Who's driving? New technologies and the collaborative state

1. Prologue: Who killed Elaine Herzberg?

Elaine Herzberg did not know that she was part of an experiment. She was walking her bicycle across the road at 10pm on a dark desert night in Tempe, Arizona. Having crossed three lanes of a four-lane highway, Herzberg was run down by a Volvo SUV travelling at 38 miles per hour. She was pronounced dead at 10:30pm.

The next day, the officer in charge of the investigation rushed to blame the pedestrian. Police Chief Sylvia Moir told a local newspaper, 'It's very clear it would have been difficult to avoid this collision… she came from the shadows right into the roadway… the driver said it was like a flash.' According to the rules of the road, Herzberg should not have been there. Had she been at the crosswalk just down the road, things would probably have turned out differently.

Rafaela Vasquez was behind the wheel of the Volvo, but she wasn't driving. The car, operated by Uber, was in autonomous mode. Vasquez's job was to monitor the computer that was doing the driving and take over if anything went wrong. A few days after the crash, the police released a video from a camera on the rear-view mirror. It showed Vasquez looking down at her knees in the seconds before the crash and for almost a third of the 21-minute journey that led up to it. Data taken from her phone suggested that she had been watching an episode of 'The Voice' rather than the road. Embarrassingly for the police chief, her colleagues' investigation calculated that, had Vasquez been looking at the road, she would have been able to stop more than 40 feet before impact.[1]

Drivers and pedestrians make mistakes all the time. More than 90% of crashes are blamed on human error. The Tempe Police report concluded that the crash had been caused by human frailties on both sides: Herzberg should not have been in the road; Vasquez for her part should have seen the pedestrian, she should have taken control of the car and she should have been paying attention to her job. In the crash investigation business, these factors are known as 'proximate causes'. But if we focus only on proximate causes, we fail to learn from the novelty of the situation. Herzberg was the first pedestrian to be killed by a self-driving car. The Uber crash was not just a case of human error. It was also a failure of technology.

Here was a car on a public road in which the driving had been delegated to a computer. A thing that had very recently seemed impossible had become, on the streets of Arizona, mundane, so mundane that the person who was supposed to be checking the system had, in effect, switched off.[2] The car's sensors – 360-degree radar, short- and long-range cameras, a lidar laser scanner on the roof

and a GPS system – were supposed to provide superhuman awareness of the surroundings. The car's software was designed to interpret this information based on thousands of hours of similar experiences, identifying objects, predicting what they were going to do next and plotting a safe route. This was artificial intelligence in the wild: not playing chess or translating text but steering two tonnes of metal.

When high-profile transport disasters happen in the US, the National Transportation Safety Board is called in. The NTSB are less interested in blame than in learning from mistakes to make things safer. Their investigations are part of the reason why air travel is so astonishingly safe. In 2017, for the first time, a whole year passed in which not a single person died in a commercial passenger jet crash. If self-driving cars are going to be as safe as aeroplanes, regulators need to listen to the NTSB. The Board's initial report on the Uber crash concluded that the car's sensors had detected an object in the road six seconds before the crash. The software classified Herzberg 'as an unknown object, as a vehicle, and then as a bicycle', in the NTSB's words, but the car continued. A second before the car hit Herzberg, the driver took the wheel but swerved only slightly. Vasquez only applied the brakes after the crash.

As well as the proximate causes, Elaine Herzberg's death was the result of a set of more distant choices about technology and how it should be developed. Claiming that they were in a race against other manufacturers, Uber chose to test their system quickly and cheaply. Other self-driving car companies put two or more qualified engineers in each of their test vehicles. Vasquez was alone and she was no test pilot. The only qualification she needed before starting work was a driving licence.

Uber's strategy filtered all the way down into its cars' software, which was much less intelligent than the company's hype had implied. As the company's engineers worked out how to make sense of the information coming from the car's sensors, they balanced the risk of a false positive (detecting a thing that isn't really there) against the risk of a false negative (failing to react to an object that turns out to be dangerous). After earlier tests of self-driving cars in which software overreacted to things like steam, plastic bags and shadows on the roads, engineers retuned their systems. The misidentification of Elaine Herzberg was partly the result of a conscious choice about how safe the technology needed to be in order to be safe enough. One engineer at Uber later told a journalist that the company had 'refused to take responsibility. They blamed it on the homeless lady [Herzberg], the Latina with a criminal record driving the car [Vasquez], even though we all knew Perception [Uber's software] was broken.'[3]

The companies who had built the hardware also blamed Uber. The president of Velodyne, the manufacturer of the car's main sensors, told Bloomberg, 'Certainly, our lidar is capable of clearly imaging Elaine and her bicycle in this situation. However, our lidar doesn't make the decision to put on the brakes or get out of her way.'[4] Volvo made clear that they had nothing to do with the experiment. They provided the body of the car, not its brain. An automatic braking system that was built into the Volvo – using well-established technology – would almost

certainly have saved Herzberg's life, but this had been switched off by Uber engineers, who were testing their own technology and didn't want interference from another system.

We don't know what Elaine Herzberg was thinking when she set off across the road. Nor do we know exactly what the car was thinking. Machines make decisions differently from humans and the decisions made by machine learning systems are often inscrutable. However, the evidence from the crash points to a reckless approach to the development of a new technology. The company shouldered some of the blame, agreeing an out-of-court settlement with the victim's family and changing their approach to safety. But to point the finger only at the company would be to ignore the context. Roads are dangerous places, particularly in the US and particularly for pedestrians. A century of decisions by policymakers and carmakers has produced a system that gives power and freedom to drivers. Tempe, part of the sprawling metropolitan area of Phoenix, is car-friendly. The roads are wide and neat and the weather is good. It is ideally suited to testing a self-driving car. For a pedestrian, the place and its infrastructure can feel hostile. Official statistics bear this out. In 2017, Arizona was the most dangerous state for pedestrians in the US.

Members of Herzberg's family sued the state government on the grounds that, first, the streets unsafe for pedestrians and, second, policymakers were complicit in Uber's experiments. In addition to the climate and the tidiness of the roads, Uber had been attracted to Tempe by the governor of Arizona, Doug Ducey. The company had started their testing in San Francisco, near their headquarters. But when one of their self-driving cars ran a red light, California regulators told Uber that they needed a $150 permit. Uber objected and Ducey seized his opportunity. With the Governor's blessing, the company had already been testing in secret on the streets of Phoenix. Ducey could now go public and claim that he had tempted a big tech company away from Silicon Valley. He tweeted 'This is what over-regulation looks like #ditchcalifornia' and 'Here in AZ we WELCOME this kind of technology & innovation! #ditchcalifornia #AZmeansBIZ'. With almost no oversight, Uber moved their experiments to Arizona in 2016. When Herzberg was killed less than 18 months later, Ducey's enthusiasm collapsed and Uber were thrown out of their new laboratory.

When technologies fail, it is often hard to find the person responsible and easy for those involved to blame others or claim it was a freak occurrence. It's a symptom of a wider problem, which is that we aren't clear who is in control of the development of new technologies. When technological dreams meet the real world, the results are often disappointing and occasionally messy. Policymakers are often seduced by the promise of new technologies, which arrive without instructions for how they should be governed. It is all too common for regulation to be an afterthought. In the world of aviation, it's called a tombstone mentality: defects are noticed, lessons are learned and rules are written in grim hindsight. In Arizona, policymakers allowed a private experiment to take place in public, with citizens as unwitting participants. It ended badly for everyone involved. Tragedies are opportunities for learning, opportunities to challenge claims made about technology and opportunities to think about alternatives. We should ask if

a technology is safe enough, but this means also asking, Safe enough for what? Why are self-driving cars being developed? Where are they taking us? As politicians compete in their enthusiasm for innovation for innovation, such questions often go unasked. Two months after the crash, the Governor of Ohio saw his opportunity and announced plans to make his state 'the wild, wild West' for unregulated self-driving car testing.[5]

It is vital to scrutinise technologies at an early stage, before they become just another fact of life. If we agree that technology is too important to be left to technology companies, we are left with the challenge of how to democratise innovation. New technologies should prompt us to update the question posed by political scientist Robert Dahl (1961): 'Who governs?'[6] If we are to hang onto democracy in the 21st Century, we should keep asking 'Who's driving?'

The hope that powers this book is that we can do better and imagine a proactive role for policymakers and citizens. I want to make the case that new technologies can and should be redirected towards public benefit. In the next chapter, I ask why we fall into the trap of thinking that technologies are autonomous and inevitable. Chapter three is about the politics of innovation: who benefits and who decides? Chapter four focuses on hype, which is used by technology developers as a way to stake claims on the future. The final chapter is an argument for collaboration between the developers of technology, governments and citizens. The book uses self-driving cars as an example, but its messages are relevant to other new technologies, particularly those in the area of artificial intelligence, whose rules have not yet been written and whose destinies have not yet been set. If our technological future is to incorporate democratic values, policymakers will need to challenge the story of technological inevitability, resist hype, target the needs of those who often lose out from innovation and enable new collaborations to flourish.

[1] Tempe police department general offense report, March 2018

[2] Research with autopilots in aircraft has revealed the detrimental effects of automation on people's concentration in the short term and skills in the long term. The idea of the human-on-the-loop, overseeing an automated system, looks near-fatally flawed. The findings are elegantly summarised by Nicholas Carr (2014)

[3] Some Uber insiders later told journalists that the company was fixated on a smooth ride that would impress senior executives, rather than safety (Uber insiders describe infighting and questionable decisions before its self-driving car killed a pedestrian. Julie Bort Nov. 19, 2018, Business Insider, https://www.businessinsider.com/sources-describe-questionable-decisions-and-dysfunction-inside-ubers-self-driving-unit-before-one-of-its-cars-killed-a-pedestrian-2018-10)

[4] Sensor Supplier to Self-Driving Uber Defends Tech After Fatality, Keith Naughton, Bloomberg, 23 March 2018, https://www.bloomberg.com/news/articles/2018-03-23/sensor-supplier-to-self-driving-uber-defends-tech-after-fatality

[5] Governor John Kasich, Ohio, 9 May 2018. Video: https://www.nbc4i.com/news/politics/ohio-statehouse-newsroom/gov-kasich-wants-to-make-ohio-the-wild-west-of-automated-vehicle-testing/1169957664

[6] Thanks to Declan McDowell-Naylor (2018) for drawing the connections between Dahl's work and self-driving cars.