Michelle Landry, MP
Standing Committee on Industry, Innovation, Science and Resources
PO Box 6021
House of Representatives
Parliament House
Canberra ACT 2600

RE: Inquiry into the social issues relating to land-based driverless vehicles in Australia

Dear Ms Michelle Landry MP,

We write in relation to the recent call for submissions to the parliamentary enquiry into the social issues relating to land-based driverless vehicles in Australia. Our submission addresses the following terms of reference: 1a, social acceptance levels; 1b, passenger and non-passenger safety; 1c, legal responsibility and insurance; 1d, potential impacts on employment and different industry sectors; 1e, access and equity issues; 1f, potential public transport applications and; 2, the opportunities and challenges for each social issue. We offer several recommendations to progress action on the social issues identified.

Yours sincerely

Professor Robert Sparrow and Dr Mark Howard

Tel.: [redacted]
E-mail address: [redacted]
Department of Philosophy
School of Philosophical, Historical and International Studies
Monash University

Social issues relating to land-based driverless vehicles in Australia

We believe that public policy in relation to driverless vehicles should be developed in the light of the following considerations: we have highlighted considerations particularly relevant to the terms of reference of the Committee via the inclusion of text in bold.

At some point in the future, autonomous driving systems will be safer than human drivers. At that point, the rapid introduction of driverless vehicles on Australian roads will save thousands of lives. Once autonomous driving systems are safer than human beings when it comes to the risk to third parties, *human beings will be the moral equivalent of "drunk robots"* and it should therefore be illegal for human beings to be in control of a powered vehicle on a public road.

While the number of years it will take to reach this standard of performance is contestable, we believe that it will eventually be reached [**technological readiness]**. There is no reason to believe that human beings represent the optimal solution to the task of controlling a vehicle. Computers already outperform human beings in multiple domains.

Moreover, until autonomous driving systems reach this standard, it should be illegal to employ them on public roads.

In short, either fully autonomous vehicles will be safer than human beings, in which case humans shouldn't be allowed to drive, or they won't be safer than human beings in which case they shouldn't be allowed on the roads.

The development and installation of various forms of "driver assist" systems, such as Tesla's "autopilot" might be thought to offer a way of gradually introducing more and more "autonomy" into motor vehicles without confronting this dilemma.

However, we believe there are significant limits on the extent to which these systems, which rely upon supervision from a human driver in order to handle rare events or road conditions outside of the capacity of the system to respond to effectively, may be introduced without creating new risks to drivers and third parties **[passenger and non-passenger safety]**. There is extensive evidence from studies of human computer interaction that human beings quickly cease to pay attention to matters that are not directly relevant to the tasks in which they are engaged. Thus, if a human driver is required to retake control of a vehicle when a normally reliable driver assist system fails at high speeds, they may be ill placed to do so. Similarly, human beings typically over-rely on systems that perform well in ordinary circumstances. Thus, we anticipate that once driverless vehicles are reliable enough, people will let the car drive itself while they respond to their email, fall asleep, or enjoy a few drinks. They may even let the car drive their children to school without them. Consequently, if the system does suddenly require input from a human "supervisor" in an unusual circumstance there may be no one capable of playing that role in time. Once enough driverless vehicles are on the road, of course, unusual circumstances will arise every day. It is also worth noting that relying regularly on driver assist systems is likely to produce significant "deskilling" of drivers with an accompanying decrease in their performance when it is required.

Systems that require drivers to actively supervise the driving task and that are capable of detecting when the driver is no longer paying attention will reduce the risks of accidents produced by this dynamic but at a significant cost to the utility of driverless vehicle technology. The effectiveness of these systems will depend upon them not being able to be hacked or bypassed. They are also unlikely to reduce the extent of driver deskilling that will

occur when drivers regularly rely upon computers for the routine operations of vehicles. Requiring a driver to play an active role supervising the driving of a vehicle will radically reduce the benefits that driverless vehicle technology might otherwise offer to people in advanced old age, those with disabilities, or those suffering cognitive impairments **[access and equity issues]**. Finally, we expect that consumers will rapidly turn against "driverless" vehicle technology that fails to free them from the driving task **[general social acceptance levels]**.

These dynamics provide a very strong incentive for manufacturers to move quickly to fully autonomous vehicles. Such vehicles need not be "perfectly" safe in order to be ethical. As noted above, we believe the ethical standard for allowing such vehicles on the roads, at least at a level of public policy, is when their introduction would reduce the road toll. That is, when they produce fewer road fatalities on average than human drivers. Data regarding this level of performance will quickly be available, given that the average number of fatalities per vehicle hour caused by human drivers is relatively easily calculable and that autonomous vehicles may be expected to return data about their performance to their manufacturers as a matter of course.

The potential of driverless vehicle technology to reduce the road toll establishes a strong ethical and public policy imperative to move as quickly as possible to fully driverless fleet as soon as driverless vehicle technology reaches this standard.

Moreover, we anticipate that there will be significant public support for removing human drivers from the roads as soon as it can be shown that the risks to 3rd parties posed by driverless vehicles are less than the risks posed by human drivers **[general social acceptance levels]**. Just as the public has become increasingly hostile to people who kill others by driving a vehicle when they are drunk, we expect that they will become hostile to those who kill people by taking the wheel when an autonomous driving system would have avoided the accident. Pressure on governments to prohibit human driving may also arise via court decisions that conclude that drivers involved in accidents were negligent in taking the wheel or that manufacturers who produce vehicles that allow drivers to take the wheel are manufacturing unsafe products.

However, even if governments wish to move rapidly to a fully driverless vehicle fleet, they may struggle to achieve this given the relatively slow rate at which households replace their vehicles. For this reason, we suggest the most plausible public policy to achieve the goal of a fully driverless fleet would be to require all vehicles sold beyond a certain date to have the capacity for fully autonomous operations. More ambitiously, state governments might declare some years in advance that it will be illegal to be in direct control of a vehicle on public roads beyond a certain date **[role of government]**.

The likelihood that the vehicle fleet will remain mixed and consist of driverless and ordinary vehicles for a number of years will itself delay the introduction of driverless vehicles significantly. Predicting the actions of human drivers is itself one of the most difficult tasks for driverless vehicles. Similarly, human drivers are likely to struggle with sharing the road with autonomous vehicles given that the latter will often fail to provide the subtle interpersonal communication signals that allow human beings to anticipate each other's actions. In contrast, driverless vehicles will be able to communicate electronically with other vehicles in order to avoid accidents and share the roads much more efficiently than vehicles driven by human beings can. It is also worth noting that if driverless vehicles are programmed to avoid collisions it will be relatively easy for human beings to assert their own right of way over driverless vehicles by adopting a driving style that requires the latter to take evasive action and therefore defer to them.

Given that the introduction of driverless vehicles as soon as possible would save thousands of lives, the difficulties human drivers pose to the performance of driverless vehicles adds further weight to the case that governments should be working to remove human drivers from the roads as soon as possible once driverless vehicle technology matures to the point that a driverless fleet would produce fewer road fatalities than a mixed fleet.

The introduction of driverless vehicle technology is likely to be extremely disruptive both to the motor vehicle insurance industry and to the motor vehicle industry as a whole.

Once vehicles are driven by software, the manufacturer of that software will effectively become the driver of all the vehicles that use it **[legal responsibility and insurance]**. This shift in moral responsibility for road accidents from drivers to engineers explains philosophical and popular interest in the question of how driverless cars will resolve various sorts of "trolley problems" that will inevitably occur once sufficient driverless vehicles take to the roads. It also suggests that eventually insurance to protect against the cost of motor vehicle accidents will be primarily purchased by manufacturers rather than individual owners, thus radically reshaping the current motor vehicle insurance industry**.**

Once people cease to drive their cars, we believe that they will be much less likely to identify with them as a source of consumer satisfaction. Instead, many people will come to view cars merely as a source of transportation. Moreover, when cars can drive themselves private motor vehicle ownership makes much less sense. A car that can drive itself can be earning income when its owner does not require its services as long as the owner is willing to rent it out to others. Such arrangements will be subject to significant economies of scale. Essentially, companies providing "transport services", wherein subscribers purchase a guarantee of particular trip times over particular distances and are then collected by an autonomous vehicle at a location they nominate on their phone and taken to their destination, will be able to outcompete private motor vehicle ownership. If people are willing to share vehicles with others who are travelling along the same route this will reduce the cost of their transport even further. This will obviously be extremely disruptive to the motor vehicle industry as a whole by virtue of greatly reducing demand for private vehicles **[potential impacts on employment and different industry sectors]**.

The potential for driverless vehicle technology to massively reduce the number of vehicles on the roads is perhaps the most significant benefit after their potential to reduce the road toll. The social, environmental, and economic costs of high levels private vehicle ownership are enormous. Again, we believe that this establishes a significant public policy incentive to encourage the adoption of driverless vehicle technology and especially its adoption for ride sharing services of the sort intimated above.

Relatedly, the introduction of driverless vehicles will offer remarkable opportunities to redesign urban landscapes and transport infrastructures to produce much more liveable environments. If cars can drive themselves there is no need for them to be parked at either the beginning or the end of a trip. Instead, cars could travel from nearby locations to collect passengers and park themselves (or indeed go on to collect other passengers) after dropping them off. Private houses would no longer need garages and popular destinations would no longer need to be surrounded by asphalt. Congestion on the roads would be reduced not only by the decrease in the number of vehicles but also by using inter-vehicle communication to dynamically plan routes to facilitate more efficient traffic flows. Indeed, more radically there would be no need for the majority of commutes into and out of urban centres to be made by motor vehicle at all. Instead, people could be picked up by a fleet of autonomous vehicles associated with the local railway station and delivered to the station to continue their journey by public transport **[potential public transport applications]**. At the end of their commute

on public transport they could be taken by driverless vehicle to their ultimate destination. Presuming that people nominate their pickup address and destination using an app on the phone, or individuals travelling similar routes could be transported by a single driverless minibus assigned to get them where they are going on time. Driverless vehicle technology therefore has the potential to solve the "last mile" problem that currently bedevils public transport systems.

Unfortunately these most profound social benefits that might be produced by the introduction of driverless vehicles are unlikely to be realised without significant input from government.

In particular, as the technology for an autonomous driving nears fruition **[recommendations to progress action on the social issues identified]**, we believe governments should consider:

- making it illegal, beyond a certain date, to sell vehicles that allow human beings the option of assuming manual control and/or making it illegal for a human being to drive a vehicle that has an autonomous driving capability.
- Support the development of technical standards to facilitate communication amongst autonomous vehicles so that they may better avoid collisions, travel in convoys, and reduce congestion by participating in collective dynamic route planning.
- Research and develop infrastructure for the use of autonomous vehicles to provide "last mile" solutions for public transport
- supporting and perhaps even subsidising ridesharing services provided by autonomous vehicles

and

- discouraging individual private ownership of passenger vehicles intended for use on public roads.

An ambitious policy agenda, perhaps. Nevertheless, such steps are necessary if Australia is to realise the full social and economic benefits that driverless vehicles make possible.

[This submission draws heavily on research carried out in the course of writing an academic paper, "When human beings are like drunk robots", co-authored by Professor Robert Sparrow and Dr Mark Howard, and currently under consideration for publication in *Transportation Research Part C*. However, it does not reproduce significant portions of the text of that manuscript. The full manuscript is available upon request from the authors.]