Hitesh Nandakumar
Professor Caron
ISLA 123
06/07/2026

Final Project

Waymo's autonomous vehicles have logged over 170 million fully driverless miles and provide half a million paid rides a week, a scale that makes the company's safety record genuinely impressive by the numbers. But in a single month in 2026, a software defect led Waymo to recall 3,791 vehicles after one robotaxi entered floodwater and was swept into a creek, and a separate vehicle accelerated through a San Francisco construction zone with terrified passengers on board while police vehicles gave chase (Del Rey; Tangermann). These are not just isolated accidents; they show that Waymo's development process is missing input from the people who actually live with the risks. The ethical issues they raise are best examined through the lens of deontology, and best addressed through Constructive Technology Assessment paired with design justice.

I. Technical Components and the Ethical Dilemma

Waymo's vehicles operate through a combination of sensor fusion, machine learning, and remote oversight. LiDAR, radar, cameras, and GPS continuously send data into a decision-making system trained across millions of miles to anticipate road conditions, other drivers, and pedestrian behavior. When the system encounters something it cannot resolve on its own, a remote operator can communicate with passengers through an in-cabin audio system, though the operator cannot always take direct physical control.

The recall affecting 3,791 vehicles exposes what the system was never trained to handle: flooded roads. In San Antonio, a vehicle ended up carried into a creek at highway speed because the automated driving system did not treat standing water as a serious enough hazard to override normal operation (Del Rey). Flooding in San Antonio is not an edge case; it is a documented and recurring feature of the climate there. The freeway incident reveals a different

gap. When Elliot Slade and his fiancée were riding in a Waymo that accelerated into a construction zone with police vehicles behind it, neither passenger had any mechanism to stop the car. When it finally pulled over, a remote operator came over the audio and asked if they wanted to continue the trip. Waymo later offered $40 in credits (Tangermann).

The moral dilemma here is not whether autonomous vehicles are good or bad. It is a conflict between two values that both carry real moral weight: the aggregate public benefit of safer roads, greater mobility for elderly and disabled riders, and fewer drunk drivers on one side; and the right of individuals and communities to participate meaningfully in decisions about technologies that operate on their bodies and in their neighborhoods on the other. Waymo is not doing something obviously immoral; it is doing something that produces real benefits while quietly offloading the costs of development onto people who never agreed.

II. Applying Ethical Frameworks

Utilitarianism rests on three main ideas. The first is consequentialism: what makes an action right or wrong is entirely what it produces, not the intent behind it or what kind of act it is. The outcome is everything. The second is hedonism: among all possible consequences, what makes them good or bad is the balance of happiness and pain they produce. The third is the aggregate: utility is calculated by adding the happiness and pain across everyone affected, and the correct action is the one that maximizes that total happiness. Taken together, those three ideas make the utilitarian argument for Waymo pretty strong. Waymo's internal research reports a 57% reduction in police-reported crashes compared to human drivers, putting their accident rate at 2.1 incidents per million miles versus 4.85 for humans (Waymo). Across millions of rides, that gap translates into real lives saved and injuries prevented. A utilitarian would conclude that continued deployment is justified even when specific incidents cause localized harm, because the pleasure produced across the full population of riders outweighs the pain of the incidents. The floodwater recall and the couple trapped on a freeway are real harms, but in the utilitarian math they are weighed against an enormous aggregate benefit, and the math favors the technology.

Deontology evaluates actions differently. Rather than asking what consequences an action produces, it asks whether the action itself respects the dignity of the people it involves. Kant's Formula of End in Itself holds that we must act so that we treat others always as an end and never merely as a means. Applied to Waymo, both incidents raise serious concerns under this principle. Elliot Slade and his fiancée did not meaningfully consent to riding a vehicle with no passenger override mechanism on a freeway; they were placed in a live test of an unresolved edge case, with no way out, used as a means to Waymo's commercial and engineering ends. The same analysis holds for San Antonio residents whose roads became an unacknowledged testing environment for flood scenario handling. Neither group was treated as persons with inherent worth whose informed participation in that risk mattered; both were treated as background conditions for Waymo's operations. Deontology gives those violations a clear and unconditional name regardless of how the broader numbers appear.

III. The Strongest Ethical Framework

Deontology handles this case better than utilitarianism, and the reason is what each framework is actually measuring. Utilitarianism's failure here is structural, not incidental. Because it evaluates actions entirely through their consequences, it can always absorb individual violations into the aggregate if the overall numbers hold up. Waymo's 170 million miles become a defense. Statistically, the technology is saving lives, so the suffering of specific passengers and specific communities gets outmatched. Hedonism makes this worse, not better: since what matters is the balance of happiness and pain across all affected people, the happiness produced by thousands of safe, convenient rides can simply outweigh the pain of the couple trapped on the freeway. A framework that can rationalize almost any individual harm as long as the total utility is positive is not a reliable check on how companies treat the people they test on.

Deontology does not have that weakness. The Formula of End in Itself does not ask whether the aggregate came out positive; it asks whether the people involved were treated as ends in themselves or merely as means. That question has a direct and unconditional answer in this case. Slade and his fiancée were means: they were inside a vehicle they could not control,

absorbing a risk Waymo had not resolved, with no recourse beyond a remote operator asking if they wanted to keep going. The $40 credit Waymo offered afterward is the clearest possible illustration of what happens when a company treats a dignity violation as a consequence to be compensated rather than a wrong to be answered for. Under the categorical imperative, the principle behind Waymo's approach cannot be let to stand. If every technology company operated by the logic of deploy now, compensate later, informed consent to the risks of new technology would become meaningless. Communities would simply be testing grounds for products they had no role in shaping. Deontology names that as wrong in a way that utilitarianism, working from consequences and aggregates, cannot.

IV. Technology Assessment Paradigms

The Office of Technology Assessment model relies on expert-driven risk analysis. A team of specialists examines a technology, maps its failure modes, and produces concrete policy recommendations grounded in technical evidence. Applied to Waymo, an OTA process could have generated a rigorous assessment of AV vulnerabilities to adverse weather conditions before deployment expanded to cities like San Antonio, and produced enforceable testing standards to address them. OTA's strength is its depth and credibility; because it draws on expert knowledge, its findings carry significant weight with regulators and can translate directly into safety requirements.

Public Technology Assessment opens that process up. Instead of keeping the conversation among specialists, it brings in residents and riders; the people who actually live with the technology day to day. For Waymo, that could have meant public hearings in San Antonio and San Francisco where community members flagged local hazards and raised questions about what passengers can do when something goes wrong. Policy that comes out of a PTA process carries democratic backing, not just technical authority, and it surfaces the kind of local knowledge that no expert panel is guaranteed to have.

Constructive Technology Assessment takes that logic further by moving community participation out of the regulation phase and into the design phase itself. Rather than deliberating over a finished product, CTA treats the people most affected by a technology as contributors to how it gets built. The Design Justice Network's principles capture this directly: design processes should center those directly impacted by their outcomes, and community impact should take precedence over a designer's intentions. Applied to Waymo, CTA would have required San Antonio residents and actual riders to be part of the development process, treating their knowledge of local conditions and their experience as passengers as inputs the system needed to incorporate before deployment.

V. The Best Technology Assessment Paradigm

CTA is the right paradigm for this case, and the argument for it is not just that it is more inclusive than OTA or PTA; it is that it is the only one that operates at the level where the failures actually originated. The Waymo incidents were not only policy failures; they were design-process failures. An OTA report or PTA forum might have identified flash flooding eventually, but CTA is more likely to catch that kind of issue earlier because it brings local knowledge into design before deployment. Flooding is not an obscure risk that required expert analysis to surface; it is something San Antonio residents deal with regularly and could have flagged immediately if anyone had asked them. The fact that a vehicle ended up in a creek is evidence that nobody did. OTA assumes experts can anticipate community-specific risks from the outside; this case proves they cannot always. PTA assumes that post-design public deliberation can fix mistakes; but by the time a PTA forum is convened, the software is already written.

CTA is different. It does not add community input as a layer on top of the existing development process; it makes community input a part of the development process itself. The remote operator asking whether Slade wanted to continue the trip after the freeway incident is a perfect illustration of what accountability looks like when it is designed as an afterthought; under CTA, those accountability structures would have been designed in from the start, because

the people most affected by their absence would have had a seat at the table when the system was being built.

VI. The Ideal Solution

An ideal solution would not rely on Waymo choosing to do better; it would create specific regulatory requirements that force accountability into the process before deployment and after incidents. Three concrete measures would accomplish this. The first is a mandatory pre-deployment community review process. Before an AV company can operate in a new city or expand to a new use like freeway deployment, it would have to file a public impact assessment mapping the affected communities, identifying local hazards, and documenting how community input shaped the plan. That assessment would be open for public comment, and operating approval would require showing that concerns were actually addressed; filing a document and ignoring it would not be enough. The second is a strict liability standard for AV-related accidents and injuries. Right now, someone injured by a Waymo vehicle has to prove the company was negligent, which is nearly impossible when the company owns all the data about what its software decided and why. Strict liability cuts through that problem: if an AV causes an injury, the company is responsible regardless of whether the injured person can prove negligence. The third is a federal requirement that all AV passenger cabins include a physical override mechanism allowing riders to bring the vehicle to a safe stop. Making it a federal safety standard, the way seatbelts and airbags are, treats passenger control not as a product feature but as a requirement. Together these three measures give CTA regulatory power: community input becomes a legal condition of deployment, companies bear the actual cost of design failures, and passengers have a guaranteed floor of safety regardless of what the company decided to prioritize.

VII. Roadblocks and How to Overcome Them

The problem is straightforward: none of what the previous section proposes is currently required of anyone. NHTSA's oversight of autonomous vehicles is built almost entirely on voluntary participation; its main program for AV oversight, AV STEP, allows companies to opt in

rather than compelling them to comply, and its other mechanisms follow the same logic (Grigorian and Englund). A company that chooses not to engage faces no consequences for that choice. Waymo can deploy faster, fix problems after they surface, and absorb the cost of recalls more cheaply than it could build genuine community participation into its process. The incentive structure actively works against the solution.

Overcoming that requires changing what NHTSA can demand. The measures proposed above only have power if they are backed by a federal mandate that ties deployment approval to compliance. That is a political challenge as much as a regulatory one; the current NHTSA framework is explicitly oriented toward removing barriers to deployment, not adding them (Grigorian and Englund). Waymo logged 170 million miles, published strong safety statistics, and still had to recall thousands of vehicles after one entered floodwater and two passengers were left with no way out of a runaway car. Volume and speed are not the same thing as accountability. The only way to close that gap is to make community engagement and passenger protection conditions of being on the road, not commitments a company can choose to skip.

Works Cited

Del Rey, Michelle. “Thousands of Waymo Autonomous Driving Cars Recalled. Here’s Why.” *Usatoday.com*, USA TODAY, 12 May 2026, www.usatoday.com/story/cars/recalls/2026/05/12/waymo-car-recall/90042845007/.

Design Justice Network. “Design Justice Network Principles.” *Design Justice Network*, 2018, designjustice.org/read-the-principles.

Grigorian, Christopher, and Nicholas Englund. “Driving into 2026: The State of NHTSA and the Future of Vehicle Safety Regulation | Foley & Lardner LLP.” *Foley & Lardner LLP*, 20 Nov. 2025, www.foley.com/insights/publications/2025/11/driving-into-2026-the-state-of-nhtsa-and-the-future-of-vehicle-safety-regulation/.

Tangermann, Victor. “Waymo Pulled Its Cars from the Freeway after One Fled Police with Horrified Couple on Board.” *Futurism*, 30 May 2026, futurism.com/advanced-transport/waymo-pulled-cars-freeway-fled-police. Accessed 31 May 2026.

Waymo. “Safety Impact.” *Waymo*, 2024, waymo.com/safety/impact/.