

Safety Assurance for Automated Driving Systems: Consultation Regulation Impact Statement

Please find following our submission to the National Transport Commission's consultation regulation impact statement: "Safety Assurance for Automated Driving Systems." While we acknowledge the decision of the National Transport Commission to exclude ethical considerations from this consultation, whom suggest the requirement for automated driving systems entities (ADSEs) to explicitly include ethical criteria in their statement of compliance will not provide additional safety benefits, we argue that the question of how we facilitate the inclusion of autonomous vehicles into the transport system is necessarily an ethical one. While we believe that ethics is inextricably linked to the question of safety, we also argue that the larger ethical challenge posed by driverless vehicles for transport policy is to ensure that the most socially and environmentally beneficial outcomes are realised.

In response to this RIS, we believe questions 6: 'Are the proposed safety criteria and obligations on ADSEs sufficient, appropriate and proportionate to manage the safety risk?'; and 7: 'Are there any additional criteria or other obligations that should be included?' cannot be adequately answered if ethical considerations are quarantined from discussion. Briefly here, for we discuss the issue more broadly in Appendix A, ethics is inherent to the concept of safety, which must stipulate acceptable levels of harm. In crash scenarios, which will be unavoidable, ADS will encounter a duty to act. This duty is discharged in relation to persons and their wellbeing, and as such, is an ethical one. Further, overarching directives mentioned by the NTC, such as human life is to be valued above all else, and that programming ADS to limit the harm done may be acceptable, are fundamentally ethical in nature. Also, the prohibition of programming crash scenarios that discriminate against individuals based on characteristics such as age and gender, while welcome, is itself an ethical incursion in real world scenarios. Such ethical decisions should be transparent and explicitly acknowledged within the vehicles duty to act. During the duty to act, the question of how best to protect vulnerable road users will also arise, and the appropriate response may in fact require discrimination based on levels of vulnerability to deliver the greatest benefit. This form of discrimination may also involve a heightened risk for vehicle passengers, who will, at times be the least vulnerable party in a crash scenario.

Finally, we hope the attached paper, published in *Transportation Research Part C: Emerging Technologies*, speaks to some of these issues. In addition, we believe our article helps situate questions regarding the regulation of automated vehicle safety within the broader context of the ethical debate concerning the future impact of driverless vehicles on urban planning, and transport policy.

Thank you for considering our submission. Please feel free to contact Professor Robert Sparrow if you would like to further discuss our response.

Yours sincerely

Professor Robert Sparrow and Dr Mark Howard

Tel.: +61 3990 53244

E-mail address: robert.sparrow@monash.edu

Department of Philosophy

School of Philosophical, Historical and International Studies

Monash University

Appendix A

When human beings are like drunk robots: driverless vehicles, ethics, and the future of transport

Robert Sparrow ^{a*} and Mark Howard ^a

*Corresponding author: Tel.: +61 3990 53244; E-mail address robert.sparrow@monash.edu

^a Department of Philosophy

Monash University

Victoria 3800, Australia

Abstract

It is often argued that driverless vehicles will save lives. In this paper, we treat the ethical case for driverless vehicles seriously and show that it has radical implications for the future of transport. After briefly discussing the current state of driverless vehicle technology, we suggest that systems that rely upon human supervision are likely to be dangerous when used by ordinary people in real-world driving conditions and are unlikely to satisfy the desires of consumers. We then argue that the invention of fully autonomous vehicles that pose a lower risk to third parties than human drivers will establish a compelling case against the moral permissibility of manual driving. As long as driverless vehicles aren't safer than human drivers, it will be unethical to sell them. Once they are safer than human drivers when it comes to risks to 3rd parties, then it should be illegal to drive them: at that point human drivers will be the moral equivalent of drunk robots. We also describe two plausible mechanisms whereby this ethical argument may generate political pressure to have it reflected in legislation. Freeing people from the necessity of driving, though, will transform the relationship people have with their cars, which will in turn open up new possibilities for the transport uses of the automobile. The ethical challenge posed by driverless vehicles for transport policy is therefore to ensure that the most socially and environmentally beneficial of these possibilities is realised. We highlight several key policy choices that will determine how likely it is that this challenge will be met.

Keywords

Autonomous cars

Driverless vehicles

Ethics

Urban planning

Safety

Transport policy

Introduction

Every presentation by an engineer on the topic of driverless vehicles that we have seen has begun with a statistic about road crash fatalities. Worldwide, 1.25 million people are killed annually by people driving cars (World Health Organization, 2015, p. 2). An oft-cited statistic is that over 90 percent of all road traffic accidents are a result of human error and behaviour (Anderson et al., 2016; Gao et al., 2014; National Highway Traffic Safety Administration, 2008). Driverless cars, we are solemnly informed, can do better. They will save lives (Anderson et al., 2016; Garza, 2011, p. 606). The case for driverless vehicles is ultimately, then, an ethical one. In this paper, we want to take this claim seriously and show that it has radical implications for the future of transport. Our argument proceeds via a simple dialectic. As long as driverless vehicles aren't safer than human drivers, it will be unethical to sell them (Shladover, 2016). Once they are safer than human drivers when it comes to risks to 3rd parties, then it should be illegal to drive them: at that point human drivers will be the moral equivalent of drunk robots. Freeing people from the necessity of driving, though, will transform the relationship people have with their cars, which will in turn open up new possibilities for the transport uses of the automobile. The challenge posed by driverless vehicles for transport policy is therefore to ensure that the most socially and environmentally beneficial of these possibilities is realised.

In the first section of the paper, we briefly discuss the current state of driverless vehicle technology and argue that systems that rely upon human supervision are likely to be dangerous when used by ordinary people in real-world driving conditions and are unlikely to satisfy the desires of consumers.¹ For this reason, we anticipate that the future of driverless vehicles will be vehicles that are fully autonomous and do not require human supervision.² In section 2, we argue that the invention of fully autonomous vehicles that pose a lower risk to third parties than human drivers will establish a compelling case against the moral permissibility of manual driving. We also highlight two plausible mechanisms whereby this ethical argument may generate political pressure to have it reflected in legislation. In section 3, we acknowledge some complexities that section 2 neglects for the sake of ease of exposition. In the fourth section, we offer some predictions about the implications of the adoption of autonomous vehicles for the future of the automobile and of the city. In the fifth and final section, we highlight several key policy choices that will determine whether or not the best outcomes made possible by the development of autonomous vehicles will be realised.

¹ The SAE international standard for defining levels of automation, which we adopt in this article, ranges from no automation (level 0) to full automation (level 5). Levels 1-3 require human supervision and human takeover when the driving task is beyond the capacity of the automated vehicle system, while levels 4-5 operate without human assistance (level 4 in certain environments): see (NHTSA, 2016) and (Shladover, 2016).

² Unless otherwise noted, when discussing fully autonomous vehicles we are referring to vehicles with SAE level 4 automation, where "an automated system can conduct the driving task and monitor the driving environment, and the human need not take back control ..." (NHTSA, 2016, p. 9), capable of driving in North American and/or European cities and on sealed country roads. While also relevant, achieving SAE level 5 automation where "the automated system can perform all driving tasks, under all conditions that a human driver could perform them," (NHTSA, 2016, p. 9) is unnecessary for the success of our argument. By way of explanation, an SAE level 4 automated vehicle may be capable of performing all driving tasks on all designated roads, while being incapable of off-road driving, with the latter preventing its designation as SAE level 5.

1. Safer at Any Speed: The Argument for Driverless Vehicles

Driverless vehicles are widely anticipated to represent the future of transport (Bamonte, 2013; Bilger, 2013; Burns, 2013). The date at which we may expect their arrival, however, is the subject of some dispute. According to some authorities, driverless cars are possible already and will be on the market by 2020 (Fagnant and Kockelman, 2014, p. 1; Burns, 2013, p. 182). Other writers are more cautious and suggest that it will be perhaps another two or three decades before the technology of driverless vehicles is mature enough to be suitable for widespread use (Litman, 2015, p. 1; Shladover, 2016). The explanation for this difference in expectations regarding the timeframe for the introduction of driverless vehicles lies in a difference in opinion regarding the extent of the technological challenges that will need to be overcome in order for them to safely integrate into the transport environment.

One popular theory about how driverless vehicles will take their place on the roads involves gradually increasing levels of automation of tasks currently performed by humans while still retaining human supervision of the driving task. According to this way of thinking, cruise control and anti-skid braking systems, already widely adopted, represent the beginnings of autonomous driving (Garza, 2011, p. 584; Fagnant and Kockelman, 2015a, p. 168; Gordon and Lidberg, 2015, p. 959); lane change assist and automated freeway driving represent the next step, with Tesla's "autopilot" most of the way to realising this latter goal (Singhvi and Russell, 2016; Vlasiv, 2016). By progressively extending the capacities of these technologies, engineers will be able to produce vehicles that can drive in a larger and larger set of environmental and road conditions, as long as a human being is available to step in should conditions exceed the capacities of the vehicle to handle them safely (Gordon and Lidberg, 2015).³

The driverless vehicles that are being tested on the roads today almost all rely upon supervision of a human being in order to drive at speed under real-world conditions in urban environments (Thrun, 2010; Goodall, 2014a, p. 58; Google, 2016). That is to say, they have only Level 3 or Level 4 autonomy. However, there is an obvious problem with this approach, which we believe will be extremely difficult to overcome: at some point in the not too distant future, when driver assist systems become sufficiently reliable, the human "supervisor" will stop paying attention. Human beings quickly cease to pay attention to — or even to perceive — features of their environment that are not relevant to the tasks in which they are engaged (Merat and Jamson, 2009). To a certain extent this dynamic may be resisted by trained professionals, such as airline pilots, or by engineers who are being paid to monitor the activities of driverless vehicles in order to improve them. Nevertheless, an ordinary person "supervising" the activities of an autonomous vehicle that is 99.9% reliable will quickly cease to be engaged in this task. If the vehicle then requires them to take control quickly when circumstances exceed the capacity of the driving software to manage them safely, the human supervisor is unlikely to be able to do so effectively (Shladover, 2016). It may take more than 30 seconds — and will at minimum take 2 seconds — for a human being to regain situational awareness when required to do so (Lin, 2015, p. 72; Kirkpatrick, 2015, p. 20; Radlmayr et al., 2014).

It might be thought that as long as the autonomous driving systems fail gradually and gracefully it will be possible to arrange for a human driver to retake control of the vehicle as required.⁴ Driverless vehicles need not be able to cope with all road conditions, for instance, in order to be useful. For

³ This appears to be the pathway to the introduction of driverless vehicle technology being pursued by Tesla: see (Bhuiyan, 2016).

⁴ For a discussion of 'safe' handover scenarios see (Gold et al., 2013)

example, a car that is capable of driving on freeways but not local roads might simply alert the driver and require them to retake control as it approaches the end of the freeway.⁵ Equally well, however, if conditions may change rapidly in a given geographic area, this will cause problems. A sudden hailstorm that renders the roads slippery, for instance, or unexpected roadworks, may rapidly transform safe into unsafe driving conditions.⁶

One problem with relying on human drivers to take the wheel when driving conditions change, over the longer term at least, is that the improved reliability of driverless vehicle's is likely to be accompanied by a loss of skill amongst human drivers. As more and more of our daily commute is handled by driverless vehicles our opportunities to practice and even learn how to drive will disappear. Moreover, it is hard to see how it would be possible to rule out situations that required human decision-making within a fraction of a second once vehicles began operating at high speeds. A child will run out from behind a car, for instance, or wildlife will enter the road,⁷ or a vehicle will unexpectedly veer out of the opposite lane as a result of sudden mechanical failure or of the driver having a heart attack. Put enough driverless vehicles on the roads and even the rarest of circumstances becomes inevitable. Unless driverless vehicles are capable of handling them more reliably than human beings, we can expect such circumstances to generate numerous accidents (Goodall, 2014a; Goodall, 2014b).⁸

Furthermore, once driverless vehicles become reliable enough, people are likely to rely on them completely. Busy parents will put their children in the car and instruct it to drive them to school. Pet owners will send their pets to the vet without making the trip themselves. People will fall asleep, or read a book, enjoy being drunk, or choose to have sex, while the car drives itself. Moreover, driverless vehicles will provide autonomy of movement to individuals with physical or cognitive constraints that currently exclude them from driving cars. Even if the vehicle can provide several seconds of warning before requiring manual control to be re-engaged, then, there may be no one in a fit state to do so.

One solution would be to provide the vehicle with the capacity to monitor whether the "supervisor" was paying attention to the route, traffic, and road conditions. Vehicles could sense whether the driver had their hands on the steering wheel, track their gaze, and monitor signs of physiological

⁵ Such scenarios, where an autonomous vehicle detects that it is operating outside of its 'operational design domain' (ODD) and alerts the driver to resume control, are identified as an element of the 'fall back' or 'minimal risk condition' appropriate to automated vehicles in the U.S. Department of Transport's Federal automated vehicles policy (NHTSA, 2016, p 30). A potential 'fall back' option is that of 'teleoperation' where control of the vehicle is handed over to a remote operator situated in a centralised control centre: Kim and Ryu (2013) offer an explanation of a potential teleoperation system.

⁶ The problems posed by changes in driving conditions due to weather are evidenced by a recent future automobile technology competition in South Korea: See (Ackerman, 2014)

⁷ In the US in 2000, there were approximately 256, 000 crashes involving wildlife (Goodall, 2014b, p. 95)

⁸ While some proponents of automated vehicles promote vehicle teleoperation as a possible solution in critical handover situations, we find this implausible. Effective teleoperation (where the controller will presumably be in a 1 to many relation with an automated fleet) in response to dangerous situations brought about by automated system malfunction or unforeseen hazards in the driving environment may prove even more problematic than takeover control by the human driver. The concerns with the human supervisor's delayed return to situational awareness in handover situations that we discussed in section 1 persist, and may even be exacerbated, in the case of takeover via remote control, because the remote operator will have had even less opportunity than a person in the vehicle to develop situational awareness. Further, scaling up teleoperation strategies to meet the demands of a large fleet of driverless vehicles, and consumer response to vehicle remote control, would seem to make this option unlikely.

arousal. If the human driver ceases to meet some predefined set of conditions, the vehicle could insist on manual control or safely slow to a halt. This is, for instance, the approach taken by the Tesla autopilot feature, especially after the notorious crash in Florida (Lambert, 2016; Tesla, 2016). The feasibility of this solution requires that technologies not emerge to allow vehicle owners to hack or otherwise bypass these systems.

However, at a deeper level, the problem with such a “solution” is that it jeopardises most of what people are likely to consider the advantages of a driverless vehicle. What is the point of a driverless vehicle where one still needs to pay attention to the driving task? If I need to be paying attention to the road conditions in order for the vehicle to function, then I might as well be driving. We anticipate that consumers will quickly turn against autonomous driving systems that do not in fact provide the benefits they promise in terms of freeing the occupants of vehicles from the necessity of driving them.

For these reasons, we believe that the safest way to bring a driverless vehicle to market — both in terms of commercial risks and road accident fatalities — would be to design and manufacture a vehicle that possessed SAE level 4 (or 5) autonomy and did not require those inside it to pay any attention to the road conditions encountered on their daily commute in order to safely travel from point A to B.⁹ This is a much more challenging task because such a vehicle would need to be capable of handling sudden changes in road and environmental conditions and a wide range of unusual circumstances as well as human beings do. While driverless vehicles are some way from meeting this standard yet (Goodall, 2014a, p. 990; Gordon and Lidberg, 2015), as per the frequent prognostications of engineers it seems highly likely that autonomous vehicles will *eventually* prove much safer on both sealed rural roads and city streets than vehicles driven by human beings.¹⁰ There is no reason to believe that the driving record of even the best human being represents the upper limit of performance in this task. Moreover, we do not believe that an autonomous vehicle needs to be “perfect” — in the sense of capable of always avoiding collisions — in order to be ethical. If driverless vehicles would produce fewer road fatalities than would human drivers driving at the same speeds, this will establish a strong public policy case for trying to shift to a fully driverless vehicle fleet as soon as possible.¹¹ The human cost of the road toll is tremendous. The social and economic cost of caring for those who are the victims of road trauma in both the short and the long term is a significant drain on the public purse (World Health Organization, 2015, p. ix). Governments will have strong ethical, public policy, and economic reasons to incentivise the uptake of driverless vehicles once their introduction would reduce the road toll.

In short, we are inclined to believe that the story engineers tell about driverless vehicles — that they will reduce the road toll by virtue of being safer than human drivers — will eventually prove true. The implications of this denouement when it arrives are, however, more radical than is usually acknowledged.

⁹ According to (Wakabayashi, 2016), this is the approach Google (now Waymo) is adopting to the development of driverless vehicle technology.

¹⁰ *When* this will occur is another matter: our argument is agnostic on this question.

¹¹ Note that one way that introduction of driverless vehicles will facilitate a reduction in the road toll is by moving the occupants of vehicles that have a sole occupant on a particular journey into the rear seat. This simple change in location should be expected to significantly reduce the road toll.

2. When human beings are like drunk robots

While individuals arguably have a right to risk their own lives, their right to impose risks on non-consenting third parties is much weaker. It would be tempting to say that they have *no* such right were it not for the fact that our current and historical practice regarding the automobile gives lie to this claim so obviously. Every time a driver gets behind the wheel of a car they put the lives of other people at risk. Moreover, numerous laws and engineering standards tolerate and endorse the imposition of risks on third parties well beyond those that would be imposed by alternative institutions (for instance, enforcing a lower speed limit).

Nevertheless, there is clear evidence in the history of the evolution of the law regulating driving that the public resents the imposition of risk that is not reasonably understood to be necessary to securing the goods provided by motor travel (namely, transport and convenience). In particular, the robust public support for laws that prohibit driving whilst under the influence of drugs or alcohol testify to the fact that people don't like it when drivers place them at an elevated risk of death or injury.¹² Hostility towards the imposition of risks on third parties may also be perceived in the increasing support for laws regulating passive smoking (Borland et al., 2006, p. 39).

If vehicles without a human being at the controls are safer than vehicles with a human being at the controls, then the moment a human being takes the wheel they will place the lives of third parties — as well as their own lives — at risk. Moreover, imposing this extra risk on third parties will be unethical: the human driver will be the moral equivalent of a drunk robot. Eventually, we believe, the compelling moral argument against human drivers will be reflected in law: driving will be made illegal.

We can imagine at least two circumstances that will provide an institutional impetus to prohibit human drivers.

First, and most obviously, a test case might arise in which a human being is in control of a vehicle that is also fitted with a state-of-the-art autonomous driving system when it causes someone other than the driver to be severely injured. The injured party — perhaps an occupant of another vehicle, a pedestrian, or a passenger in the vehicle — will then sue the driver, insisting that the latter was negligent in taking the wheel. If it can be shown either that the autonomous driving system was much less likely to cause a crash than the human driver, in general, or in this particular situation, then the plaintiff's case is likely to be granted.¹³ In the short term, this will establish a precedent that to take control of a vehicle fitted with autonomous driving software is to open oneself to significant legal jeopardy. In the long-term, it is likely to generate legal and political pressure to make it illegal to take the wheel of a vehicle on a public road.¹⁴

¹² See for example: (National Highway Traffic Safety Administration, 2011); (Department for Transport, 2015)

¹³ Data on the relative safety of driverless versus manual vehicles will very quickly become available once driverless vehicles are on the roads. Moreover, given the amount of data and telemetry they will generate, information about the performance of driverless vehicles in various road and traffic conditions will be extremely fine-grained. As we have argued here, if driverless vehicles are *not* safer than human drivers they should not be on the roads at all.

¹⁴ It might be objected that concern about the road toll has not generated public support for the mandatory introduction of other technologies, such as speed limiters, which would also greatly reduce the number of people killed on the roads. We believe that driverless vehicle technology is likely to be treated very differently to speed limiters by the public. Where speed limiters appear as limits on the functionality of cars, which disadvantage the driver, driverless vehicle technology presents itself —

Second, an injured party or their relative, or perhaps the government itself, might take the manufacturers of vehicles *without* autonomous driving capability and/or those that allow optional human control of the vehicle to court for manufacturing an unsafe product. Once autonomous driving systems become safer and more reliable than human drivers, placing a human being at the control of the vehicle will produce casualties that were foreseeable. Indeed, it would be unreasonable to believe that this would *not* place people's lives at risk.¹⁵ At some point, then, providing the option of manual control will fail the 'risk-utility test' and will be a violation of the standard of care that manufacturers owe consumers (Marchant and Lindor, 2012, pp. 1323-1324).¹⁶

It is worth noting that financial incentives may play an independent role in motivating the removal of the option for manual driving once driverless cars are safer than human drivers. By virtue of being less prone to accidents, driverless vehicles will be cheaper to insure (Fagnant and Kockelman, 2015a, p. 174).¹⁷ Hire companies will prefer to rent out cars without steering wheels and these will be cheaper to rent than vehicles that allow the option of driving. Similarly, insurance policies will be cheaper for car owners or users who renounce the option of driving. Finally, it will be cheaper to manufacture vehicles that only offer autonomous driving than those that offer the option of autonomous or manual operations, simply because the latter requires extra systems and components. Even if driving remains legal for some time, then, the fact that it will be more expensive than letting the machines take us from place to place may lead to a rapid decline in the number of vehicles that include the option of driving.

3. Some subtleties

While the broad brush-strokes of the account we have offered here are accurate there are two subtleties that we have neglected thus far for the sake of ease of exposition. The first concerns the standard of performance required of driverless vehicles before it will be ethical to promote them to consumers. The second concerns the policies that are likely to emerge as driverless vehicles come onto the market.

Because a small percentage of drivers, who are young or reckless, are responsible for most road crash fatalities (World Health Organization, 2015; Jonah, 1986), driverless vehicles may be safer than

and is promoted— as an enhancement of the functionality of cars which free drivers from the burden of having to drive. Consequently, risking the lives of third parties by refusing to use driverless technology is likely to be viewed with much less sympathy by other members of the public than refusal to fit a speed limiter. Our thanks to an anonymous referee for encouraging us to respond to this objection.

¹⁵ Under the stated conditions, designing a level 4 (or 5) automated vehicle with the capacity for human takeover seemingly conflicts with the U.S. Department of Transportation's 'Federal automated vehicles policy' with regard to risk assessment of 'safety-critical system functionality': see (NHTSA, 2016, p. 21).

¹⁶ This violation is interpreted in accordance with the third category of defect—design defect—outlined in § 2 of the *Restatement (third) of Torts: Products liability* (American Law Institute, 1998). See also (Villasenor, 2014). Note that in this scenario, the aggrieved party might even be the driver of a vehicle that allowed the option of taking the wheel. While an appropriately worded and strong end user agreement might protect manufacturers from suits by those who had purchased their products, it's hard to see how they could prevent injured third parties from taking them to court to argue that manual vehicles are inherently unsafe.

¹⁷ There is already some evidence that requiring all vehicles to be fitted with systems that achieve only a partial automation of the driving task would produce a net social benefit. See: (Harper et al., 2016)

the average driver without being safer than most drivers. This means that our answer as to the threshold of safety required for the introduction of driverless vehicles to be ethical turns on a larger question about the kind of ethical theory appropriate to public policy. For consequentialists, who believe that the test of public policy is its overall social consequences, the calculation will be straightforward and as we described it above: driverless vehicles should be introduced the moment doing so will reduce the road toll. That is to say, they should be introduced when they are safer than the average driver. However, this conclusion requires a willingness to trade-off the interests of different individuals of the sort that utilitarianism is notorious for and for which it has been roundly criticised (Rawls, 1971, pp. 22-29; Williams, 1973). If we focus on the situation of the vast majority of drivers who might actually be exposed to a (slightly) higher risk of death or injury in a driverless vehicle it becomes apparent that we would be trading off their safety – and occasionally their lives – for benefits that will accrue to others. If they were adequately informed about the implications for their personal safety, then, presumably they would not wish to travel in a driverless vehicle. Those inclined to endorse a more “Kantian” ethics, which emphasises our obligations to particular individuals understood as ends in themselves, may therefore be moved to conclude that we should not encourage the introduction of driverless vehicles until they are safer than most drivers — or perhaps not until they are safer than even the best human driver.

While this is an interesting irruption of a philosophical debate into a real-world policy context, this particular problem is unlikely to delay the introduction of driverless vehicles — or the arrival of the date at which it should be made illegal to drive — for very long. If it’s possible to manufacture a driverless vehicle that is safer than the average human driver then it will be possible to manufacture a driverless vehicle that is safer than even the best human driver. While the challenges involved in overcoming the first hurdle are profound, the difference between the first and the second hurdle are merely a matter of degree: given the nature and rate of progress in computer science and information technology, we should expect this gap to be closed quickly.¹⁸

Even if this turns out not to be the case, there are strong reasons for thinking that when it comes to questions of public policy — and especially questions related to public health — governments should be inclined to reason as utilitarians (Goodin, 1995). Any transport policy places some people at risk: delaying the introduction of driverless vehicles that were safer than the average driver because they weren’t yet as safe as the best driver would risk the lives of many to avoid risking the lives of a few.

The implications of a second epicycle on the larger trajectory of the argument above are more substantial. Even if governments and insurers would prefer everyone to use driverless vehicles they will not be able to achieve this overnight. Because that the average car owner waits approximately a decade before replacing their car (Litman, 2015, p. 12; European Automobile Manufacturers Association, 2016; IHS Markit, 2014), most cars on the road will still require a human driver for many years after driverless cars become available.¹⁹ Some drivers will remain emotionally or aesthetically attached to their cars and will want to keep driving them even when new models are much safer. It is therefore hard to imagine that governments will be able to secure political support for making manual driving *simpliciter* illegal even when doing so would lower the road toll. For these reasons, it

¹⁸ Ford Motor Company aim to launch commercial SAE level 4 autonomous vehicles, which have no option for human take-over, by 2021 (Bhuiyan, 2016). BMW have nominated a similar release date for their fully autonomous vehicle, and the recent advances in machine learning has seen multiple companies producing “autonomy technology” comparable to that of Google and Tesla (Condliffe, 2016). For a list of recent predictions regarding the launch dates of autonomous vehicles, see http://www.driverless-future.com/?page_id=384.

¹⁹ Cars in the EU are on average 9.8 years old (European Automobile Manufacturers Association, 2016) and in America 11.4 years old (IHS Markit, 2014).

is most plausible to think that laws will be passed to require all *new* vehicles to be provided with the capacity for autonomous driving and with no option for manual driving. Alternatively, it might simply be made illegal, while on a public road, to manually drive a car that had the capacity to operate autonomously.

Both of these options concede that the vehicle fleet will remain mixed for a significant period of time after the date on which driverless vehicles become safer than vehicles driven by human beings. This concession is regrettable because there are strong reasons for moving to an entirely driverless vehicle fleet as soon as possible.

One of the most difficult challenges for SAE Level 4 and Level 5 autonomous vehicles is predicting and responding to the actions of human drivers. In comparison, the actions of other driverless vehicles will almost certainly be predictable and their intentions transparent to driverless vehicles, which will most likely be in constant communication with each other in order to share information about traffic conditions and coordinate their use of the roads (Ferrerias, 2013, pp. 409-410; Fagnant and Kockelman, 2015a, pp. 170-171). The presence of human drivers on the roads will also pose another — more profound — challenge to driverless vehicles, owing to the likelihood that human drivers will quickly learn to take advantage of the programmed behaviour of their driverless counterparts. If autonomous vehicles are designed to prioritise the safety of their occupants as well as of other persons, then human beings may learn to cut in front of autonomous vehicles — or not to give way to them — on the assumption that the autonomous vehicle will give way every time in order to avoid an accident.²⁰ For both these reasons, driverless cars would be much safer if they did not need to worry about the foibles of human drivers.

Mixing driverless vehicles and human drivers also poses challenges for human drivers. Most of the collisions that have occurred in trials of driverless vehicles to date have occurred where driverless vehicles stop unexpectedly, for example at an amber traffic light or congested intersection, when human beings do not anticipate that they will (Richtel and Dougherty, 2015). People rely on an array of cues beyond the trajectory of the vehicle to safely navigate traffic (Richtel and Dougherty, 2015, p. 5; Sivak and Schoettle, 2015). We communicate our intentions through an array of bodily performances and we can perceive when a driver may be distracted by the crying child beside them, the ringing phone, or their friend in the back seat and take account of these factors (Joint Standing Committee on Road Safety, 2016, pp. 22, 39). The absence of such cues from driverless vehicles poses a problem for humans (drivers, cyclists, pedestrians) trying to anticipate the actions of a driverless vehicle.²¹

The difficulty of safely integrating driverless vehicles with vehicles driven by human beings in comparison with ensuring passenger safety in a transport fleet composed entirely of driverless

²⁰ This phenomenon, where individuals change their risk-taking behaviour as a result of being in an altered environment that makes them feel safer, is known as “offsetting behaviour” or “risk compensation” see (Litman, 2015, p. 5; Summala, 1996; Noland, 1995). A related problem is that pedestrians may learn that they can safely step out in front of autonomous vehicles in the expectation that the vehicles will stop or swerve. This problem obviously cannot be solved just by removing human drivers from vehicles (indeed it is caused by it) and so represents a deeper challenge for the introduction of autonomous vehicles. It is, for instance, hard to see how we could justify the collective decision that pedestrians crossing outside of authorised crossings should expect to be struck by vehicles. Yet absent this expectation, there is little to disincentivise pedestrian interference with the operations of autonomous vehicles.

²¹ The problem of how driverless cars will signal their intentions to other users of the transport environment is part of the larger ‘human machine interface’ challenge associated with automated vehicle performance: see (NHTSA, 2016, p. 22).

vehicles is likely to significantly defer the date by which we will realise the full benefits of autonomous vehicles. If governments were willing to make manual driving illegal as soon as driverless vehicles would be safer than human drivers are today, in a circumstance in which there were only driverless vehicles on the roads, this would save tens of thousands of lives. This is clearly unlikely, not least because driverless vehicles are currently being tested in real-world driving conditions alongside human drivers in large cities already and also because it would be difficult to establish what road conditions would be like with only other autonomous vehicles on the road. Nevertheless, the potential to drastically reduce road fatalities by making manual driving illegal under the conditions mentioned above, is still a significant policy consideration.

4. Implications

Let us now turn to consider some of the implications of the introduction of autonomous vehicles for the future of the automobile and for the future of the city. Our claims here are, of necessity, somewhat speculative but are, we believe, reasonable and are each supported by observations made elsewhere in the literature about driverless vehicles.

The first, which we will mention only briefly, concerns the future of motor vehicle insurance. Insofar as driverless vehicles are in fact “driven” by software all the vehicles running the same software will effectively have the same driver. That is, the engineering team responsible for the software will be responsible — ethically if not legally — for all of the accidents in which those vehicles are involved (Marchant and Lindor, 2012, p. 1326). This fact explains philosophers’ fascination with the question as to how autonomous vehicles should resolve this “trolley problem” that will inevitably arise once enough autonomous vehicles are on the roads, wherein the computer must decide who lives or dies in circumstances where the death of someone is unavoidable (Lin, 2015, p. 78). It also suggests that eventually it will be the vehicle manufacturers, or perhaps their software providers, rather than individual vehicle owners, who will be insured against damages arising from vehicle accidents. Corporations producing driverless vehicles will have a much greater understanding of the risks to which they are exposed than individual consumers, especially given that they will have access to very reliable data about the performance of their vehicles and the rate of accidents, which in turn suggests that they will be much better placed in negotiations with insurers regarding the price of insurance. The introduction of driverless vehicles is therefore likely to be extremely disruptive to the motor vehicle insurance industry.

Indeed, second, the advent of autonomous vehicles is likely to be extremely disruptive to the motor vehicle industry as a whole. For many people, especially men, the pleasures of car ownership are intimately connected with the experience of driving. When a driver is behind the wheel of a car it becomes an extension of them and driving a means whereby they can exercise agency and express their personality (Gao et al., 2014, p. 1; Litman, 2015, p. 6). Once the car becomes a transport “pod” we expect that this link will be broken; people will simply pay much less attention to the vehicles in which they travel. We anticipate that this will greatly reduce consumer demand for private motor vehicles. More importantly, the economics of owning a car will change dramatically once cars can drive themselves. Private motor vehicles are currently unproductive assets 95% of the time, when they are not being driven (Bates and Leibling, 2012, p. 24). Autonomous vehicles, though, even if privately owned, can be rented out to produce income by driving other people around, when they are not needed by the owner. In effect, even private motor vehicles will become part of a general transport pool. Access to this pool will, we believe, be facilitated by various companies selling

contracts for “transport services”, which will also own fleets of autonomous vehicles themselves.²² Such contracts will guarantee consumers certain trip times over certain distances at particular times of the day. When an individual drops a destination pin on an “app” on her phone and nominates a particular arrival or departure time, an autonomous vehicle will be dispatched to provide the necessary transport. Depending on the contract, the vehicle might also collect and drop off other people travelling similar routes. Signing up with such a provider will offer most of the convenience of the motorcar at a fraction of the expense. Perversely, then, the introduction of driverless vehicles may be extremely bad for the motor industry greatly undercutting demand for their products.²³

Third, the development of sophisticated Level 4 and Level 5 autonomous vehicles will have dramatic implications for the future of cities, architecture and urban planning. When a car can come to you, parking at your home or destination are unnecessary. Vehicles could be stored in central locations from which they drive themselves when required. As a consequence, private homes could convert garages into living or storage space, and governments and urban planners could therefore aspire to higher housing densities (Guerra, 2016, p. 215). Office towers, cinemas, and shopping centres could do away with the carparks that currently encompass them and isolate them from the surrounding community, providing instead a small plaza where people could be dropped off and collected by their vehicles. All of these benefits would be realised to an even greater degree if, as we have suggested is likely to be the case, the advent of driverless vehicles reduces the total number of vehicles in the vehicle fleet.

Fourth, as their numbers increase relative to ordinary vehicles, driverless vehicles will also enable innovative uses of existing roads to decrease journey time *or* allow higher traffic densities.²⁴ Driverless vehicles in communication with each other will be able to safely reduce the distance between vehicles because all members of the resulting “convoy” will be able to brake or accelerate simultaneously as required (Shladover et al., 2012; Fagnant and Kockelman, 2015a, p. 170; Gordon and Lidberg, 2015, p. 970). Central route allocation services will be able to distribute traffic across the road network to minimise congestion (Fagnant and Kockelman, 2015a, p. 170). Eventually, when there are no human drivers, traffic lights, give way signs, street signs, speed limit signs, and a good deal of other physical infrastructure will no longer be necessary, being replaced with digital versions in a “cloud”: cars will adjust their speeds automatically to safely handle road and traffic conditions and will “bid” for the right to cross intersections at a particular time in an electronic auction conducted amongst all the vehicles approaching the intersection (Zohdy and Rakha, 2016; Levin and Boyles, 2015; Levin et al., 2016). Importantly, however, these benefits will only be realised if governments and manufacturers can agree upon a set of standards that will allow vehicles built by different manufacturers to communicate safely and effectively.²⁵

Fifth, although, as we have argued, the advent of driverless vehicles seems likely to reduce the total number of vehicles in the fleet, it will also increase the average number of hours each day those

²² For discussion of shared autonomous vehicle modelling see: (Fagnant and Kockelman, 2015b; Fagnant and Kockelman, 2014).

²³ The opportunity to massively reduce the number of vehicles in the fleet is, as we will argue further below, potentially one of the most profound benefits made available by this technology. For some discussion of why this would be a benefit, see: (Chakraborty, 2009); (Künzli et al., 2000).

²⁴ As we discuss further below, which of these outcomes will be realised will depend crucially upon the public policy choices governments make regarding the transport infrastructure and the legal and financial incentives that will shape the public's adoption of driverless vehicle technology.

²⁵ The process of development of such standards is ongoing. See SAE J2735 (http://standards.sae.org/j2735_201603/) and the United States Department of Transport connected vehicles program (https://www.its.dot.gov/cv_basics/).

vehicles that do exist are on the roads. Once cars can drive themselves they are likely to make more — and longer — journeys. In particular, relieving human beings of the driving task will make transport by light motor vehicle significantly more available to those who are currently restricted by an inability to drive or a lack of confidence with driving, including those without driver's licenses, the elderly, people with disabilities, and children. A recent paper estimated that uptake of transport using driverless vehicles by individuals in the first three groups might increase the vehicle miles travelled by up to 14% (Harper et al. 2016). Various other shifts in commuter and consumer behaviour may also increase the number of miles travelled by driverless cars compared to existing vehicles. Every adult in what are currently single car households may choose to travel via driverless vehicles provided by transport services companies. If people are free to concentrate on other tasks while a driverless vehicle takes them to and from work, they may be willing to tolerate longer commutes. People may order more consumables online once the items they order can be delivered by driverless cars rather than collected from stores or post offices. Absent government intervention to encourage other ways of using driverless vehicle technology to travel, all of these dynamics may be expected to reduce — although not entirely eliminate — many of the social and environmental benefits that might otherwise accrue from the adoption of driverless vehicle technology.

Finally — and conversely — perhaps the most dramatic implication for the future of the city is the potential that driverless vehicles offer to allow a universal and efficient public transport system. Trains and light rail can carry high volumes of passengers but only between stations, which means that most people must travel from their home and workplace to the station by another means. The inconvenience this involves is often enough to discourage people from using trains at all (Mitchell, 2007). Combining a rail network with a fleet of autonomous buses or minibuses serving passengers arriving at and departing from each station could overcome this barrier. Passengers setting off on a journey could use a phone app to request to be collected from their home or starting point and taken to the station. Passengers arriving at a station could nominate a final destination within a certain distance. A central booking system would group individuals who have similar routes and assign them to an autonomous minibus that would then collect them or drop them off within some specified period. Combined in this fashion, trains plus autonomous vehicles could offer a “door to door” public transport service, overcoming the “last mile” problem that currently bedevils public transport systems (Ferrerias, 2013, p. 410; Mitchell, 2007).

5. Realising the benefits of a driverless world

The contemporary excitement about autonomous vehicles relates to proclamations that they will save lives, enable swifter and more efficient transport, and enhance our standard of living. However, the radical implications of the development of this technology for the future of transport and the city are not yet widely understood. In particular, as we have argued here, once autonomous vehicles pose a lower risk to third parties than human drivers there is a compelling case against the moral permissibility of humans driving.²⁶

We have identified two ways in which there may arise political pressure to have this ethical argument reflected in legislation making it illegal for human beings to drive. In a mixed fleet environment, assuming manual control of a vehicle will open individuals to legal jeopardy while

²⁶ Although we have confined our discussion to the issues raised by the development of driverless cars, we believe our conclusions generalise to other passenger vehicles, once driverless versions there-of become safer than vehicles driven by human beings.

manufacturing a vehicle that allows a human being to drive will be in breach the standard of care that manufacturers owe consumers. Nevertheless, the length of time it ordinarily takes to replace the vehicle fleet as well as the emotional commitment of (some) individuals to driving will make it politically difficult for governments to mandate autonomous driving as quickly as a concern for public safety would suggest.

We believe that the social and economic benefits of a large reduction in the road toll establish a compelling moral case for governments to provide regulatory support for a swift transition to an exclusively driverless vehicle environment. We have, further, emphasised the opportunities that driverless vehicle technology creates to establish a universal and efficient public transport system and to create more liveable urban spaces by reclaiming public space currently dedicated to carparks and freeways. We have also acknowledged that these benefits are by no means guaranteed and that the introduction of driverless vehicle technology might instead simply lead to an increase in the number of miles travelled in cars, even if the total fleet size is reduced. However, we hold that the environmental and public health benefits associated with the former scenario establish a compelling case for governments to try to bring it about. This will require both investments in physical infrastructure and regulatory incursions into the legal and economic framework surrounding private ownership and use of passenger vehicles. In particular, we believe governments should consider:

- making it illegal to manufacture vehicles that allow for manual driving beyond a certain date and/or making it illegal, while on a public road, to manually drive a vehicle that has the capacity for autonomous operations;
- investing in the provision of appropriate infrastructure for, “last mile” solutions using autonomous vehicles for public transport.
- facilitating — and perhaps even subsidising — the provision of ride-sharing services using autonomous vehicles

and, ultimately,

- working to develop incentive structures to discourage ownership of passenger motor vehicles intended for private individual use by their owners on public roads.²⁷

These suggestions undoubtedly constitute an ambitious policy agenda, especially the last item. Each and every one of these will be controversial for reasons that we do not have the space to consider here. Nevertheless, as we have argued, the case against human drivers in the future is compelling and ultimately, we suspect, public opinion against people behaving like “drunk robots” will drive governments to ban human drivers. The question then will be whether society will reap the full range of benefits made possible by the development of autonomous vehicles or whether our continuing thrall to the idea of individual private motor vehicle ownership will prevent us from doing so. This, we believe, is the real ethical challenge posed by the development of autonomous vehicles.

²⁷ Maximising the social and environmental benefits of driverless vehicle technology will require minimising a potential “rebound effect” wherein people revert to private vehicle use as mass uptake of a more efficient public transport systems supported by driverless vehicles relieves traffic congestion and therefore improves travel times for those travelling solely by car (Litman, 2015, p. 5; Ohnsman, 2014).

Acknowledgement

We would like to thank Prof. Graeme Hodge, Prof. Michael Regan, Prof. Tom Drummond, Dr. Angela Daly, Carl Liersch, Matthew Leyson, and three anonymous reviewers for comments and discussion that informed this manuscript.

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

References

- Ackerman, E. 2014. 'Korean competition shows weather still a challenge for autonomous cars', *IEEE Spectrum*, 11 November, viewed 4 December 2016, <http://spectrum.ieee.org/carsthat-think/transportation/advanced-cars/japan-competition-shows-weather-still-a-challenge-for-autonomous-cars>.
- American Law Institute 1998. *Restatement (third) of Torts: Products liability*, St. Paul, American Law Institute.
- Anderson, J. M., Nidhi, K., Stanley, K. D., Sorensen, P., Samaras, C., Oluwatola, O. A., 2016. *Autonomous vehicle technology: A guide for policymakers*, Rand Corporation, Santa Monica.
- Bamonte, T., 2013. 'Autonomous vehicles: Drivers for change', *Transport Management and Engineering*, 24 July, viewed 23 November 2016, <http://www.tmemag.com/autonomous-vehicles-drivers-change>.
- Bates, J., Leibling, D., 2012. *Spaced out: Perspectives on parking policy*, London, Royal Automobile Club Foundation.
- Bhuiyan, J., 2016. Ford plans to have a fleet of fully autonomous cars operating in a ride-hail service by 2021', *Recode*, 16 August, viewed 24 March 2017, <http://www.recode.net/2016/8/16/12500632/ford-mark-fields-self-driving-cars-2021?dom=pscau&src=syn>.
- Bilger, B., 2013. 'Auto Correct: Has the self-driving car at last arrived', *The New Yorker*, 25 November, viewed 19 November 2016, <http://www.newyorker.com/magazine/2013/11/25/auto-correct>.
- Borland, R., Yong, H.-H., Siahpush, M., Hyland, A., Campbell, S., Hastings, G., Cummings, K. M., Fong, G. T., 2006. 'Support for and reported compliance with smoke-free restaurants and bars by smokers in four countries: Findings from the International Tobacco Control (ITC) Four Country Survey', *Tobacco control*, 15, no. suppl 3, iii34-iii41.
- Burns, L. D., 2013. 'Sustainable mobility: a vision of our transport future', *Nature*, 497, no. 7448, 181-182.
- Chakraborty, J., 2009. 'Automobiles, air toxics, and adverse health risks: Environmental inequities in Tampa Bay, Florida', *Annals of the Association of American Geographers*, 99, no. 4, 674-697.
- Conliffe, J., 2016. '2021 may be the year of the fully autonomous car', *MIT Technology Review*, 17 August, viewed 25 March 2017, <https://www.technologyreview.com/s/602196/2021-may-be-the-year-of-the-fully-autonomous-car/>.
- Department for Transport, 2015. *British Social Attitudes Survey 2014: Public attitudes towards transport*, British Parliament.
- European Automobile Manufacturers Association., 2016. *Average vehicle age*, viewed 18 November 2016, <http://www.acea.be/statistics/tag/category/average-vehicle-age>.
- Fagnant, D. J., Kockelman, K., 2015a. 'Preparing a nation for autonomous vehicles: Opportunities, barriers and policy recommendations', *Transportation Research Part A: Policy and Practice*, 77, 167-181.
- Fagnant, D. J., Kockelman, K. M., 2014. 'The travel and environmental implications of shared autonomous vehicles, using agent-based model scenarios', *Transportation Research Part C: Emerging Technologies*, 40, 1-13.

Fagnant, D. J., Kockelman, K. M., 2015b. 'Dynamic ride-sharing and optimal fleet sizing for a system of shared autonomous vehicles', *Transportation Research Board 94th Annual Meeting*, 11-15 January 2015, Washington.

Ferreras, L. E., 2013. 'Autonomous vehicles: a critical tool to solve the XXI century urban transportation grand challenge', In: Jones, S. L., (ed.), *Third international conference on urban public transportation systems*, 17-20 November, Paris, 405-412.

Gao, P., Hensley, R., Zielke, A., 2014. 'A road map to the future for the auto industry', *McKinsey Quarterly*, no. October, 1-11.

Garza, A. P., 2011. 'Look Ma, no hands: Wrinkles and wrecks in the age of autonomous vehicles', *New England Law Review*, 46, 581-616.

Gold, C., Damböck, D., Lorenz, L., Bengler, K., 2013. '"Take over!" How long does it take to get the driver back into the loop?', *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*: Sage Publications, 1938-1942.

Goodall, N., 2014a. 'Ethical decision making during automated vehicle crashes', *Transportation Research Record: Journal of the Transportation Research Board*, no. 2424, 58-65.

Goodall, N. J., 2014b. 'Machine ethics and automated vehicles', In: Meyer, G., Beiker, S., (eds.) *Road Vehicle Automation*, Springer, Urdorf, 93-102.

Goodin, R.E., 1995. *Utilitarianism as a public philosophy*. Cambridge University Press.

Google., 2016. *Google self-driving car project*, viewed 11 December 2016, <https://www.google.com/selfdrivingcar/where/>.

Gordon, T., Lidberg, M., 2015. 'Automated driving and autonomous functions on road vehicles', *Vehicle System Dynamics*, 53, no. 7, 958-994.

Guerra, E., 2016. 'Planning for cars that drive themselves: Metropolitan planning organizations, regional transportation plans, and autonomous vehicles', *Journal of Planning Education and Research*, 36, no. 2, 210-224.

Harper, C. D., Hendrickson, C.T., Mangones, S., Samaras, C., 2016. 'Estimating potential increases in travel with autonomous vehicles for the non-driving, elderly and people with travel-restrictive medical conditions', *Transportation Research Part C: Emerging Technologies* 72, 1-9.

Harper, C. D., Hendrickson, C. T., Samaras, C., 2016. 'Cost and benefit estimates of partially-automated vehicle collision avoidance technologies', *Accident Analysis & Prevention* 95, 104-115.

IHS Markit., 2014. *Average age of vehicles on the road remains steady at 11.4 years, according to IHS Automotive*, IHS, viewed 11 December 2016, <http://news.ihsmarkit.com/press-release/automotive/average-age-vehicles-road-remains-steady-114-years-according-ihs-automotive>.

Joint Standing Committee on Road Safety., 2016. *Driverless vehicles and road safety in NSW*, Sydney, Parliament of New South Wales.

Jonah, B. A., 1986. 'Accident risk and risk-taking behaviour among young drivers', *Accident Analysis & Prevention*, 18, no. 4, 255-271.

Kim, J.S., Ryu, J.H., 2013. Shared teleoperation of a vehicle with a virtual driving interface. In *Control, Automation and Systems (ICCAS), 13th International Conference*, October, 851-857.

Kirkpatrick, K., 2015. 'The moral challenges of driverless cars', *Communications of the ACM*, 58, no. 8, 19-20.

Künzli, N., Kaiser, R., Medina, S., Studnicka, M., Chanel, O., Filliger, P., Herry, M., Horak, F., Puybonnieux-Textier, V., Quénel, P., 2000. 'Public-health impact of outdoor and traffic-related air pollution: A European assessment', *The Lancet*, 356, no. 9232, 795-801.

Lambert, F., 2016. *Tesla will soon introduce new Autopilot safety restrictions after recent accidents*, Electrek, viewed 11 December 2016, <https://electrek.co/2016/08/28/tesla-autopilot-safety-restrictions-v8-0-accidents/>.

Levin, M. W., Boyles, S. D., 2015. 'Intersection auctions and reservation-based control in dynamic traffic assignment', *Transportation Research Record: Journal of the Transportation Research Board*, no. 2497, 35-44.

Levin, M. W., Boyles, S. D., Patel, R., 2016. 'Paradoxes of reservation-based intersection controls in traffic networks', *Transportation Research Part A: Policy and Practice*, 90, 14-25.

Lin, P., 2015. 'Why ethics matters for autonomous cars', In: Maurer, M., Gerdes, J. C., Lenz, B., Winner, H., (eds.) *Autonomous driving: Technical, legal and social aspects*, Springer, Berlin, 69-85.

Litman, T., 2015. *Autonomous vehicle implementation predictions: Implications for transport planning*, Victoria Transport Policy Institute.

Marchant, G. E., Lindor, R. A., 2012. 'The coming collision between autonomous vehicles and the liability system', *Santa Clara Law Review*, 52, 1321-1340.

Merat, N., Jamson, A. H., 2009 'How do drivers behave in a highly automated car', *5th International Driving Symposium on Human Factors in Driver Assessment, Training and Vehicle Design*, 22-25 June, Big Sky, 514-521.

Mitchell, W. J., 2007. 'Intelligent cities', *UOC papers*, 5, 3-9.

National Highway Traffic Safety Administration, 2008. 'National motor vehicle crash causation survey: Report to congress', *National Highway Traffic Safety Administration DOT HS*, 811.

National Highway Traffic Safety Administration., 2011. 'National survey of drinking and driving attitudes and behaviors: 2008', *Annals of Emergency Medicine*, 57, no. 4, 405-406.

National Highway Traffic Safety Administration, 2016. 'Federal Automated Vehicles Policy: Accelerating the Next Revolution in Roadway Safety', *National Highway Traffic Safety Administration DOT HS*, 812 329

Noland, R. B., 1995. 'Perceived risk and modal choice: risk compensation in transportation systems', *Accident Analysis & Prevention*, 27, no. 4, 503-521.

Ohnsman, A., 2014. 'Automated cars may boost fuel use, Toyota scientist says', *Bloomberg Technology*, 17 July, viewed 12 December, www.bloomberg.com/news/2014-07-16/automated-cars-may-boost-fuel-use-toyota-scientist-says.html.

Radlmayr, J., Gold, C., Lorenz, L., Farid, M., Bengler, K., 'How traffic situations and non-driving related tasks affect the take-over quality in highly automated driving', *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*: Sage Publications, 2063-2067.

Rawls, J., 1971. *A theory of justice*, Harvard University Press, Cambridge, MA.

Richtel, M., Dougherty, C., 2015. 'Google's driverless cars run into problem: Cars with drivers', *The New York Times*, 1 September, viewed 22 November 2016, <http://www.nytimes.com/2015/09/02/technology/personaltech/google-says-its-not-the-driverless-cars-fault-its-other-drivers.html>.

Shladover, S., 2016. 'The Truth about "Self-Driving" Cars', *Scientific American*, 314(6), 52-57

Shladover, S., 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*, no. 2324, 63-70.

Singhvi, A., Russell, K., 2016. 'How self-driving cars work', *The New York Times*, 14 December, viewed 14 December 2016, http://www.nytimes.com/interactive/2016/12/14/technology/100000004821211.mobile.html?_r=1.

Sivak, M., Schoettle, B., 2015. *Road safety with self-driving vehicles: General limitations and road sharing with conventional vehicles*, Ann Arbor, University of Michigan Transportation Research Institute.

Summala, H., 1996. 'Accident risk and driver behaviour', *Safety Science*, 22, no. 1, 103-117.

Tesla., 2016. *Tesla press information: Autopilot*, viewed 11 December 2016, https://www.tesla.com/en_AU/presskit/autopilot?redirect=no#autopilot.

Thrun, S., 2010. *What we're driving at*, viewed 12 December 2016, <https://googleblog.blogspot.com.au/2010/10/what-were-driving-at.html>.

Villasenor, J., 2014. 'Products liability and driverless cars: Issues and guiding principles for legislation', 24 April, viewed 11 December 2016, <https://www.brookings.edu/research/products-liability-and-driverless-cars-issues-and-guiding-principles-for-legislation/>.

Vlasiv, B., 2016. 'Tesla will make its cars fully self-driving, but not turn the system on yet', *The New York Times*, 19 October, viewed 14 December 2016, <http://www.nytimes.com/2016/10/20/business/tesla-autonomous-autopilot-vehicles.html>.

Wakabayashi, D., 2016. 'Google parent company spins off self-driving car business', *The New York Times*, 13 December, viewed 14 December, <http://mobile.nytimes.com/2016/12/13/technology/google-parent-company-spins-off-waymo-self-driving-car-business.html>.

Williams, B., 1973. 'A critique of utilitarianism', In: Smart, J. J. C. & Williams, B. (eds.) *Utilitarianism: For and Against*, Cambridge University Press, Cambridge, 75-150.

World Health Organization, 2015. *Global status report on road safety 2015*, Switzerland, World Health Organization.

Zohdy, I. H., Rakha, H. A., 2016. 'Intersection management via vehicle connectivity: The intersection cooperative adaptive cruise control system concept', *Journal of Intelligent Transportation Systems*, 20, no. 1, 17-32.