced visibility conditions. Transportation Research Record, 2672, 91-101.

Rahmati, Y., Hosseini, M. K., Talebpour, A., Swain, B., Nelson, C. 2019. Influence of autonomous vehicles on car-following behavior of human drivers. Transportation Research Record, 2673, 367-379.

Rosén, E., Sander U. 2009. Pedestrian fatality risk as a function of impact speed. Accident Analysis and Prevention, 41, 536-542.

Schmitt, B. 2019. Top OTA Expert Shows How State Actors Hack into Your Car and What Happens Next: 'People Will Die'. *The Drive*. Retrieved from https://www.thedrive.com/tech/29120/top-ota-expert-shows-how-state-actors-hack-into-your-car-and-what-happens-next-people-will-die

Shabanpour, R., Golshani, N., Shamshiripour, A., Mohammadian, A. 2018. Eliciting preferences for adoption of fully automated vehicles using best-worst analysis. Transportation Research Part C, 93, 463-478.

Shepherd, S. P. 2014. A review of system dynamics models applied in transportation. Transportmetrica B: Transport Dynamics, 2, 83-105.

Shi, Y., He, Q., Huang, Z. 2019. Capacity analysis and cooperative lane changing for connected and automated vehicles: entropy-based assessment method. Transportation Research Record, 2673, 485-498.

Shin, K. J., Tada, N., Managi, S. 2019. Consumer demand for fully automated driving technology. Economic Analysis and Policy, 61, 16-28.

Shladover, S. E., Su, D., Lu, X.-Y. 2012. Impacts of Cooperative Adaptive Cruise Control on Freeway Traffic Flow, Transportation Research Record: Journal of the Transportation Research Board, 2324, 63–70. doi: 10.3141/2324-08.

Slowik, P., Sharpe, B. 2018. Automation in the long haul: Challenges and opportunities of autonomous heavy-duty trucking in the United States. Working paper 2018-06. International Council on Clean Transportation.

Soare, B. 2019. Are Hackers Threatening the Adoption of Self-Driving Cars? And all about the current state of autonomous vehicles. Retrieved from https://heimdalsecurity.com/blog/hackers-self-driving-cars/

Soteropoulos, A., Berger, M., Ciari, F. 2019. Impacts of automated vehicles on travel behaviour and land use: an international review of modelling studies. Transport Reviews, 39, 29-49.

Sparrow, R., Howard, M. 2017. When human beings are like drunk robots: Driverless vehicles, ethics, and the future of transport. Transportation Research Part C, 80, 206-215.

Standards Norway. 2014. Societal security – protection against undesirable actions – requirements for security risk analysis. Oslo, Standards Norway NS 5832:2014

Stanek, D. et al. 2018. Measuring Autonomous Vehicle Impacts on Congested Networks Using Simulation, in Transportation Research Procedia.

Steck, F., Kolarova, V., Bahamonde-Birke, F., Trommer, S., Lenz, B. 2018. How autonomous driving may affect the value of travel time savings for commuting. Paper submitted to 97th TRB Annual Meeting.

Tainio, M. 2015. Burden of disease caused by local transport in Warsaw, Poland. Journal of Transport and Health, 2, 423-433.

Tencent Keen Security Lab. 2019. *Experimental Security Research of Tesla Autopilot*. Retrieved from https://keenlab.tencent.com/en/whitepapers/Experimental_Security_Research_of_Tesla_Autopilot.pdf

Tennøy, A., Tønnesen, A., Gundersen, F. 2019. Effects of urban road capacity expansion – experiences from two Norwegian cases. Transportation Research Part D, 69, 90-106.

Thomson, I. 2018. Say what you will about self-driving cars - the security is looking 'OK'. *The Register*. Retrieved from https://www.theregister.co.uk/2018/08/10/autonomous_car_hacking/

Tientrakool, P., Ho, Y-C., Maxemchuk, F. 2011. Highway capacity benefits froim using vehicle-to-vehicle communication and sensors for collision avoidance. IEEE proceedings of Vehicular Technology Conference, San Francisco.

Tirachini, A., Antoniou, C. 2020. The economics of automated public transport: Effects on operator cost, travel time, fare and subsidy. Economics of Transportation, 21, 100151.

Vasebi, S., Hayeri, Y. M., Samaras, C., Hendrickson, C. 2018. Low-level automated light-duty vehicle technologies provide opportunities to reduce fuel consumption. Transportation Research Record, 2672, 60-74.

Virdi, N., Grzybowska, H., Waller, S. T., Dixit, V. 2019. A safety assessment of mixed fleets with connected and autonomous vehicles using the surrogate safety assessment module. Accident Analysis and Prevention, 131, 95-111.

Wadud, Z. 2017. Fully automated vehicles: A cost of ownership analysis to inform early adoption. Transportation Research Part A, 101, 163-176.

Wadud, Z., MacKenzie, D., Leiby, P. 2016. Help or hindrance? The travel, energy and carbon impacts of highly automated vehicles, Transportation Research Part A: Policy and Practice, 86, 1–18. doi: 10.1016/j.tra.2015.12.001.

Wang, L., Zhong, H., Abdel-Aty, M., Park, J. 2020. How many crashes can connected and automated vehicle technologies prevent: A meta-analysis. Accident Analysis and Prevention, 136, 105299.

Wiedemann, R. 1974. Simulation des Strassenverkehrsflusses. Schriftenreihe des Instituts für Verkehrswesen des Universtät Karlsruhe, Heft 8.

Willumsen, L., Kohli, S. 2016. Traffic forecasting and autonomous vehicles. Proceedings of AET Conference.

Yang, Z., Wang, X, Pei, X., Feng, S., Wong, S, C. 2018. Longitudinal safety analysis for heterogeneous platoon of automated and human vehicles. IEEE Transactions on Intelligent Transport Systems (downloaded from Research Gate).

Ye, L., Yamamoto, T. 2018. Modeling connected and autonomous vehicles in heterogeneous traffic flow. Physica A, 490, 269-277.

Zamora, W. 2018. Under the hoodie: why money, power, and ego drive hackers to cybercrime Retrieved from https://blog.malwarebytes.com/cybercrime/2018/08/under-the-hoodie-why-money-power-and-ego-drive-hackers-to-cybercrime/

Zhang, B., Winter. J. de, Varotto, S., Happee, R., Martens, M. 2019. Determinants of take-over time from automated driving: A meta-analysis of 129 studies. Transportation Research Part F, 64, 285-307.