

2020

How Reporters Can Evaluate Automated Driving Announcements

Bryant Walker Smith

Follow this and additional works at: https://scholarcommons.sc.edu/law_facpub

 Part of the [Law Commons](#)

HOW REPORTERS CAN EVALUATE AUTOMATED DRIVING ANNOUNCEMENTS

BRYANT WALKER SMITH

Cite as: Bryant Walker Smith, *How Reporters Can Evaluate Automated Driving Announcements*, 2020 J. L. & MOB. 1.

This manuscript may be accessed online at <https://futurist.law.umich.edu> and at <https://repository.law.umich.edu/jlm/>.

This article identifies a series of specific questions that reporters can ask about claims made by developers of automated motor vehicles (“AVs”).¹ Its immediate intent is to facilitate more critical, credible, and ultimately constructive reporting on progress toward automated driving. In turn, reporting of this kind advances three additional goals. First, it encourages AV developers to qualify and support their public claims. Second, it appropriately manages public expectations about these vehicles. Third, it fosters more technical accuracy and technological circumspection in legal and policy scholarship.

This third purpose goes to the core of this interdisciplinary journal. Legal and policy scholarship about emerging technologies often relies at least in part on popular reporting. On one hand, this reporting can provide timely and accessible insights into these technologies, particularly when the scientific literature cannot. On the other hand, this reporting can reflect misconceptions based on incomplete information supplied by self-interested developers—misconceptions that are then entrenched through legal citation. For example, I have pushed back against claims that automated driving will be a panacea,²

1. These questions first appeared in *Questions to Ask About AV Announcements*, LAW OF THE NEWLY POSSIBLE, www.newlypossible.org/wiki/index.php?title=Questions_to_Ask_About_AV_Announcements (last updated Oct. 14, 2019). This article updates and explains them.

2. See Bryant Walker Smith, *How Governments Can Promote Automated Driving*, 47 N.M. L. REV. 99 (2017); Bryant Walker Smith, *Managing Autonomous Transportation Demand*, 52 SANTA CLARA L. REV. 1401 (2012).

that its technical challenges have long been “solved,”³ and that nontechnical issues involving regulation, liability, popularity, and philosophy are therefore the paramount obstacles to deployment.⁴

Common to many of these misconceptions is the question of whether automated driving is finally here. AVs were 20 years away from the late 1930s until the early 2010s and have been about five years away ever since. This is clearly a long history of misplaced optimism, but more recent predictions, while still moving targets, are now proximate enough to realistically drive decisions about investment, planning, and production. Indeed, of the companies that claim to be even closer, some really are—at least to automated driving of some kind.

The “what” of these predictions matters as much as the “when,” and the leading definitions document for automated driving—SAE J3016—is helpful for understanding this what.⁵ The figure below offers a gloss on these definitions, including the widely (mis)referenced levels of driving automation. No developer has credibly promised level 5 (full automation) anytime soon. But many are working toward various applications of level 4 (high automation), which could, depending on their implementation, include everything from low-speed shuttles and delivery robots to traffic jam automation features and automated long-haul trucks. When anything approaching level 5 does become a reality, it will likely be an afterthought in a world that has already been revolutionized in a hundred other ways.

3. See Bryant Walker Smith, *Automated Driving and Product Liability*, 2017 MICH. ST. L. REV. 1, (2017); Bryant Walker Smith, *A Legal Perspective on Three Misconceptions in Vehicle Automation*, in LECTURE NOTES IN MOBILITY: ROAD VEHICLE AUTOMATION 85 (Gereon Meyer & Sven Beiker eds., 2014).

4. See Bryant Walker Smith, *Automated Vehicles Are Probably Legal in the United States*, 1 TEX. A&M L. REV. 411 (2014); Bryant Walker Smith, *supra* note 3 (discussing product liability); Bryant Walker Smith, *The Trolley and the Pinto: Cost-Benefit Analysis in Automated Driving and Other Cyber-Physical Systems*, 4 TEX. A&M L. REV. 197 (2017).

5. SAE INT'L, J3016, TAXONOMY AND DEFINITIONS FOR TERMS RELATED TO DRIVING AUTOMATION SYSTEMS FOR ON-ROAD MOTOR VEHICLES (last updated June 15, 2018), https://www.sae.org/standards/content/j3016_201806 [hereinafter SAE J3016]. The term “automated vehicle” deviates slightly from SAE J3016 but is nonetheless widely accepted. See, e.g., U.N. ECON. COMM’N FOR EUR., RESOLUTION ON THE DEPLOYMENT OF HIGHLY AND FULLY AUTOMATED VEHICLES IN ROAD TRAFFIC (Oct. 2019), unece.org/trans/resources/publications/transwp1publications/2019/resolution-on-the-deployment-of-highly-and-fully-automated-vehicles-in-road-traffic/doc.html; U.S. DEP’T OF TRANSP., USDOT AUTOMATED VEHICLES ACTIVITIES (last updated Feb. 7, 2020), <https://www.transportation.gov/AV>; Final Act, With Comments: Uniform Automated Operation of Vehicles Act (2019), <https://www.uniformlaws.org/HigherLogic/System/DownloadDocumentFile.ashx?DocumentFileKey=a78d1ab0-fac8-9ea1-d8f2-a77612050e6e&forceDialog=0>. However, the levels of automation generally describe features on vehicles rather than the vehicles themselves. See SAE J3016.

Figure: A Gloss on SAE J3016⁶

Your role in driving automation

Driving involves paying attention to the vehicle, the road, and the environment so that you can steer, brake, and accelerate as needed. If you're expected to pay attention, you're still driving -- even when a vehicle feature is assisting you with steering, braking, and/or accelerating. (Driving may have an even broader legal meaning.)

Types of trips

- A. You must drive for the entire trip
- B. You will need to drive if prompted in order to maintain safety
- C. You will need to drive if prompted in order to reach your destination
- D. You will not need to drive for any reason, but you may drive if you want
- E. You will not need to drive for any reason, and you may not drive

Types of vehicles

- A. Vehicles you can drive
- B. Vehicles you can't drive

Types of vehicle features

These are the levels of driving automation. They describe features in vehicles rather than the vehicles themselves. This is because a vehicle's feature or features may not always be engaged or even available.

The operational design domain ("ODD") describes when and where a feature is specifically designed to function. For example, one feature may be designed for freeway traffic jams, while another may be designed for a particular neighborhood in good weather.

By describing a feature's level of automation and operational design domain, the feature's developer makes a promise to the public about that feature's capabilities.

6. This first appeared at *Automated Driving Definitions*, LAW OF THE NEWLY POSSIBLE, http://newlypossible.org/wiki/index.php?title=Automated_Driving_Definitions (last updated Aug. 1, 2018).

Assisted driving features

L0: You're driving

L1: You're driving, but you're assisted with either steering or speed

L2: You're driving, but you're assisted with both steering and speed

Automated driving features

L3: You're not driving, but you will need to drive if prompted in order to maintain safety

L4: You're not driving, but either

a) you will need to drive if prompted in order to reach your destination (in a vehicle you can drive) or

b) you will not be able to reach every destination (in a vehicle you can't drive)

L5: You're not driving, and you can reach any destination

As the following questions for reporters make clear, automated driving is much more than just a level of automation. The questions, which fall into five overlapping categories (human monitoring, technical definitions, deployment, safety, and reevaluation), are:

1. Human monitoring

1.1. Is a person monitoring the AV from inside the vehicle? Why? Are they always paying attention? How can they intervene? How often do they intervene? How are they supervised?

1.2. Is a person monitoring the AV from outside the vehicle? Why? Are they always paying attention? How can they intervene? How often do they intervene? How are they supervised?

1.3. Is a person monitoring the AV from a remote center? Why? Are they always paying attention? How can they intervene? How often do they intervene? How are they supervised?

1.4. What are specific examples of difficult scenarios in which a person did not intervene? In which a person unnecessarily intervened? In which a person necessarily intervened? What form did this intervention take?

1.5. At any moment, what is the ratio between the number of people who are monitoring and the number of AVs that are deployed?

2. Technical definitions

2.1. What level of automation corresponds to the design intent for the AV? What level of automation corresponds to how the AV is actually being operated?

2.2. In what environment is the AV operating? On roads open to other motor vehicles? To bicyclists? To pedestrians?

2.3. What infrastructure, if any, has been changed or added to support the AV in this environment?

2.4. If the AV perceives that its path is obstructed, what does it do? For example, does it wait for the obstruction to clear, wait for a person to intervene, or plan and follow a new path?

3. Deployment

3.1. What is the AV's deployment timeline? For how long will it be deployed? Is this a temporary or permanent service?

3.2. Who can buy the AV or its automated driving feature? Under what conditions?

3.3. Who can ride in, receive products or services from, or otherwise use the AV? Under what conditions?

3.4. As part of the deployment, who is paying whom? For what?

3.5. What promises or commitments has the developer of the AV made to governments and other project partners?

3.6. What previous promises, commitments, and announcements has the developer made about their AVs? Have they met them? Do they still stand by them? What has changed, and what have they learned? Why should we believe them now?

4. Safety

4.1. Why do the developer of the AV and any companies or governments involved in its deployment think that the deployment is reasonably safe? Why should we believe them?

4.2. What will the developer of the AV and any companies or governments involved in its deployment do in the event of a crash or other incident?

5. Reevaluation

5.1. Might the answers to any of these questions change during the deployment of the AV? How and why? What will trigger that change?

The remainder of this article explores these questions with a view toward assessing the reality behind a given automated driving announcement or activity. To this end, it is important to understand that a vehicle that requires an attentive safety driver is not truly an automated vehicle. Aspirational, yes. But actual, no. This point underlies many of the questions that follow.

HUMAN MONITORING

Is a person monitoring the AV from inside the vehicle? Why? Are they always paying attention? How can they intervene? How often do they intervene? How are they supervised?

Imagine that as you are boarding a plane, the captain announces that “I’ll be using autopilot today. We’ll be pushing off shortly. Have a nice flight.” How do you feel?

Now imagine that the captain instead announces that “*You’ll* be using autopilot today, because I’m getting off. *You’ll* be pushing off shortly. Have a nice flight.” How do you feel now?

Just as there is a significant difference between these two scenarios, automated driving under the supervision of a safety driver is not the same as automated driving without this supervision. Yet news headlines, ledes, and even entire articles often describe only “driverless” vehicles—even when those vehicles are supervised by at least one trained safety driver who is physically present for every trip.

This confusion has consequences. Casual readers (and even reporters) may believe that an automated driving project is far more technically advanced or economically feasible than it really is. They may therefore be more likely to look for nontechnical explanations for the seemingly slow rollout of automated vehicles. Ironically, they may also discount truly significant news, such as Waymo’s recent decision to remove safety drivers from some of its vehicles.⁷

Reporters should therefore ask whether an automated vehicle is being

7. Dan Chu, *Waymo One: A year of firsts*, WAYMO, (Dec. 5, 2019), <https://blog.waymo.com/2019/12/waymo-one-year-of-firsts.html>.

operated with or without a safety driver inside it, and they should include the answer to this question in the first rather than the final paragraph of their stories. Related questions can then provide further context. Is the safety driver seated in the traditional driver's seat (if there is one) or elsewhere in the vehicle? Can they immediately brake, steer, and accelerate the vehicle? And, in the interest of safety, how are they supervised? As Uber's 2018 fatal crash tragically demonstrated, a system's machine and human elements can both be fallible.⁸

For the most part, an AV developer that uses safety drivers is not yet confident that its vehicles can reliably achieve an acceptable level of safety on their own. This is still true even if a vehicle completes a drive without any actual intervention by that safety driver. At least in the United States, alternative explanations for retaining the safety driver—to comply with ostensible legal requirements, to reassure passengers, or to perform nondriving functions—are generally lacking.

At the same time, AV developers might reach different conclusions about the requisite level of safety or the requisite level of confidence in that safety. To use a very limited analogy: A rock climber's rejection of ropes and harnesses probably says more about the climber's confidence than about their skill.

Is a person monitoring the AV from outside the vehicle? Why? Are they always paying attention? How can they intervene? How often do they intervene? How are they supervised?

A safety driver might be present near rather than inside a vehicle. For example, a demonstration of a small delivery vehicle that is not designed to carry people may nonetheless involve a safety driver seated in a car that trails the delivery vehicle. Reliance on such a safety driver places a significant technical and economic asterisk on claims about the capabilities of these delivery vehicles. Because reliance on safety drivers also involves reliance on a robust communications system, reliance on them also introduces an additional issue of safety.

Tesla's recent introduction of its Smart Summon feature also shows why unoccupied does not necessarily mean driverless.⁹ This feature does not reach the threshold for automated driving—and certainly not “full self-

8. In short: Both the design and the driver were lax on the assumption that the other would not be. Cf. NAT'L TRANSP. SAFETY BD., NTSB - ADOPTED BOARD REPORT HAR-19/03 (Dec. 12, 2019), <https://dms.nts.gov/pubdms/search/document.cfm?docID=479021&docketID=62978&mkey=96894> (describing the factors that contributed to the crash).

9. *Introducing Software Version 10.0*, TESLA BLOG (Sept. 26, 2019), <https://www.tesla.com/blog/introducing-software-version-10-0>.

driving”—because it is designed with the expectation that there will be a human driver who will supervise the vehicle from the outside and intervene to prevent harm. Emphasizing that the user is still a driver may help to temper claims and assumptions that could lead to the dangerous misuse of this driver assistance feature.

Is a person monitoring the AV from a remote center? Why? Are they always paying attention? How can they intervene? How often do they intervene? How are they supervised?

For years, one of the more contentious issues in the automated driving community has involved what might be neutrally termed “remote facilitation of the driving task.” This phrase encompasses a broad spectrum of potential roles performed by actors outside the vehicle—roles that are important to understanding the technical and safety claims made by developers of automotive technologies.

On one side of the spectrum lies remote driving, in which a human driver who may be many miles away from a vehicle uses a communications system to perceive the vehicle’s driving environment and to steer, accelerate, and brake in real time—what SAE J3016 calls “performance of the dynamic driving task.”¹⁰ This remote driving is orthogonal to automated driving (in other words, neither its synonym nor its antonym). Indeed, some automated driving developers skeptical of remote driving are eager to differentiate the two in both language and law.

On the other side of the spectrum lies network monitoring. An automated driving company might maintain a facility in which human agents collectively monitor its AVs, communicate with the users of those vehicles, and coordinate with emergency responders. While stressing that their human agents are not performing the dynamic driving task, some AV developers have been vague about what specifically these agents are and are otherwise not doing.

Journalists, however, can be concrete in their questioning. They can ask whether there is a remote person assigned to or available for each vehicle, what that person does during the vehicle’s normal operation, and what that person does in less common situations. For example, imagine that an AV approaches a crash scene and concludes that it cannot confidently navigate by itself. What role might a remote agent play? Might this person give the vehicle permission to proceed? Might they manually identify roadway objects that the AV could not confidently classify? Might they sketch a rough travel path for the AV to follow if the AV agrees? Might they direct the AV to follow the path even if the AV would otherwise reject it? Or might they

10. SAE J3016, *supra* note 5.

actually relay specific steering, accelerating, and braking commands to the AV?

How a company answers these questions can provide insight into the maturity of its automated driving program. If the company uses physically present safety drivers in its deployments (as most still do), then these questions are largely speculative. But if the company plans to remove these safety drivers, then it should have careful and concrete answers. And if the company declines to share these answers, one might reasonably inquire why.

What are specific examples of difficult scenarios in which a person did not intervene? In which a person unnecessarily intervened? In which a person necessarily intervened? What form did this intervention take?

While anecdotes alone are not enough to establish reasonable safety, they can be helpful in measuring progress. An automated driving developer that has been testing its vehicles will have stories about unusual situations that those vehicles (and their safety drivers) encountered. Many of these developers may be happy to share situations that the automated vehicle handled or could have handled without intervention. But pairing these with situations in which human intervention was necessary provides important context. And a company's willingness to share these more challenging stories demonstrates its trustworthiness.

At any moment, what is the ratio between the number of people who are monitoring and the number of AVs that are deployed?

Economic feasibility offers another metric for automated driving—and one that is intertwined with technical feasibility. Economically, automated driving is both attractive and controversial in large part because, true to its name, it promises to reduce the need for human drivers. Asking whether this is in fact happening—that is, whether the ratio of human monitors to automated vehicles is less than 1.0—is another way to assess the technical progress of an automated driving program.

This may be especially helpful with respect to pilot projects involving specialized vehicles traveling at low speeds in limited areas such as airports, downtowns, and shopping malls. There have been and will likely continue to be numerous announcements about these projects across the country. But so long as these vehicles are deployed with at least one safety driver on board, their economic viability is unclear. After all, their hosts could have achieved (and could still achieve) the same functional benefits by simply deploying conventional fleets.

TECHNICAL DEFINITIONS

*What level of automation corresponds to the design intent for the AV?
What level of automation corresponds to how the AV is actually being
operated?*

Automated driving developers are almost certainly familiar, though not necessarily proficient, with the levels of driving automation defined in SAE J3016. They may even reference these levels in their announcements—correctly or not. Understanding the levels may help to assess the claims.

Most automated driving development is focused on levels 3 and 4. On one side, levels 0, 1, and 2 are in fact driver assistance rather than automated driving, and a credible developer should not suggest otherwise. After all, features at these levels only work unless and until they don't, which is why a human driver is still needed to supervise them. On the other side, level 5 describes a feature that can operate everywhere that humans can drive today. But while this is the hope of many automated driving developers, it remains a distant one.

A confusing quirk in the levels of automation is the difference between what I call an aspirational level and what I call a functional level. The aspirational level describes what an automated driving developer hopes its system can achieve, whereas the functional level describes what the automated driving developer assumes its system can currently achieve. For example, most developers of low-speed automated shuttles envision level 4 automated driving, which would not require a human driver for safe operation. But most of these developers still keep their systems under the supervision of human safety drivers who are expected to pay attention, which corresponds to level 2 rather than level 4. Nonetheless, because SAE J3016 focuses on design intent, developers of these systems correctly characterize them as level 4 (the aspirational level) rather than level 2 (the functional level).¹¹

Similarly, California's Department of Motor Vehicles considers automated vehicles that are merely being tested to be "autonomous" even though their safe operation still requires a human safety driver.¹² Otherwise, rules requiring a safety driver absent specific permission otherwise would

11. SAE J3016, *supra* note 5 (explaining the developer of a feature determines its level of automation).

12. *Cf.*, *Key Autonomous Vehicle Definitions*, STATE OF CALIFORNIA DEPARTMENT OF MOTOR VEHICLES, <https://www.dmv.ca.gov/portal/dmv/detail/vr/autonomous/definitions> (last visited March 9, 2020) (The California DMV defines an "autonomous test vehicle" as "a vehicle that has been equipped with technology that is a combination of both hardware and software that, when engaged, performs the dynamic driving task, but requires a human test driver or a remote operator to continuously supervise the vehicle's performance of the dynamic driving task.").

apply to a null set. Because of this interpretation, companies that are testing or deploying automated driving features in California must comply with these rules, while companies that are testing or deploying mere driver assistance features need not. This is why Uber needed permission to test its automated vehicles in California, but Tesla did not need permission to make its Autopilot or Smart Summon driver assistance features available in that state.¹³ Yet, as these examples suggest, testing an automated driving feature is in many ways technically indistinguishable from using a driver assistance feature.

Asking about the aspirational level of automation invites a company to make a public characterization that has marketing and regulatory implications. And asking about the functional level of automation invites a company to temper its aspirations with the current limitations of its technologies.

References to the levels of automation may be helpful in discussions with companies but are generally not necessary or even helpful when reporting to the public. Instead, key phrases can more clearly communicate the current state of a given technology. Three of the most important are:

- “A driver assistance feature that still requires a human driver to pay attention to the road” (levels 1 and 2)
- “A vehicle that is designed to drive itself but needs a safety driver until it can reliably do so” (aspirational level 4)
- “A vehicle that drives itself without the need for a safety driver” (functional level 4)

In what environment is the AV operating? On roads open to other motor vehicles? To bicyclists? To pedestrians?

Automated vehicles have been a reality for decades: They are called elevators, escalators, people movers, and automated trains. But whereas these vehicles operate in highly controlled environments, automated motor vehicles are particularly challenging in large part because the driving environments they will face are so challenging.

Below level 5, however, these driving conditions are limited. SAE J3016 terms these driving conditions the operational design domain,¹⁴ and this ODD is essential to defining an AV’s capabilities. For example, some

13. This was understandably frustrating for Uber. See Anthony Levondowski, *Statement on Self-Driving in San Francisco*, (Dec. 17, 2016) (transcript available at Uber Newsroom). But see Bryant Walker Smith, *Uber vs. the Law*, THE CENTER FOR INTERNET AND SOCIETY: BLOG (Dec. 17, 2016), <http://cyberlaw.stanford.edu/blog/2016/12/uber-vs-law>.

14. See SAE J3016, *supra* note 5.

automated driving features may operate only on freeways, and some AVs may be restricted to certain low-speed routes within certain neighborhoods. Indeed, early automation activities are generally characterized by some combination of slow speeds, simple environments, and supervised operations.

Developers should be upfront about these limitations in their announcements—and if they are not, reporters should ask whether and how the AVs mix with other road users, including pedestrians, bicyclists, and conventional drivers. There is a big difference, for example, between deploying in complex mixed traffic and deploying on a dedicated route with no other traffic.

As an aside: State vehicle codes apply to public roads, and they may also apply to private facilities such as parking garages and private roads that are nonetheless open to the public.¹⁵ For this reason, AVs that are deployed only in privately controlled areas may still have to comply with state laws generally applicable to motor vehicles as well as state laws specific to AVs. Similarly, these laws may (or may not) also apply to delivery robots that travel on sidewalks and crosswalks.¹⁶ Developers that suggest otherwise can be asked to explain the basis for their legal conclusion.

What infrastructure, if any, has been changed or added to support the AV in this environment?

Many AV announcements involve specific tests, pilots, or demonstrations that may or may not be easily replicated in another location and scaled to many more locations. An AV that can accept today's roads as they are—inconsistently designed, marked, maintained, and operated—will be much easier to scale than one that requires the addition or standardization of physical infrastructure. Even if they would be beneficial and practical, infrastructure changes are nonetheless important considerations in evaluating scalability. For this reason, automated driving developers should be asked to identify them.

If the AV perceives that its path is obstructed, what does it do? For

15. See, e.g., N.Y. VEH. & TRAF. LAW § 1100(a) (McKinney 2019) (“The provisions of this title apply upon public highways, private roads open to public motor vehicle traffic and any other parking lot, except where a different place is specifically referred to in a given section.”).

16. E.g., N.Y. VEH. & TRAF. LAW § 144 (McKinney 2019) (“Sidewalk. That portion of a street between the curb lines, or the lateral lines of a roadway, and the adjacent property lines, intended for the use of pedestrians.”); *id.* at 159 (McKinney 2019) (“Vehicle. Every device in, upon, or by which any person or property is or may be transported or drawn upon a highway, except devices moved by human power or used exclusively upon stationary rails or tracks.”).

example, does it wait for the obstruction to clear, wait for a person to intervene, or plan and follow a new path?

Even infrastructure that is well maintained will still present surprises, and how an AV is designed to deal with these surprises provides some insight into its sophistication. Many early automated vehicles would simply stop and wait if a pedestrian stepped into their path (or a drop of rain confused their sensors). Even today, many AVs rely on frequent human intervention of some kind. This question accordingly invites a developer to describe the true capabilities of its system.

DEPLOYMENT

What is the AV's deployment timeline? For how long will it be deployed? Is this a temporary or permanent service?

Many recent AV announcements have focused less on technical capabilities and more on actual applications, from shuttling real people to delivering real products. These specific applications often involve partnerships with governments, airports, retailers, shippers, or property managers. But it can be unclear whether these applications are one-time demonstrations, short-term pilots, or long-term deployments. Querying—and, in the case of public authorities, requesting records about—the duration of these projects helps to understand their significance.

Who can buy the AV or its automated driving feature? Under what conditions?

There is an important difference between an automated driving developer that is marketing its actual system and a developer that is merely marketing itself. Yet automated driving announcements tend to conflate actual designs, promises of designs, and mere visions of designs. Automakers previewing new vehicle features, shuttle developers announcing new collaborations, and hardware manufacturers touting new breakthroughs all invite the question, “Can I actually buy this vehicle now?”

Who can ride in, receive products or services from, or otherwise use the AV? Under what conditions?

This same logic applies to announcements about services that purportedly involve automated driving. The launch of an automated pizza delivery service open to everyone in a city is much more significant than the staged delivery of a single pizza by a single AV. So too with the automation of long-haul shipping, low-speed shuttles, and taxis. Services that at least part of the public can actually and regularly use are far more significant than one-off

demonstrations.

As part of the deployment, who is paying whom? For what?

For the reasons already discussed, the economics of early deployments can be hazy. Why are automated shuttles, each with its own safety driver, more cost-effective than conventional shuttles? Why are automated trucks, each with its own safety driver, more cost-effective than conventional trucks? The financial arrangements with project partners—especially public authorities subject to open records laws—can offer some insight into whether these early deployments provide tangible benefits or are instead largely exploratory or promotional.

What promises or commitments has the developer of the AV made to governments and other project partners?

When project partners are involved for long-term rather than near-term benefit, it can be helpful to query their expectations. Imagine, for example, that an airport or retirement community announces its intent to host automated shuttles that are supervised by safety drivers. When has the developer of these shuttles suggested or promised that safety drivers will no longer be necessary? And who bears the cost of paying these drivers in the interim?

What previous promises, commitments, and announcements has the developer made about their AVs? Have they met them? Do they still stand by them? What has changed, and what have they learned? Why should we believe them now?

Because innovation is unpredictable, claims about deployment timelines may turn out to be incorrect even if they are made in good faith. However, the companies (or people) responsible for these claims should acknowledge that they were wrong, explain why, and temper their new claims accordingly. Reporters should demand this context from their subjects and report it to their audience. Of course, a commercial emphasis on speed and controversy can make this especially challenging, in which case the headline “Company X makes another claim” could at least be used for the more egregious offenders.

SAFETY

Why do the developer of the AV and any companies or governments involved in its deployment think that the deployment is reasonably safe? Why should we believe them?

While the broader topic of AV safety is beyond the scope of this article, it

should occupy a prominent place in any automated driving announcement. For years, I have encouraged companies that are developing new technologies to publicly share their safety philosophies—in other words, to explain what they are doing, why they think it is reasonably safe, and why we should believe them. Journalists can pose these same questions and push for concrete answers.

The phrasing of these questions matters. For example, a company might explain that its AV testing is reasonably safe because it uses safety drivers. But it should also go further by explaining *why it believes* that the presence of safety drivers is sufficient for reasonable safety. Conversely, if a company does not use safety drivers, it should explain why it believes that they are not necessary for reasonable safety. And in answering these questions, the company may also have to detail its own view of what reasonable safety means.

In this regard, it is important to recognize that safety is not just a single test. Instead, it includes a wide range of considerations over the entire product lifecycle, including management philosophy, design philosophy, hiring and supervision, standards integration, technological monitoring and updating, communication and disclosure, and even strategies for managing inevitable technological obsolescence. In this way, safety is a marriage rather than just a wedding: a lifelong commitment rather than a one-time event.

What will the developer of the AV and any companies or governments involved in its deployment do in the event of a crash or other incident?

Safety is not absolute. Indeed, just because an AV is involved in a crash does not mean that the vehicle is unsafe. Regardless, an AV developer should have a “break-the-glass” plan to document its preparation for and guide its response to incidents involving its AVs. (So too should governments.) How will it recognize and manage a crash? How will it coordinate with first responders and investigators? A developer that has such a plan—and is willing to discuss the safety-relevant portions of it—signals that it understands that deployment is about more than just the state of the technologies.

REEVALUATION

Might the answers to any of these questions change during the deployment of the AV? How and why? What will trigger that change?

This article ends where it began: Automated driving is complex, dynamic, and difficult to predict. For these reasons, many of an AV developer’s answers to the questions identified here could evolve over the course of a deployment. On one hand, the realities of testing or deployment may demand

a more cautious approach or frustrate the fulfilment of some promises. On the other hand, developers still hope to remove their safety drivers and to expand their operational design domain at some point. How—and on what basis—will they decide when to take these steps? Their answers can help to shift discussions from vague and speculative predictions to meaningful and credible roadmaps.