

PANEL 4
New Ventures
in Tort Theory





HOLISTIC TORT THEORY

PROFESSOR ALAN CALNAN

FRONTIERS OF TORTS CONFERENCE; FEBRUARY 7, 2020

SUMMARY

- Traditional Tort theories are siloed and restricted
- Legal institutions (like Torts) form and operate holistically
- Many Tort features reflect the law's holism
- Holism exposes anomalies in other Tort features
- Holistic tort theory fills voids in Tort jurisprudence

TRADITIONAL TORT THEORIES & ALTERNATIVES

Agency Theories

Humanities

Analytical &
Historical

Law &
Economics

Civil Recourse

Social Theories

Social Sciences

Empirical &
Qualitative

Compensation
& Deterrence

Social Justice &
Pluralism

Natural Law Theories

Humanities

Analytical &
Conceptual

Corrective
Justice

Naturalized Theories

Natural
Sciences

Empirical &
Quantitative

Reciprocity

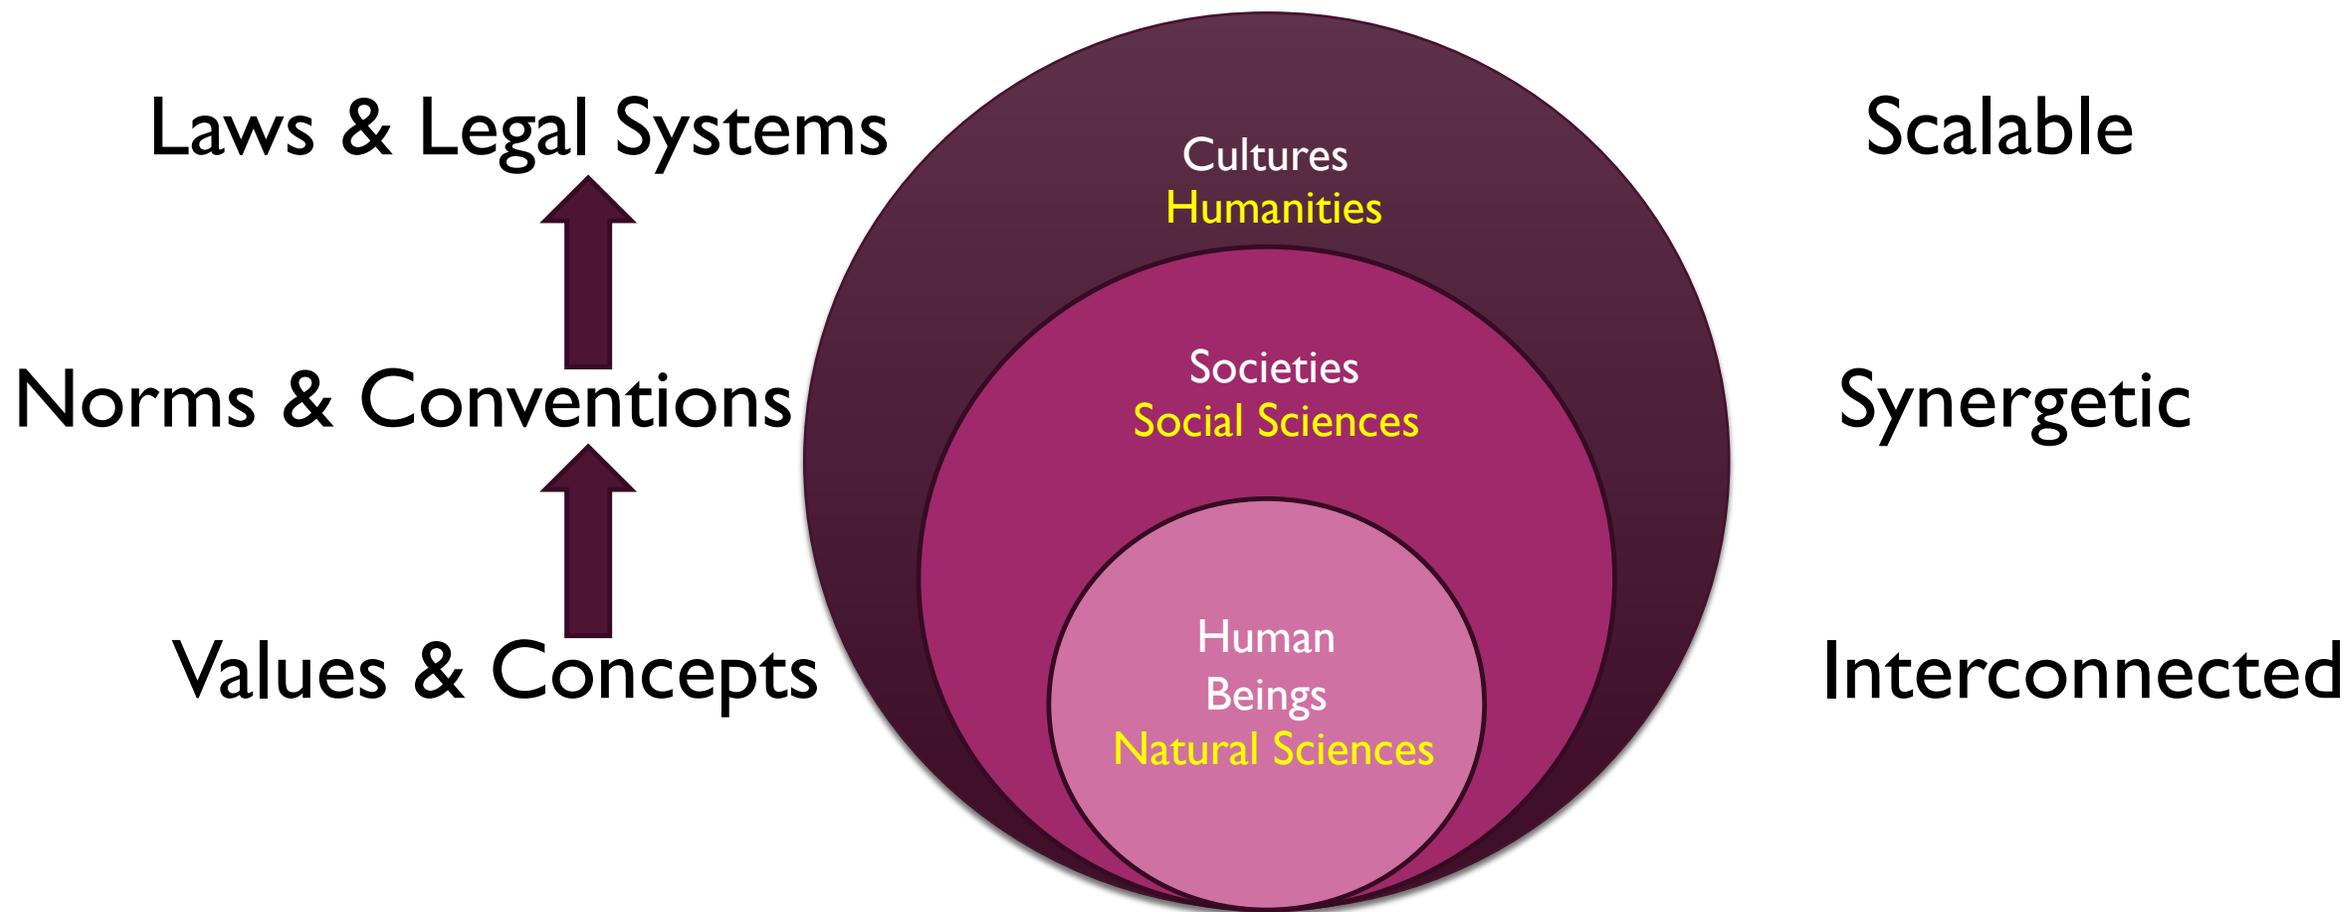
Holistic Theory

All Knowledge

Analytical &
Scientific

Torts as
Systems

LAW'S HOLISM



TORTS' HOLISTIC FOUNDATIONS

Tort System



- Cooperation & Conflict Reconciliation
- Sociocultural Homeostasis
- Universality

System Structure



- Dispute Resolution + Lawmaking + Social Oversight
- Complex Adaptive Systems
- Coordination Dynamics

Tort Concepts



- Rules, Standards, & Principles
- Instinctive & Conventional Wrongs
- Multifactor Balancing Analyses

TORT ANOMALIES EXPOSED BY HOLISM

Reasonableness

Concept: Rational
NOT Emotional

Reality: Emotion
and/before Reason

Wrongs

Concept: Wrongs
are Mental

Reality: Wrongs
ALSO are Personal

Atomism

Concept: Judgment
is Atomistic

Reality: Judgment is
Holistic

CONCLUSIONS

Methodology

- A holistic approach **broadens** the knowledge domains informing our Tort theories

Substance

- A holistic theory **deepens** our understanding of Torts' values, structure, content, and processes

Talk

The Coming Collision

Autonomous Vehicles and Individual Fault Liability

1. *The Allure of Autonomous Vehicles?* The impending emergence of autonomous vehicles will be a shock of seismic proportions for the American automobile liability system. At the moment, it also seems to be a sure thing—a matter of when, not if. Some of this may be marketing hype, but it isn't all hype.
 - 1.1. *Perfected driverless cars* have an enormous advantage over human drivers—no lapses, no bad moods, no anti-social behavior.
 - 1.2. *Fully autonomous vehicles* aren't erratic in the way that human beings are; they don't get distracted; they aren't afflicted by road rage or other moods; they don't drink and drive; they don't forget to take their medications or take the wrong ones before getting behind the wheel.
 - 1.3. *Common sense* tells us that these are huge advantages; the data on existing auto accidents is even more overwhelming.
2. *Eliminating Lapses.* According to the National Highway Traffic Safety Administration (NHTSA), approximately 2.4 million people are injured in auto accidents each year in the United States, and almost 40,000 are killed. Driver error causes an overwhelming majority of these cases. The NHTSA pegs the proportion at 94%. Not everyone agrees with that number but all the statistical estimates I've seen classify driver error as *the* leading cause of auto accidents.
 - 2.1. But eliminating driver error—and *only* eliminating driver error—is *future promise* not present performance.
 - 2.2. *Present performance* is quite a different matter.
3. *Are Machine Error Free?* So far machines (autonomous vehicles) are not error free. Indeed, the “transition” may be exceedingly bumpy. So far, autonomous vehicles seem well designed to follow the rules of the road, but to flounder in idiosyncratic situations. They follow algorithms almost flawlessly but struggle with interpreting context.

3.1. *A fender bender.* In one accident, for instance, a truck driver backing up out of loading bay at an airport did not see a self-driving shuttle which had just that very day been put in service circling the airport. The shuttle saw the truck and stopped. But it stayed there immobile—no honking, no backing up; it just sat there. Eventually, the truck driver became aware of the shuttle—by running into it.

3.1.1. This is an instructive accident. On the one hand, the driver was at fault. On the other hand, the autonomous shuttle did not drive defensively. For a human driver avoiding this accident would have been child's play. But for the autonomous vehicle the situation was simply beyond the confines of its known universe. It hadn't been programmed to cope with this kind of situation.

3.1.2. Note, too, that human beings avoid accidents like this pretty regularly. And we have no statistics on that. Driver error is something we track; driver avoidance of accidents that might have occurred had the driver been an AV is not.

3.2. *Generalizing the problem.* When you think about it, interpreting other drivers' behavior is a large part of driving. I'm shuttling back and forth between Los Angeles and Boston this year. I'm also a cyclist. There are discernibly different driving cultures in these two cities. Boston drivers are much more likely to make left turns cutting in front of oncoming traffic the moment a light turns to green, and much more likely to make right turns after lights have changed, again trying to beat the traffic with the right of way.

3.2.1. Boston drivers, like most drivers, norm their conduct around the rules, but they also violate the rules in consistent, predictable ways. As a cyclist, I'm on guard for this, and make a habit of trying to catch the eyes of drivers' who I think might be about to do something aggressive in my vicinity. On the one hand, interpreting a driving culture is not something autonomous vehicles seem on track to do. The intelligence involved doesn't seem to lend itself to algorithmic replication. On the other hand, just what will replace "catching their eyes"—if anything—is unclear.

4. *But still no lapses.* Even so, autonomous vehicles will not be lapse-prone. Distracted driving—which now seems to be on the rise thanks to smartphones and other computing technologies—will become a thing of the past.
 - 4.1. If you eliminate human error as a cause of car crashes, you eliminate the leading cause of car crashes. The kind of safety improvement that seems probable is one that is simply unimaginable as long as people operate cars.
 - 4.2. Consequently, autonomous vehicles can be afflicted by significant imperfections and still reduce driving deaths and accidents.
5. *Eliminating Fault—and Individual Responsibility.* Autonomous vehicles are on a collision course with the present American liability system for a simple, basic reason. That system is fault-based, and focused on individual conduct. It is a system of individual responsibility and it is utterly obsolete once cars are not driven by individual human beings.
 - 5.1. *Ordinarily, in law, we wedge new phenomena into pre-existing frameworks.* If someone comes up with a complex “new” kind of real estate arrangement we still determine whether, for say tax purposes, it is a lease or a sale.
 - 5.1.1. In exceptional circumstances, we concede that a new legal construct is necessary. For example, when creative transactional lawyers created “interest rate swaps” the law struggled mightily to squeeze them into existing boxes, but eventually conceded that they were genuinely novel and treated them accordingly.
 - 5.1.2. The construction of novel legal categories and regimes, however, rarely happens easily or quickly.
 - 5.2. *Autonomous vehicles are likely to follow the same path.* Initially, we may treat cars with no drivers as if they had drivers, and judge their conduct accordingly. When pressed to assess the adequacy of some aspect of autonomous vehicle design, we are likely to turn to product liability design defect law.
 - 5.2.1. And we may be able to wedge some of the issues into design defect law, albeit with considerable difficulty.
 - 5.2.2. But, in the long run, the negligence law we apply to cars with drivers really won’t be able to accommodate autonomous vehicles.

- 5.3. *Wedging Driverless cars into negligence liability.* Consider the shuttle accident again. If we apply ordinary negligence law to the “conduct” the shuttle we judge the shuttle to be comparatively negligent, assign a percentage to its culpably responsibility for the collision, and impute that responsibility to the entity that owned the shuttle.
- 5.3.1. This all “works” in the sense that we can perform the operations involved but it’s wrongheaded. The car had no driver. Any “fault” was on the part of the manufacturer of the vehicle. The car didn’t take appropriate defensive action because it wasn’t programmed to do so.
- 5.4. *What about holding the manufacturer liable?* This seems to be on the right track, but it’s not an easy design defect case. It’s anything but obvious what we can expect from autonomous vehicle at this point in time.
- 5.4.1. The engineering is advanced and doesn’t lend itself to evaluation by lay juries, and it is evolving quickly. More on this later.
- 5.4.2. Right now, we tend to assume that different automobile manufacturers will offer different “products”—maybe. It seems a bit strange to suppose that BMW will offer “ultimate self-driving AVs” or that Volvo will offer especially sedate ones.
- 5.4.3. Still, we can imagine a situation where some manufacturers produce cars with distinctive defect, just as airplane manufacturers do now.
- 5.4.4. Down the road, however, AVs may all share a common operating system and be networked into a grid. Individual vehicles may just be nodes where the system operates. They won’t have distinctive designs for product liability law to latch onto.
- 5.4.5. Responsibility will reside in the system as whole. We might still find it useful to hold manufacturers of AVs liable for accidents involving their AVs, but the liability won’t be product defect liability. The manufacturer will be a conduit for the imposition of liability on the system as a whole. Responsibility will reside with the collective enterprise of autonomous vehicle operation, not with any individual or firm, but
- 5.5. *This is a circumstance where you can’t successfully force a new technology into the Procrustean bed of existing legal boxes—not ultimately, anyway.* You can squeeze

autonomous vehicles into the product liability box, but only with considerable difficulty and limited success. Eventually, some other legal regime will be needed.

- 5.6. *Something like this—but less extreme—has happened before.* At the end of the 19th century, workplace accidents precipitated a crisis in the common law of torts by challenging its highly individualistic form of responsibility for avoiding and repairing harm. The result, after a major political struggle, was the displacement of the common law of torts by workers' compensation.
- 5.7. *The emergence of autonomous vehicles* is an even more acute challenge because individual drivers will, eventually, no longer exist and it will become impossible to continue with our individual responsibility-centered form of liability.
- 5.8. *In broad terms*, the alternative is to switch to a system of “collective”, or “enterprise” responsibility.
6. *The Fracture at the Center of Modern American Tort Law.* Pre-modern tort law was a law of nominate, mostly intentional, wrongs. Modern tort law is mostly a law of accidents that are recurring byproducts of basic activities in an industrial and technological society.
 - 6.1. *Modern tort law* was born out of the emergence of accidental physical harm as a pressing *social* problem. As a prominent legal historian put it, modern American tort law is “a body of law created when the industrial revolution and industrial accidents began to wreak havoc on the bodies of workers and passengers.”
 - 6.2. *Oliver Wendell Holmes' famous aphorism* epitomizes the point: “[o]ur law of torts comes from the old days of isolated, ungeneralized wrongs, assaults, slanders, and the like,” whereas “the torts with which our courts are kept busy today are mainly the incidents of certain well known businesses . . . railroads, factories, and the like.”
7. *Between Individual and Collective Responsibility.* The twin preoccupations that gave birth to modern tort law—with accidental wrongs instead of intentional ones and with a social problem instead of with individual wrongs—had, and continue to have, large implications for torts as a legal field.

- 7.1. Because the accidents with which modern tort law is preoccupied are characteristically associated with activities, responsibility for those accidents may be lodged *either* with individuals or with activities.
 - 7.2. Negligence liability instantiates the individual liability alternative. Enterprise liability instantiates the collective alternative.
8. *Clinging to Individual Responsibility.* Driving as it exists in our present society is an ongoing, systemic activity, and huge. Consequently, we are able to understand accidents in broad, systemic, ways. We can say, for example, that more than 90% of auto accidents are caused by driver error. Driver error is a systemic problem. That's why autonomous vehicles are so alluring.
- 8.1. But our liability system treats auto accidents as instances of individual responsibility. Individual fault liability is the prevailing regime. When an accident happens, we try to determine which driver was responsible—whose lapse caused the crash.
 - 8.2. Sometimes, of courses, no one is responsible; or a defect in the car is responsible; or a defect in the design of the roadway; and so on. Mostly, though, we find that someone is responsible and their irresponsibility mostly consists of a lapse in attention or concentration.
9. *Collective Responsibility.* Holmes' famous observation points us towards a form of collective responsibility. The assumption is that enterprises engaged in activities on a recurring basis—*not* individuals undertaking isolated risky acts—are the most important source of accidental harm in the modern world. In a modern technological society most tortious wrongs (and most harms) are the predictable byproducts of ongoing activities. Enterprises, therefore, ought to be treated as the *fundamental units of responsibility* for the purposes of attributing accidental harm.
- 9.1. *Collective Responsibility as the Road Not Taken.* Even in our world of human operated cars we might have instituted collective responsibility by, say, adopting a New Zealand type scheme and relying on a mix of social insurance and safety regulation to govern accidents. Fundamentally, we didn't. We stuck with individual fault liability. But aspects of our system are mixed.
 - 9.2. *Just why we stuck with Fault Liability* is a long story, shaped by idiosyncratic events, no doubt. One important reason, though, has to be that individual responsibility tracks our conviction (which itself is now shaped by the

liability system) that cars are inanimate machines. They don't drive themselves, and responsibility for operating them safely lies with their drivers.

- 9.3. *Our Mixed. System.* Even though the United States has retained individual fault liability as its core system of responsibility for accidents, *it intersects with a products liability regime which is, in part, a regime of enterprise liability.*

ENTERPRISE LIABILITY (A NECESSARY DETOUR)

10. *What is "enterprise liability"?* Enterprise liability is usually explained in short slogans—"activities should bear the costs of those accidents that result from their characteristic risk impositions"; "it is only fair that an industry should pay for the injuries it causes"; "losses should be borne by the doer, the enterprise, rather than distributed on the basis of fault." These slogans are useful ways of epitomizing enterprise liability, but they are too terse to serve as satisfactory explications.
11. *Enterprise Liability as a Tort Regime.* Enterprise liability is constituted as a distinctive regime of tort liability in part because it asserts:
- 11.1. Internalization. Internalization prescribes that the costs of those accidents that are characteristic of an enterprise *should be absorbed by the enterprise* as operating expenses, not left on those whose bad luck it is to get in the enterprise's way;
 - 11.2. Loss-Dispersion. Loss dispersion asserts that the costs of enterprise-related accidents *should not be concentrated* either on the victim who originally suffered the injury, or on the particular agent who inflicted the injury; and
 - 11.3. Fairness. Fairness prescribes that the costs of accidents that are characteristic of an activity should be *distributed among those who benefit from the imposition of the enterprise's risks.*
 - 11.4. In the case of private firms, *the costs of enterprise-related harms should be distributed* among customers, employees, suppliers, and shareholders, rather than concentrated either on the victim or on the particular agent responsible for the harm at issue.

12. *Products Liability as Enterprise Liability.* The principal form of “enterprise liability” to emerge in 20th Century America was products liability. When it first arose, products liability was an almost pure incarnation of enterprise liability ideas. Justice Roger Traynor’s concurrence in Escola v. Coca-Cola was the products liability as enterprise liability in a grain of sand.

12.1. The Escola concurrence asserts that the costs of product related accidents should be internalized by product producers and be distributed across all those shareholders, employees and customers who benefit from manufacturing selling and using the defective product.

13. *Why Enterprise Liability?* The *three basic policies* of enterprise liability—accident avoidance; loss-spreading; and fairness—all support strict liability.

13.1. Accident Avoidance. By placing the responsibility for product safety in the hands of those who manufacture, distribute and market products strict liability secures maximum protection against unsafe products.

13.1.1. Why is this maximum protection? Because the firms that make up the “enterprise” are in a better position to identify and execute risk reducing measures than courts applying negligence liability are. And because making firms bear the costs of all the physical injuries caused by their defective products provides a stronger incentive to make their products safe than negligence liability does.

13.2. Loss-Spreading. Those in the chain of distribution—especially the product manufacturer—are in the best position to spread the costs of product accidents.

13.2.1. Why? Because product users are all exposed to the same product risks. From an insurance perspective, product users form a relatively homogeneous risk pool. The imposition of strict liability is de facto mandatory third-party insurance against product accidents. Because product liability cannot be disclaimed, the insurance that it provides is relatively resistant to adverse selection problems.

13.3. Fairness. Strict liability will spread the costs of product accidents across all those—consumers, producers, distributors—who benefit from the product.

14. *Reform and Backlash.* Traynor’s envisioned regime of near pure “enterprise liability” was only the beginning of modern American product liability law. For such a

young branch of law, product liability has an intense and conflicted history. Its enterprise liability moment gave way to a negligence backlash. The body of law that we now have is a mix of the two.

15. *Where are we now?* The short answer is that the United States has a product liability law that is the product of both an aborted enterprise liability revolution and an incomplete negligence counter-revolution. Consequently, the law of products liability reflects deep structural conflict.

15.1. The most important effect of this counter-revolution, for our purposes, is on design defect law. There are now two tests for design defect prominent in American law—the “consumer expectation” test and the “risk-utility” test.

15.2. The risk utility test is quite negligence-like. And the consumer expectation test is not easily applied to complex product features. For reasons I shall try to suggest, neither is likely to be easy to apply to autonomous vehicle accidents.

POSSIBLE LIABILITY REGIMES

16. *The Starting Point.* In the U.S., then, we inherit a world in which there is a regime of individual fault liability for automobile accidents, and a regime of enterprise liability for product accidents.

16.1. The temptation is to wedge autonomous vehicles into one box or the other.

16.2. If they won't squeeze into the box of individual fault liability they must be wedged into product liability

16.3. —Or we must invent something genuinely new.

17. *Under Existing Regimes 1: Fault Liability for Drivers.* At the beginning we will no doubt muddle through, applying fault liability to erratic driving by Autonomous Vehicles (AVs) and products liability to “failures” of AVs to avoid accidents.

17.1. On the one hand, we can treat AVs as if they were cars with drivers. If they run stop signs, veer out of their lanes, skid onto sidewalks, or otherwise violate the rules of the road, we can hold them liable on a negligence per se basis for not complying with applicable traffic regulations.

17.2. But who, exactly, will we be holding liable? Not the drivers? There are none. The owners? They have control over the car's operation only in the rarest of circumstances. The manufacturers? Some component part manufacturer?

17.2.1. If, for example, autonomous airport shuttles have significant accident rates and the airports that operate them are liable for the accidents where they AVs are "at fault" by human standards, they will have very good reason to avoid autonomous shuttles. Why be responsible for a technology whose risks you are in a poor position to understand, much less control?

17.2.2. Leaving losses on victims—a regime of "no liability"—can't be the best regime. Victims will bear all responsibility for accidents and injuries arising out of the operation of AVs and they will be in a very poor position to control the risks of the enterprise.

17.2.3. Thus, the manufacturer option seems the only really plausible one. Consequently, we will take a step towards enterprise liability even under existing fault liability regimes. We will gravitate towards product liability.

18. *Under Existing Regimes 2: Product Defect Liability.* As I've suggested earlier, when autonomous vehicles "fail" in some way we can apply defect liability. This is the other half of the default framework and it will almost surely be applied in the near future. Its application is likely to be very difficult, however. Recall the shuttle collision at the Las Vegas airport and consider two fatal accidents involving autonomous driving.

19. *Product Defect Liability.* Design defect liability is difficult to begin with.

19.1. Product accidents arise at the intersection of three large activities—(1) the design, manufacture, and marketing of products; (2) the purchase of products by consumers; and (3) the use of products.

19.2. The role of design defect rules is to divide accidents between the activity of the product manufacturer and seller, on the one hand, and the product user, on the other hand.

19.3. This is a challenging task. The activities of product design, manufacture and marketing on the one hand, and product purchase and use on the other, are mutually dependent and mutually aware. Products are

designed with product users in mind and product users are aware of at least some product design choices. Consequently, when design choices are at issue, determining whether a product risk is characteristic of the manufacturer's activity or characteristic of the user's activity is not easy to devise. Just how to do so well is a problem that bedevils products liability law even when we are dealing with designs far less novel and sophisticated than autonomous vehicles.

20. *Let's reconsider the shuttle accident mentioned earlier.*

20.1. *On the one hand*, the accident was the truck driver's fault. *On the other hand*, there was contributory fault on the part of the shuttle. Consequently, the defect analysis and the articulation of the right regime are not straightforward.

20.1.1. *With respect to the defect analysis*: what can we reasonably demand in the way of defensive driving? That seems like a hard question to answer.

20.2. *But this would be an easy case if the driver were human*. The autonomous shuttle couldn't "foresee" an outcome obvious even to a child too young to drive. If this shuttle had been driven by a human being we would have no difficulty concluding that they were contributorily negligent and should bear some responsibility for the accident.

20.2.1. Is it fair to judge an autonomous vehicle by the competencies of a normal human being? It's attractively easy.

20.2.2. The very thing that makes it "unfair"—the problem of "reading context" is, apparently, formidably challenging for designers of AVs also makes it difficult to analyze as a design defect case.

20.2.3. So should we take the legally easy path of judging the shuttle's performance by the standard of a reasonable person or should we take the more difficult route of determining what we can reasonably expect in the way of safe product design?

20.3. *And what about the negligence of the truck driver?* Should we recognize human error as a defense (in the form of contributory negligence)? Perhaps we should, but doing so is a retreat from the ideal of eliminating driver error as a cause of accidents, a transitional retreat at the very least. (Moreover, fitting a defect liability rule together with a defense of contributory negligence is an awkward exercise.)

21. *Fatal Accident 1: The Uber Crash.* An AV with a safety operator in the passenger seat failed to recognize a woman who was (1) walking a bicycle across the street in the dark and, (2) not at an intersection or pedestrian crossing. She may have been under the influence of drugs or alcohol. Unable to classify her correctly, the AV did not stop, and hit and killed her.

21.1. The product failure (the AV failure) was that the AV's object recognition capacities weren't sufficient to identify the person crossing its path.

21.1.1. The operating system cycled between different classifications which caused it to constantly re-calculate its path. Unable to classify the "object crossing the street" correctly the AV just kept moving forward until it struck the pedestrian and killed her.

21.2. The AV was a modified Volvo SUV. One modification had been to disable the vehicle's forward collision warning and avoidance system. Had it been active, the forward collision warning system would have intervened to slow the SUV down and probably would have averted the fatal collision.

21.3. A separate problem was that the "safety driver" in the car was not paying attention at the time—they were watching a movie on their phone, it seems.

21.4. Victim conduct was a problem, too. The victim was not crossing at a crosswalk and may have been under the influence of alcohol or drugs. In negligence terms, she may have been at fault in two ways—she chose the wrong place to cross, and she may have failed to recognize the threat posed by the AV and respond appropriately because she was impaired.

22. *What kind of object recognition can we demand?* It's an axiom of product liability law that a non-defective product isn't a perfect product. So what's a reasonable demand?

22.1. We could hold AVs to human standards. Trouble is, it's anything but obvious that AVs can always meet human standards.

22.2. Product liability law tests design defect by two tests. One is the consumer expectation test. This test asks what a reasonable consumer would expect

in the way of product performance. The other is the risk-utility test. This test asks whether the advantages of the design justify its risks. The application of this test usually requires comparing the challenged design with a feasible alternative design.

22.3. We don't seem to have any obvious basis to form a reasonable expectation of product performance short of perfection, and applying a risk-utility test is likely to be a dauntingly difficult exercise in second-guessing complex engineering at the edge of existing knowledge.

22.4. And, again, how should we compare any product defectiveness with the victim's carelessness in crossing the street as she did—and the safety operator's failure to intervene?

23. *Fatal Accident 2: The Tesla Crash.* In this instance a Tesla driving on a freeway with its autopilot engaged failed to discern the white side of a tractor-trailer's trailer and ran into it, never adjusting at all for the possibility of collision with an object.

23.1. Again there was a failure of object detection, albeit of a somewhat different sort. The problem seems to have been that the car's sensors couldn't "see" the large white side of the tractor trailer's trailer in the bright light that was shining at that moment.

23.1.1. This happened even though Tesla's auto-pilot was, at the time it was sent out into the world, "state-of-the-art." State-of-the-art, but inferior to alert human eyesight. There are a finite number of vision technologies, all with strengths and weaknesses; none is perfect; and they are not perfect in combination, either.

23.2. And, again, the operator did not intervene. He, too, was watching a movie. But suppose he had been alert and attentive. Would he have recognized the failure of the car's autopilot in time to re-take control of the car and avoid the accident?

23.2.1. Apparently, the evidence from aviation, where automated control systems are quite advanced, indicates that it is very difficult for pilots to monitor automated systems for errors, and correct in time. The recent Boeing 737-Max crashes may be cases in point.

23.3. This case is no easier than the Uber case.

24. *Generalizing*. Both cases are hard and it's not obvious why cases should become easier as technology advances.
 - 24.1. Technological advance should reduce the frequency of accidents but it may compound the difficulty of determining whether they should have been avoided by an alternative design.
 - 24.2. "State-of-the-art" design will be a moving target, particularly if the AI systems operating AVs "learn" quickly and significantly.
25. *Muddling Through*. It may help to think, for a moment, *about where we are with safety technology now, and where we will be next*.
 - 25.1. *At the moment there is quite a bit of technology in cars whose role is correct for human error*. We have antilock brakes, forward collision warning and avoidance systems, automatic braking systems, electronic stability control, blind-spot monitoring, lane-departure warning systems, adaptive cruise control, among others.
 - 25.2. *These technologies correct for our mistakes*. Transitional autonomous vehicles require us—human beings—to correct for their mistakes.
26. *Transitioning to Full Autonomy*. The technologies of mid-transition—partially autonomous vehicles—invert the relation between driver and technology. Technology does the primary driving and humans intervene when it goes awry.
 - 26.1. Think again about the Tesla and Arizona accidents. Humans failed to intervene. Seat-of-the-pants empiricism suggests that monitoring complex, automated operating systems for failure; disengaging the system; taking over control; and then averting disaster; is a very demanding task, and an unattractive responsibility for a human driver to assume.
 - 26.1.1. Experience in the aviation context supports seat-of-the-pants empiricism on this point. Indeed, one knowledgeable writer reports that "it is a strongly-held belief in the aviation community that automation, as distinct from technology more generally, has created new risks of automation dependency that may have offset the reduction in risk attributable to automation." The recent Boeing 737 crashes are sobering cases in point.

- 26.1.2. And monitoring AVs for mistakes raises another question: Should cars monitor themselves and tell us when to “wake up and have a look” at what’s going on?
- 26.2. When humans fail to intervene correctly, how should we apportion responsibility between human and design failure?
27. *Critical Mass.* At some point where there is a critical mass of autonomous vehicles on the road it will make no sense to apply fault liability to them. They don’t have individual drivers. We’re doing “enterprise liability” anyway, and the question will be how to do it.
- 27.1. *Just what counts as “critical mass”* is a question that has only a rough answer. One recent analysis suggests that when 25% of the cars on the road are genuinely autonomous vehicles, we will have a critical mass.
- 27.2. *And that raises the question* of what counts as a genuinely autonomous vehicle. The Society of Automotive Engineers International (SAE) distinguishes five levels.
- 27.2.1. **Level 0.** Human driver does everything.
- 27.2.2. **Level 1.** Automated system *sometimes assists* the driver with *some tasks*.
- 27.2.3. **Level 2.** Automated system can *perform* some tasks subject to monitoring by the human driver and to the human driver performing the other tasks.
- 27.2.4. **Level 3.** Automated system can *both perform and monitor the driving environment* but the human driver must take back control under some circumstances—per the instructions of the automated system.
- 27.2.5. **Level 4.** Automated system can do everything but only under *some* conditions.
- 27.2.6. **Level 5.** Automated system can do everything under *all* conditions.
28. *One plausible trigger* for moving away from our existing regimes is a 25% Level 4 & 5 vehicle threshold.

29. *And what do we move towards?* The most plausible option, I think, is an enterprise liability regime of manufacturer liability for all accidents modeled on Workers' Compensation.

29.1. *While this seems attractive* in part because *it avoids the formidable difficulties of applying design defect law* it would face its own difficult decisions.

29.2. The basic liability rule would impute all accidents that “arise out of” the operation of a motor vehicle to the manufacturer. Assume that we adopt this rule, how then should we treat:

29.2.1. Third-party hacking? (Strict liability for manufacturer?)

29.2.2. Poor road design?

29.2.3. Miscommunication among AVs?

29.3. Terrible weather? Suppose the car “wants” not to drive, but the driver insists?

29.4. And how should we treat standing still—the shuttle crash?

Kenneth W. Simons

Chancellor's Professor of Law and Philosophy
UC Irvine School of Law

Is current tort doctrine adequate to address the emerging opportunities and problems raised by new technology and artificial intelligence?

1. Some emerging opportunities and problems:

AI and machine learning advances are occurring in such areas as medical diagnosis, the internet of things, robots, drones, and autonomous vehicles. These advances may improve safety overall but also create distinctive risks of harm.

The criminal justice system, businesses and private individuals will increasingly use algorithms to assess dangerousness of suspected wrongdoers. The algorithms affect such decisions as whether to take a precaution, or whether to use force, or about the duration of incapacitation in a prison or mental institution.

Tort law itself can be “personalized” based on big data.

2. Recent tort literature:

Scholars offer a wide range of views on the ability of current doctrine to address these issues. Some believe that radical change to tort doctrine is required (such as broader strict liability or, at the other extreme, exemption from any liability). Others believe that tort itself is inadequate to the problems and should be replaced by a different legal regime. But others assert that no significant doctrinal change is needed.

3. My view:

In the short run, current tort doctrine is adequate to address most of the problems. But in the longer run, when autonomous entities and machine learning algorithms play a more dominant role, they will create challenging problems of opacity, justiciability, and bias. Radical revision or replacement of common law tort doctrine might indeed be required, at least in some areas.