Perspectives on highway safety: Contemporary issues and the forthcoming age of autonomous vehicles

Fred Mannering

University of South Florida

What has been happening in highway safety?

- Incredible advances in vehicle safety technologies (multiple airbags, autonomous braking, lane departure warning, etc.)
- Drunk driving policies
- Texting laws
- Red light cameras
- Etc.

The result of these:

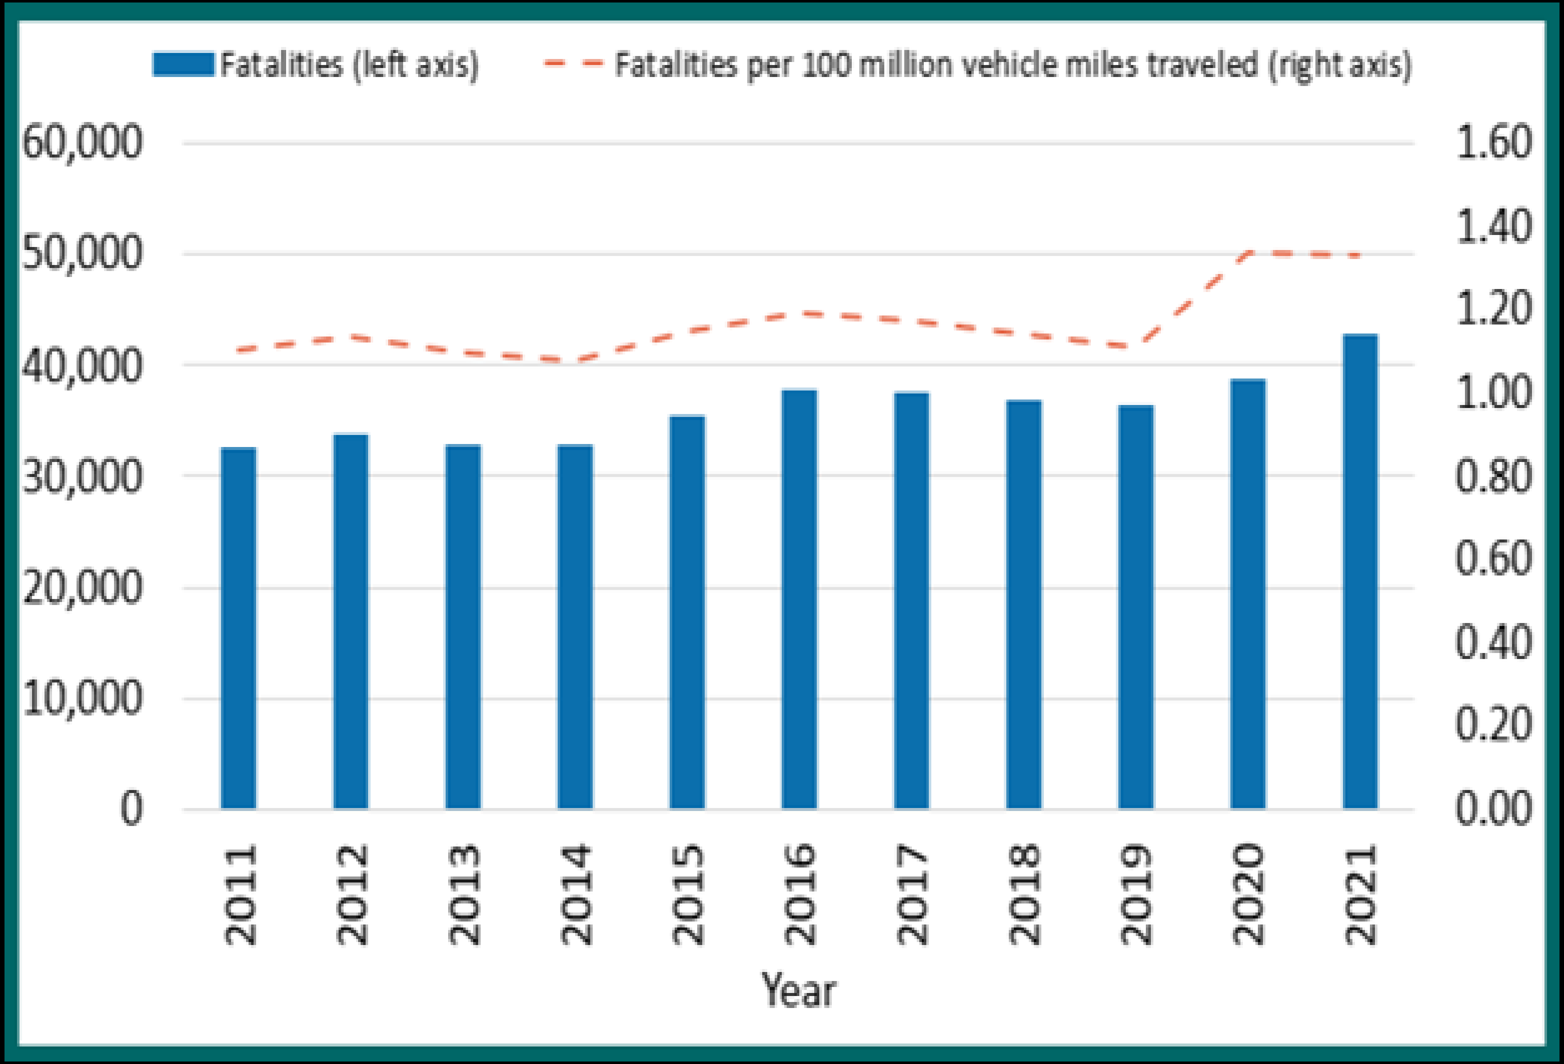

Disturbing highway safety trends:

- Bicyclist fatalities have increased 55% from their 2010 low
- Motorcycle fatalities have reached all-time highs including a 21% increase from 2019 to 2021
- Pedestrian fatalities have increased 80% from their 2009 low

Why have safety efforts not been more successful?

1. Have not accounted for the effects of safety innovations on human behavior
2. Have not considered externalities
3. Have not accounted for the evolution of human behavior
4. Have relied on flawed data analytics

The effects of safety innovations on driver behavior

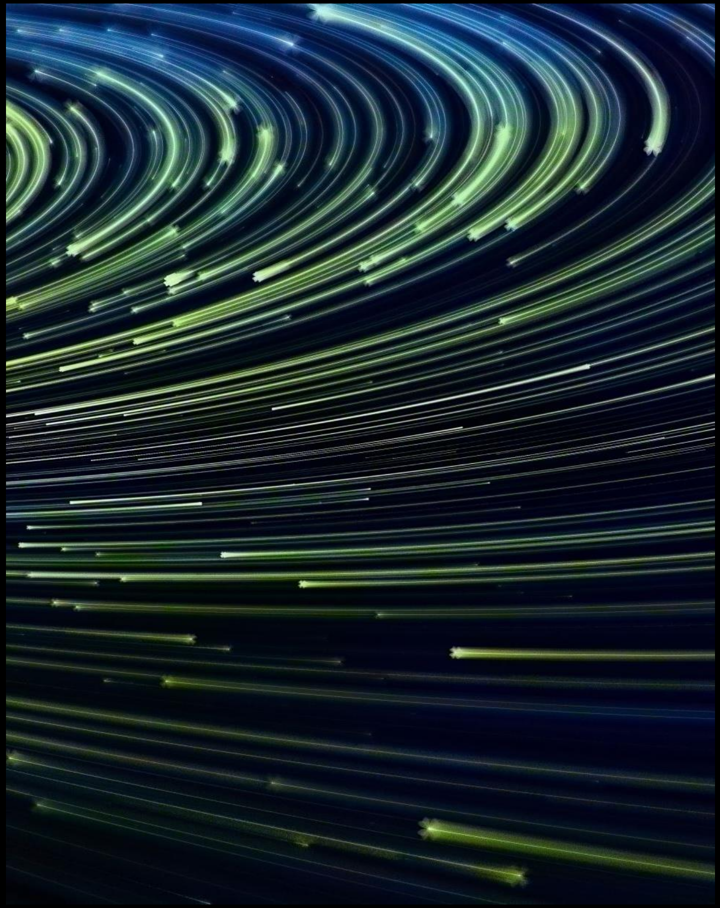

The introduction of safety innovations changes behavior in unexpected ways

Example: Consider a technological innovation that improves safety

What will happen?

Compensating behavior example:

- Uncertainty about probability curves
- Unlikely to extract all benefits in improved safety

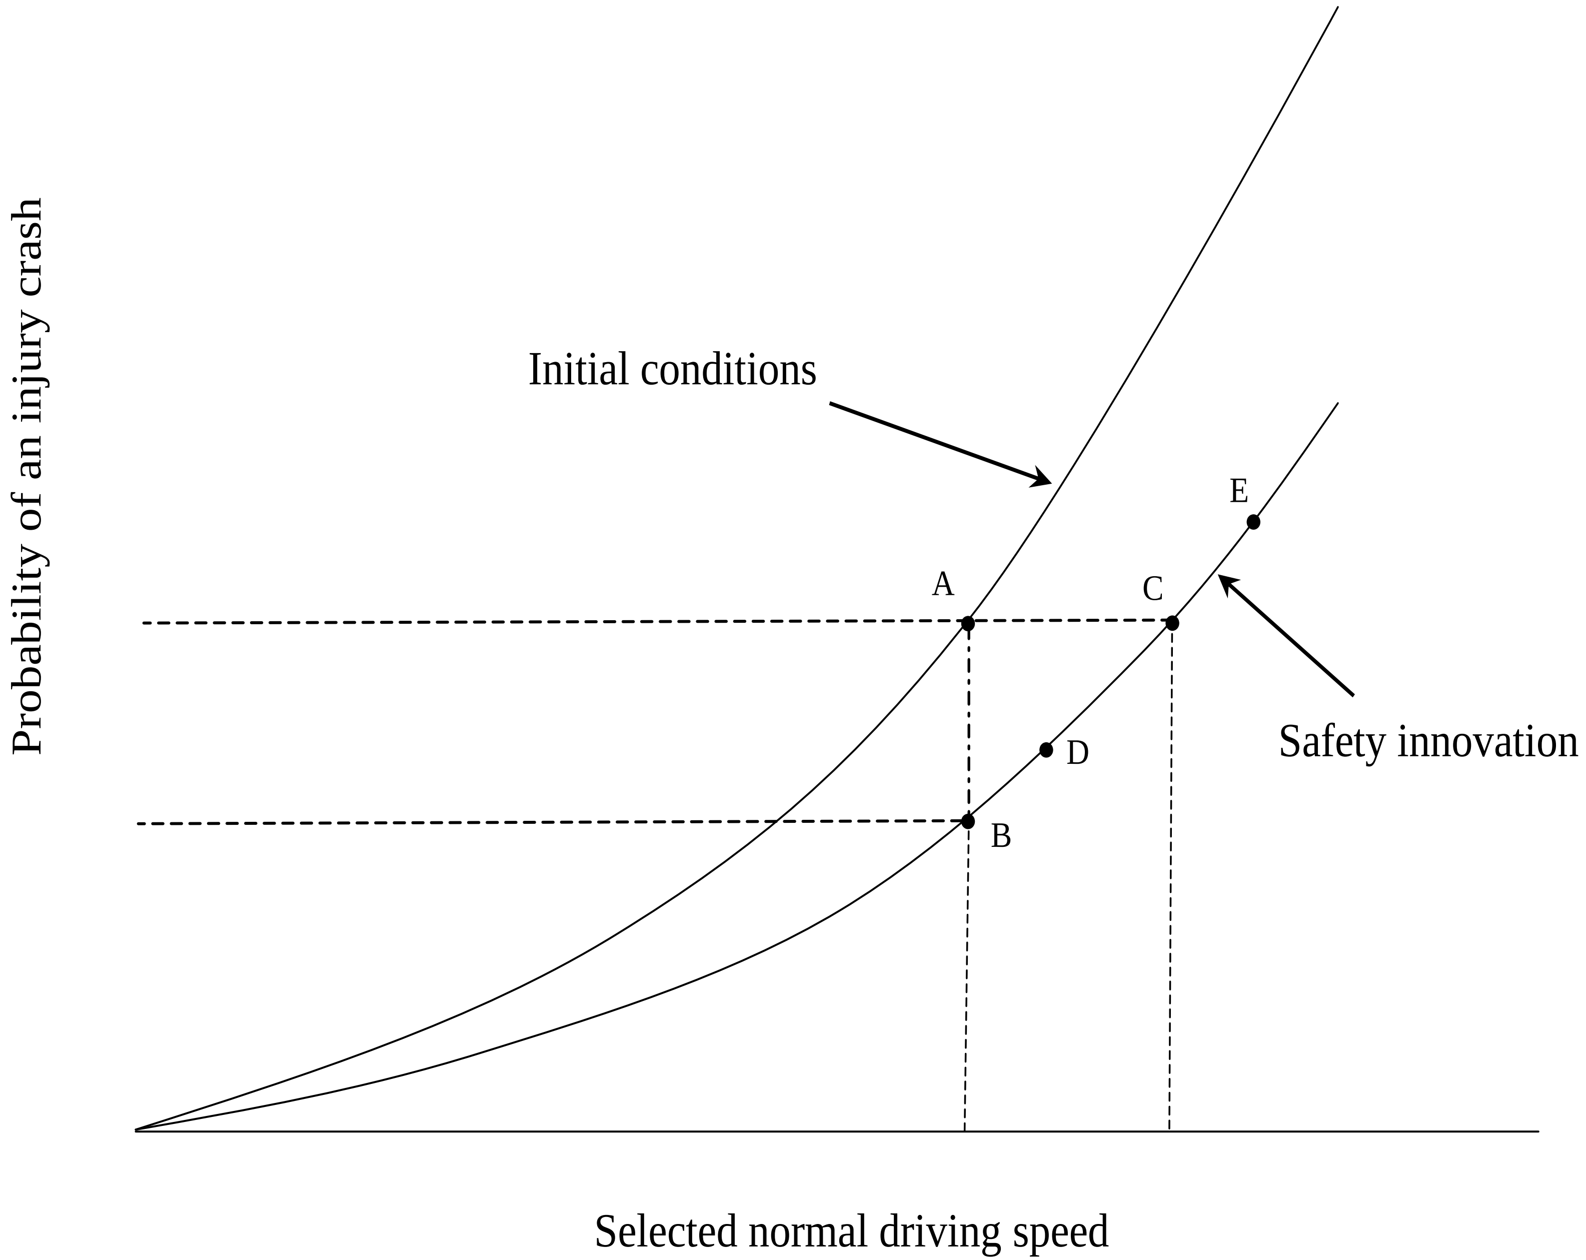

Ignoring compensating behavior:

- Leads to an overestimation of safety innovation effectiveness
- Unexpected consequences that should have been expected

Externalities

Example:
Externalities of compensating behavior

- Safety innovations often result in faster and/or more distracted driving

- Consequence:
  Vulnerable road users and drivers without advance safety features in their vehicles will be exposed to increase risks

Evolution of human behavior

Temporally shifting behavior

- Human behavior is continually evolving
- Behavioral shifts can be induced by technological innovations, changes in risk perceptions, changes in societal norms, etc.

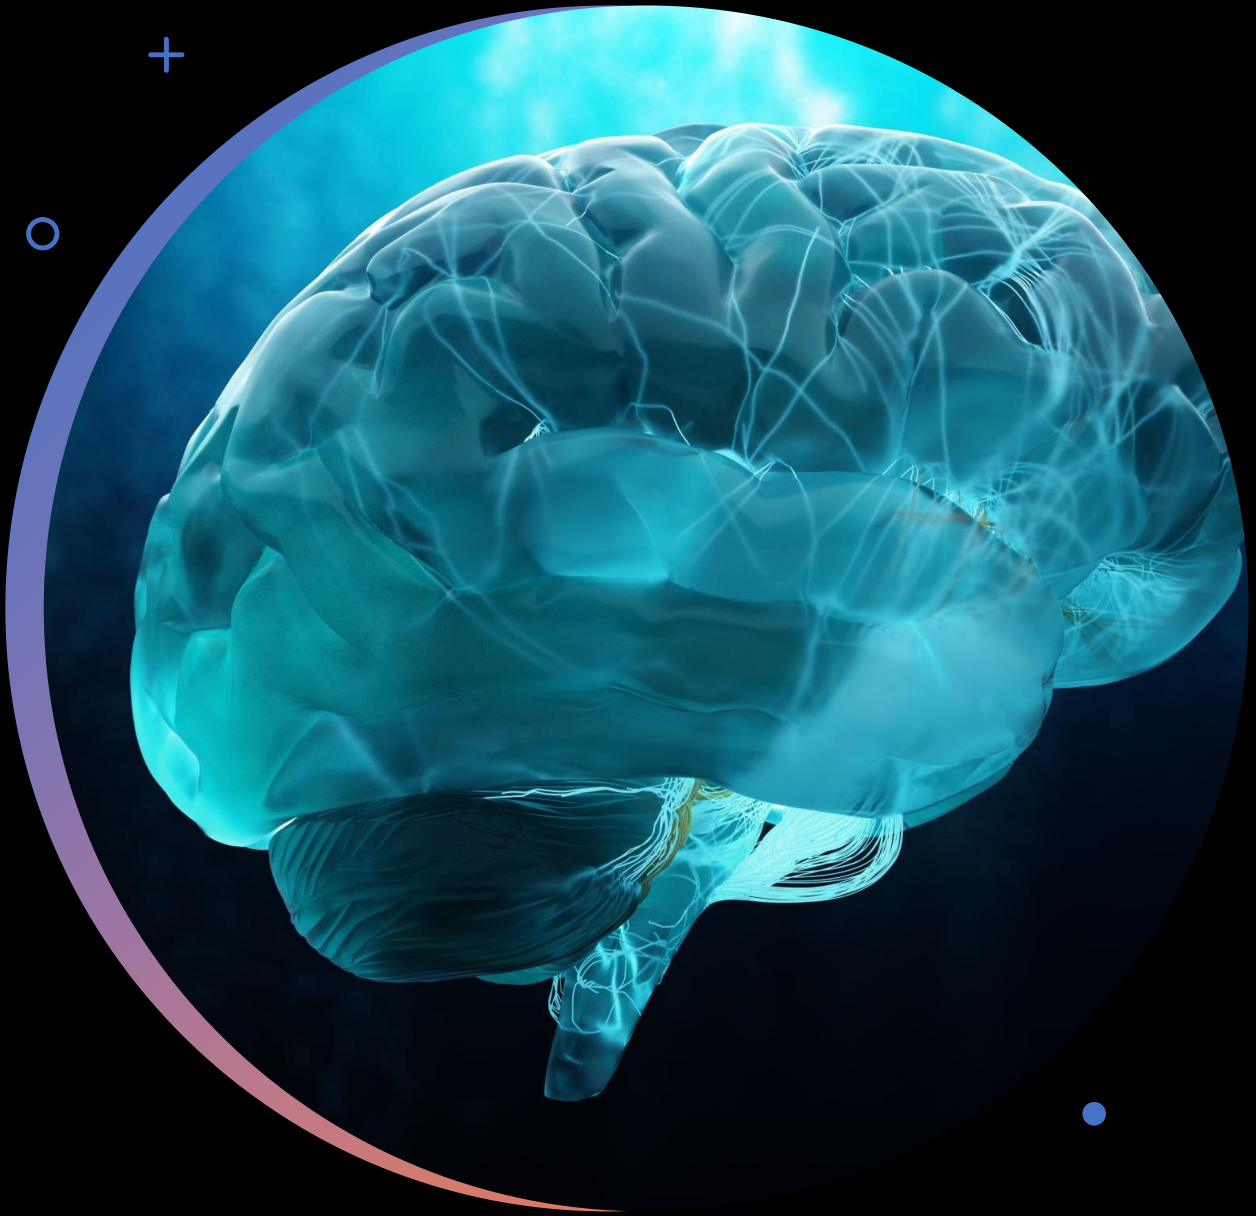

- Ample evidence in the safety literature that shows model parameters change significantly over time
- Ignoring temporal shifts leads to erroneous inferences and predictions

Flawed data analytics

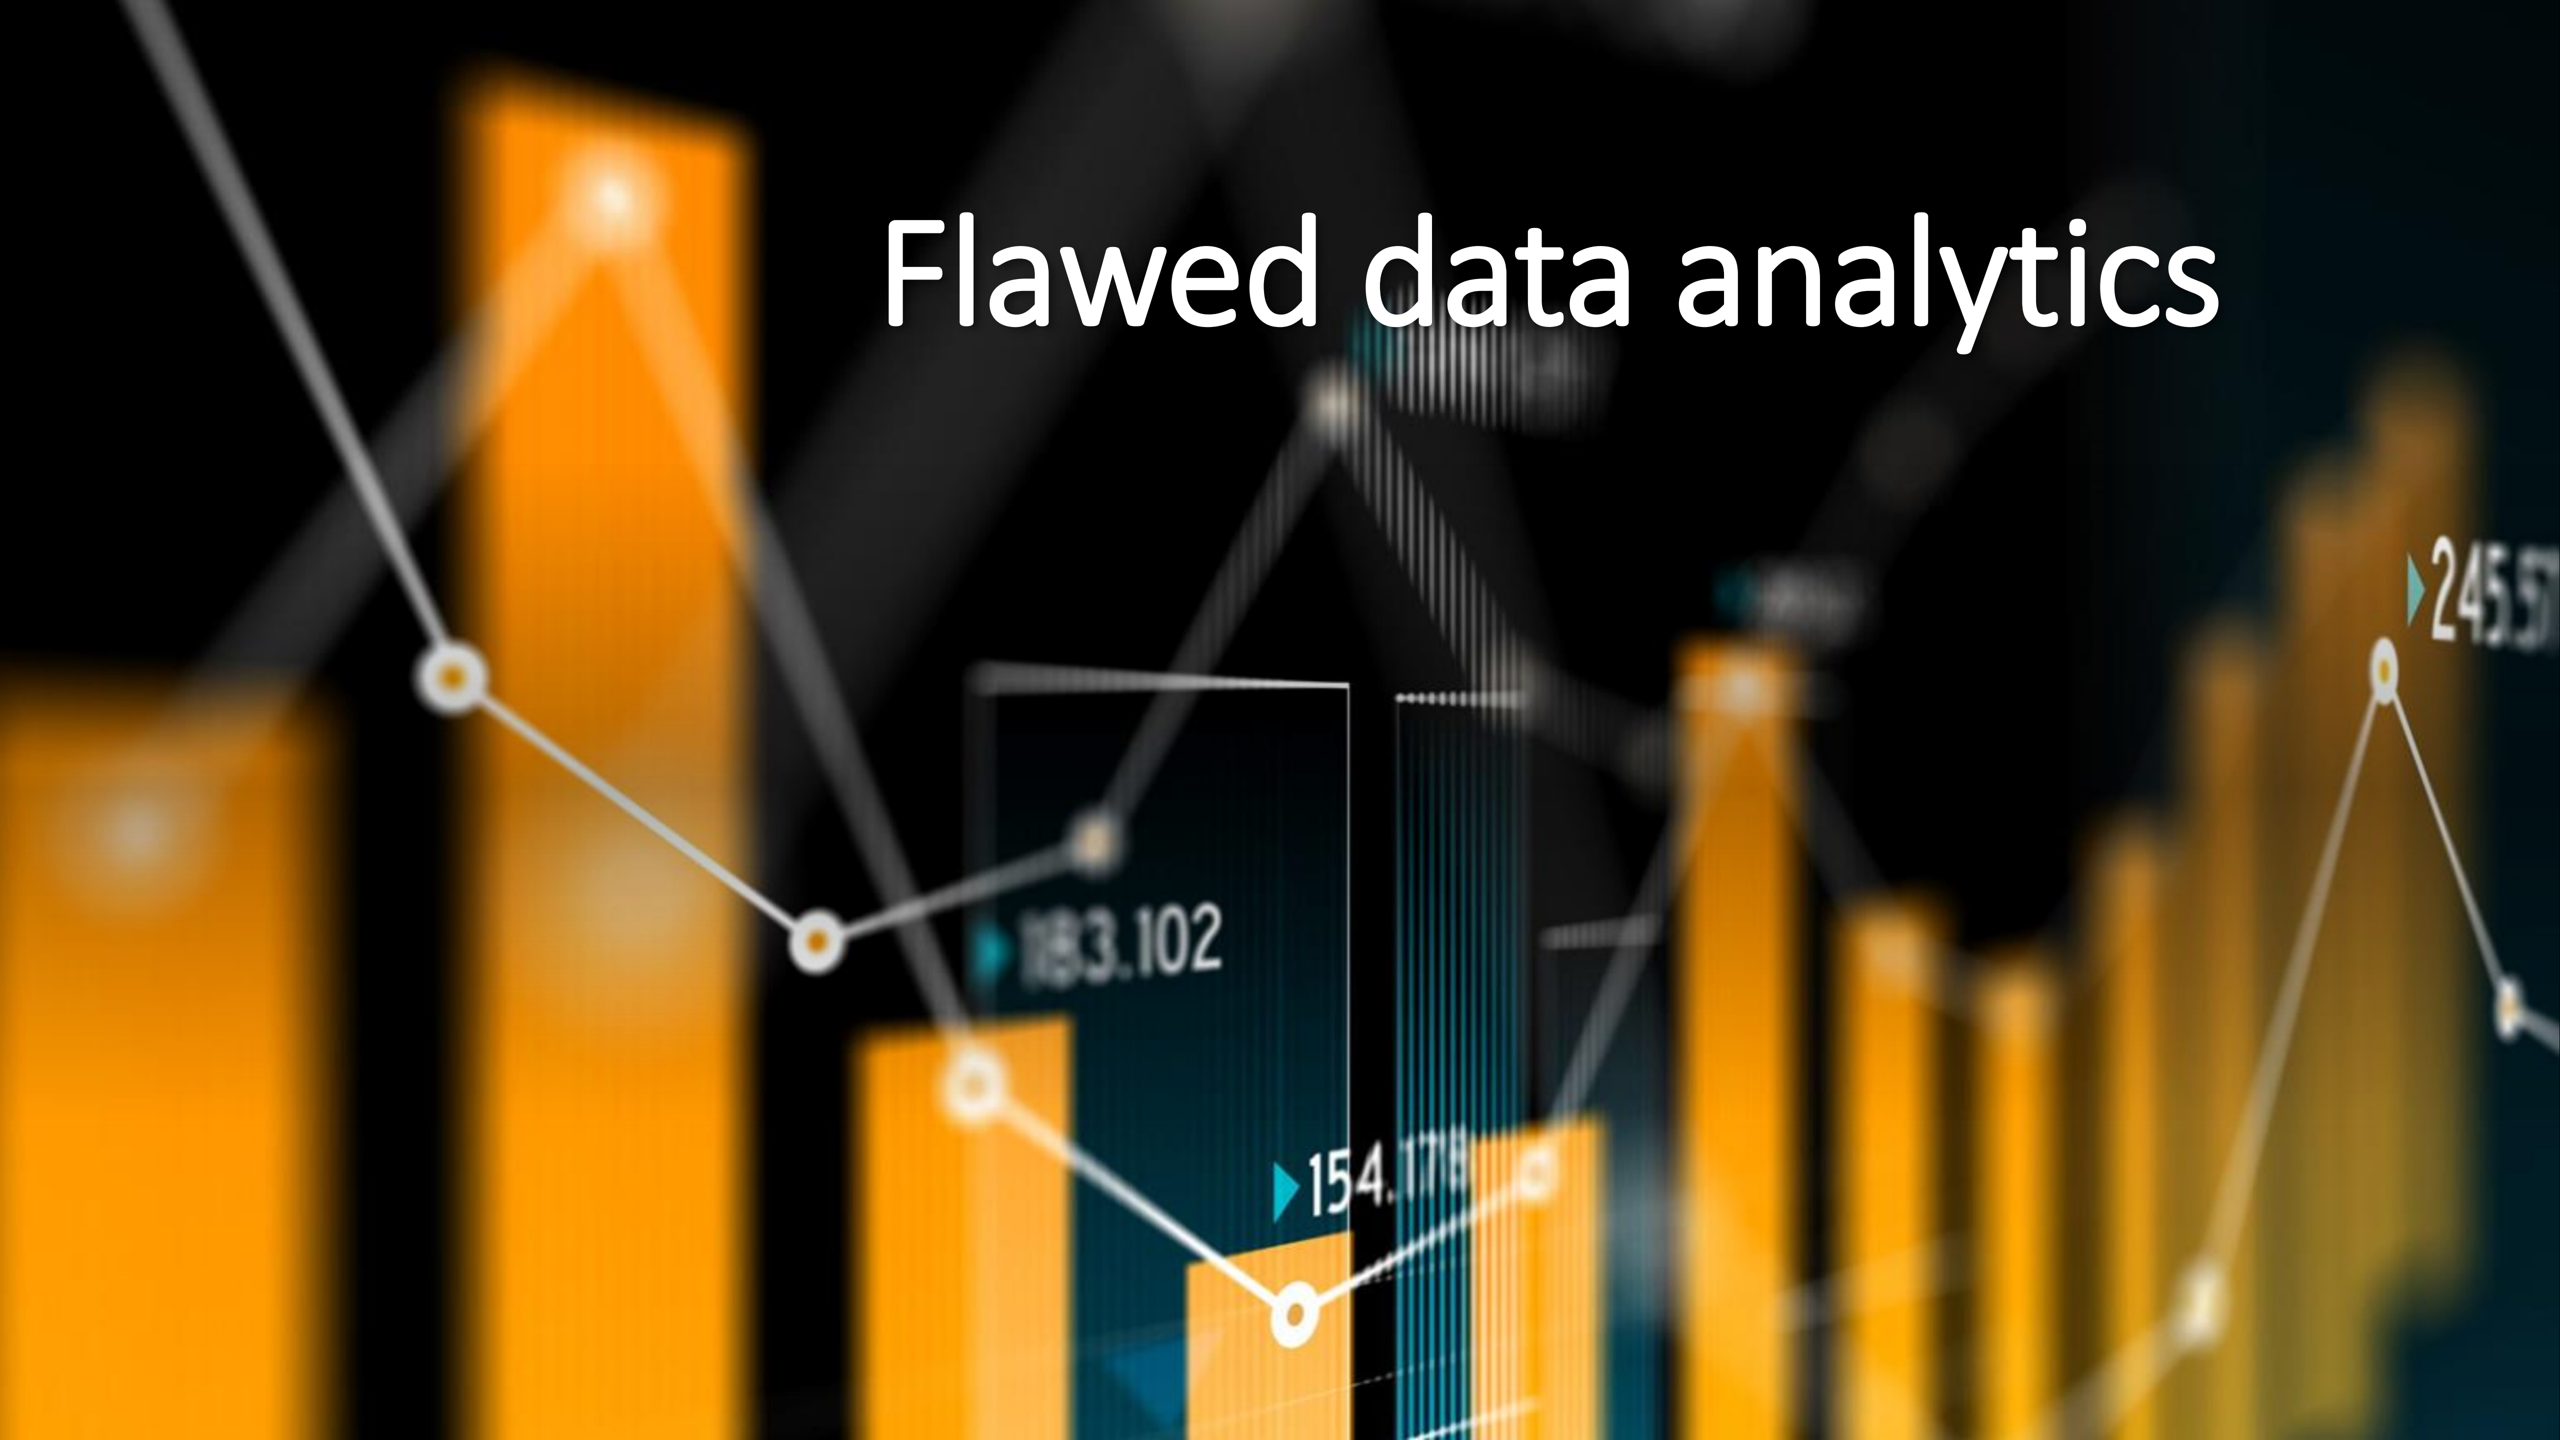

Selectivity/endogeneity/identification

The above issues that are systemic in highway safety data

Consequences of ignoring:

- Biased safety predictions
- Unable to identify underlying causality

Identification Example

Motorcycle safety:

- Model estimation shows a rain increases injury severity
- Alternate explanations:
  1. Rain does increase injury severity
  2. Only the riskiest riders ride in the rain and they have more severe crashes – the effect of the rain variable is pure selectivity and rain has no effect

Selectivity Example

Effectiveness of vehicle safety features:

- Side-impact airbags are found to reduce fatalities relative to vehicles without them
- But side-impact airbags seem to become less effective over time
- One reason: Self-selectivity. The safest drivers buy new safety features first and they tend to have less severe crashes anyway

Selectivity Example (cont.)

Side impact airbag effectiveness:

Insurance Institute for Highway Safety reported:

- 2004: 45% effective in reducing fatalities
- 2006: 37% effective in reducing fatalities
- 2008: 31% effective in reducing fatalities
- 2010: 25% effective in reducing fatalities

Again:

Drivers that first owned cars with side airbags are not a random sample of the driving population – they are safer drivers

The future of highway safety in the autonomous/human-driven vehicle environment

The need to seriously consider human behavior in data analytics

- Selectivity in data
- Temporally shifting behavior
- Compensating behavior
- Resulting externalities

Future possibilities:

- A.I. assisted driving:
  - Adjusts vehicle systems after learning driver responses to changing conditions
  - Can mitigate risk compensating behavior

Changing current autonomous vehicle thinking?

- Current approach:
  - Design autonomous vehicles to avoid collisions in response human-driven vehicle trajectories

- Alternate approach:
  - Design autonomous vehicles to interact human-driven vehicles to alter crash-critical vehicle trajectories before they happen

Likely consequences of ignoring human response to autonomous vehicles:

Human driving behavior will change around the predictable driving of autonomous vehicles

Changes in behavior around autonomous vehicles will create unsafe behaviors around human-driven vehicles increasing human-to-human crash rates

The future of highway safety?

- The fundamentals human behavior are still not being fully addressed
- While AI/ML opens new possibilities in data analytics, applications have not adequately considered human behavior
- The current approach is to simply hope that technology advances overwhelm any mitigating human behavior

Papers discussing these concepts

- Mannering, F., Bhat, C., Shankar, V., Abdel-Aty, M., 2020. Big data, traditional data and the tradeoffs between prediction and causality in highway-safety analysis. Analytic Methods in Accident Research 25, 100113.
- Mannering, F., 2018. Temporal instability and the analysis of highway accident data. Analytic Methods in Accident Research 17, 1-13.
- Mannering, F., Shankar, V., Bhat, C., 2016. Unobserved heterogeneity and the statistical analysis of highway accident data. Analytic Methods in Accident Research 11, 1-16.
- Winston, C., Maheshri, V., Mannering, F., 2006. An exploration of the offset hypothesis using disaggregate data: The case of airbags and antilock brakes. Journal of Risk and Uncertainty 32(2), 83-99.

END

Prediction/causality/data size trade-offs:

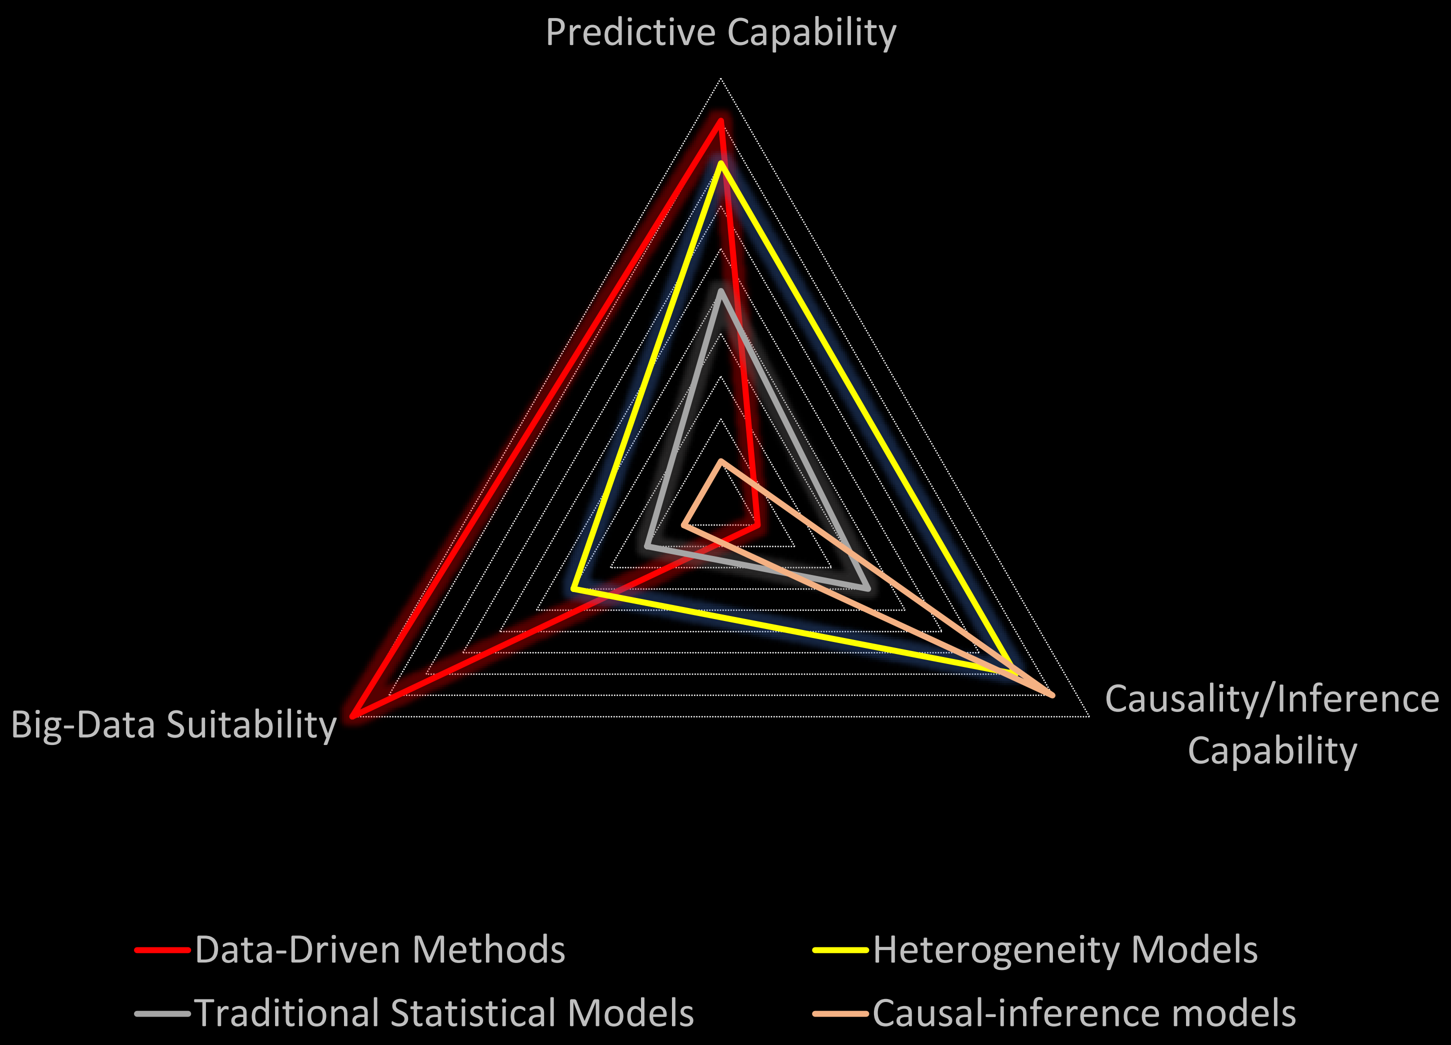

Selectivity example (with concurrent risk compensation):

Good Morning America 2006

Perspectives on Highway Safety: Contemporary issues and the forthcoming age of autonomous vehicles by Fred Mannering, 3/22/24

Matthew Webb

Part I. Literature (for further reading on why safety efforts have not been more successful)

Mannering, F., Bhat, C., Shankar, V., Abdel-Aty, M., 2020. Big data, traditional data and the tradeoffs between prediction and causality in highway-safety analysis. Analytic Methods in Accident Research 25, 100113. https://doi.org/10.1016/j.amar.2020.100113

Mannering, F., 2018. Temporal instability and the analysis of highway accident data. Analytic Methods in Accident Research 17, 1-13. https://doi.org/10.1016/j.amar.2017.10.002

Mannering, F., Shankar, V., Bhat, C., 2016. Unobserved heterogeneity and the statistical analysis of highway accident data. Analytic Methods in Accident Research 11, 1-16. https://doi.org/10.1016/j.amar.2016.04.001

Winston, C., Maheshri, V., Mannering, F., 2006. An exploration of the offset hypothesis using disaggregate data: The case of airbags and antilock brakes. Journal of Risk and Uncertainty 32(2), 83-99. https://doi.org/10.1007/s11166-006-8288-7

Part II. Recent News

Lazo, L. (2023, December 21). D.C.'s deadly streets take a growing toll on residents, victims. *The Washington Post.* https://www.washingtonpost.com/transportation/2023/12/19/dc-traffic-deaths-safety/

Leonhardt, D. (2023, December 11). The Rise in U.S. Traffic Deaths. *The New York Times*. https://www.nytimes.com/2023/12/11/briefing/us-traffic-deaths.html

Scott, D. (2024, February 12). The death of the world's best marathon runner is part of a troubling global trend. Vox. https://www.vox.com/future-perfect/2024/2/12/24071068/kelvin-kiptum-marathon-runner-car-crash-death-africa

Part III.

Q: You mentioned the aggressiveness of human visa via AV may propagate into a human to human interaction contact. Does that mean that we shouldn't design AV to be too conservative? Maybe building some degree of aggressiveness, on par with a human, would be better from the system perspective?

A: I should tell you I own 8 cars and they go from the 1960 all the way up to 2020. Now my newer cars have the full Level 1 with the collision avoidance, autonomous braking, and adaptive cruise control. And to your point that you notice how your driving behavior changes, when I start driving one of my Level 1 cars my behavior changes quite a bit, because now I know the car safer. I start messing with my Apple CarPlay and all of this distracted stuff. When I get into one of my 1960 or 1970 cars, your life is in danger and it becomes much more obvious. They don't brake as well. But we do have to be careful because this also case where we have self selectivity. If you look at the effectiveness of Level 1 safety innovations, you find they're pretty successful but you only observe the safest drivers driving these cars. As a good little anecdote, if you actually have these cars, an L1 car and a way non-L1 car (like my 1960s car that was like the first year with the shoulder belts and no safety features whatsoever), it's clear that that compensating behavior is gone on.

But to your point, it's not the way we design autonomous vehicles. We don't want them to design them so they have a higher risk of accidents, but there are ways we can keep it the same very low accident risk that still communicate to the surrounding human drivers that this car can still behave radically while not actually increasing the probability of a crash. For an example, if an autonomous vehicle is being tailgated, an algorithm that just starts flashing the brake lights, as an indication to the human driver and you don't even actually have to apply the brakes. You can just communicate with the drivers around the AV. I think this is going to be a critical safety element.

Q: You mentioned four reasons that we may make mistakes. The first one is a compensating behavior, and assuming that the degree of compensation will vary either by different technology or by different groups of people. Can you give an example where for a certain technology, some people compensate a lot while for another technology people don't compensate as much? What drives these different degrees?

A: The areas where they compensate the most are the ones that you would consider active safety features. For example, a passive safety feature would be like an airbag. You know it's there and you might do some compensation but an active safety feature that might be autonomous vehicle braking or lane departure warning. That is something that I think you have a lot of compensating behavior. To give an anecdotal example, some of my Level 1 cars have heads up display, so you just start focusing on that heads up display that has the little radar indicators that tell you if there's cars beside you and stuff. and I realize you just don't look in your side mirrors, that much anymore because everything is in the heads up display. But when you go back and use your non-Level 1 and non-heads up display vehicle, it's very dangerous. There's an adaptation period. Another example would be if you have a vehicle that has really good braking performance. Almost invariably, you'll start applying your brakes later and sort of mitigating some of the effectiveness of the braking system. So the active safety features almost certainly have compensating behavior, whereas the passive ones may have less so.

Q: So then to your second point about the externality, if at the individual level you could argue that people just have a different degree of tolerance. So it's my decision to choose where do I balance between the two. But if you choose a high risk balance, you also impose externality for other users who didn't choose to do so. When we have traffic congestion, we have the device of congestion pricing to address that. In terms of the safety externality, what are the meaningful policy interventions that we can implement to internalize this externality.

A: That's an interesting point. I'd have to think about that. I've been thinking that the most obvious solution to that sort of compensating behavior would be to develop a safety system in the car that mitigated the risk compensating behavior. I'm not sure if there's actually a policy that you could develop. There's a whole field of transportation safety that instead of looking at observed crashes instead looks at vehicle trajectories and near misses. I'll give you an example of a study that I did in real time adjustment. We did video analytics of an intersection where we looked at near misses and if we start getting a high number of near misses, then the signal timing gets changed to avoid these. You don't have to wait for a crash to occur as these things can be adapted in real time. In this case, perhaps unlike congestion, it's hard. We've made over the years many policy safety related policies, but I don't see an obvious policy to mitigate this compensating behavior because it goes on and a lot of people don't even realize that they're doing it. So there's, there's an awareness. It's not like, for drunk driving where you have to quit drinking. How do you do your risk mitigating? I guess you could develop a system in the car that determines when you're applying your breaks and gives you a warning for a late brake application or something. It looks instead like a policy that is centered around the specific technology.

Q: Your third point regarding the temporarily changing nature of human behavior, in particular you mentioned that some of the behavior changes due to the safety technology itself, but some of the human behavior evolution are totally exogenous. Can you give us some examples of that and also what are the causes for such evolution?

A: Behavior changes in response to technology, but there's a very large body of literature that has determined that behavior changes due to macroeconomic conditions. You may have noticed during the great recession in the late 2000s, there was a significant reduction in the number of crashes. Historically, what we have found looking at the aggregate data, during recessions there are fewer fatalities per mile driven and during good economic times there are many fatalities per mile driven, so if you're looking at our current economy it is doing great because we have a lot of fatalities. The people that have looked theoretically at what's going on here is there’s a lot of history that macroeconomic conditions change in fundamental ways the people's risk profiles and again there's a whole body of research that sort of look at how this effect.

Another element that inherently makes human behavior unstable is there a difference between attitudes and behavior. So, we have an attitude that, oh, I would love to save the planet, but then we have the behavior that I drive an SUV that gets 15 miles to the gallon. How do you reconcile this? As we go through life, there's always a difference between our attitudes and our behavior and there's a cognitive dissonance there, a sort of the Leon Festinger thing for the early sixties and we're constantly reconciling that dissonance and that constant reconciliation of the dissonance between your attitudes and your behavior inherently make your behavior sort of unstable. So you are adjusting your attitudes and then you're adjusting your behavior. So there's a lot going on. And if you look at the fields of economics, there's a work of Koopmans in the 1960s. There have been many people have looked at fundamental ways why people's taste and behaviors are changing over time. There's many causes of that.

Q: You contribute to both the classic statistics, econometric based methodology as well as the recent machine learning based methodology. How do they compare, particularly in the sense of falling into the victim of the four problems you illustrated in your presentation?

A: I think a lot of people that use the data driven methods like machine learning, their primary focus is on prediction. But in addition to prediction, there's causality because a lot of times just being able to predict something is usually short-term and doesn't give you any insight into what is going on, so when you think of traditional statistics and econometrics, there's certainly a prediction element, but there's also a search for the underlying causality. That is because it's the causality that will determine what an effective policy might be. I think that's one of the major problems with the application of machine learning in safety data. It is just too focused on prediction. For example, in the journal that I edit, I get many papers that show this model predicts better than the econometrics model. Yeah, of course it predicts better, but we want to know not only the prediction, we need to know causality because it's the causality that is going to determine the policy that actually affects the probability of fatality or whatever we're looking at in the safety area. That's the fundamental difference between the two and both have advantages. If you focus totally on causality, then your prediction may not be so good. You don't know what's going on. So there's a balance between those two that we have not yet struck in the field.

Q: I just watched a recent talk of yours where you categorize AI applications into three buckets. The first bucket is the area where AI will be just beautifully applied with nothing wrong with it. The second is AI with at least a moderate level of caution, and the third category is AI with extreme caution. Give some sense about the 3 categories.

A: There are some applications that have to be real time applications. You don't have time to think about the causality. You're collecting data and instantly providing a prediction. In that case, you're not so much interested in, the causality because it's like an instantaneous react. So that's an obvious application of machine learning because this can do the prediction instantaneously. When you get to longer term and think about how it might affect in the future, then you really have to think about the causal effects.

Q: One of the main points that you made was about the introduction of safety technology actually sort of changing people's behavior and causing them to embrace riskier behavior. I think it was great that you articulated the difference between active safety features on vehicles and passive safety, but I think it's worth just returning to that if you could explain that a little bit more, what's your evidence?

A: I think using the anecdotal explanation is perhaps the best, for example, in my case, when I'm driving my Level 1 autonomous car, you can notice that I'm noticeably less attentive to the roadway then I am driving one of my 1970s cars. We've developed these safety features, but we also have things like Apple CarPlay. We've developed all these safety features, but we have all these additional distractions. One of my 1960s cars was an MG, if any of you remember the English car makers, and they were so concerned about distraction that they would not provide a radio as standard equipment. They thought the radio was too distractive to driving. If we look at cars today, just think of what we have beyond the radio. Some of it is voice activated, but I can tell to get my Apple CarPlay to try and play Amazon music or trying to have it play one of my band’s CDs or something - I'm continuously distracted. We developed the safety in one path and developed the amenities, like Apple CarPlay, in another path. We're not thinking about how those two interact.

Q: I'm glad you raised distraction because that's exactly what I was driving at was if it's even possible to separate these two effects? The effect of increasing safety technology features, which as you argue, is driving riskier behavior on the one hand and distraction, which of course has been a major focus of road safety agencies around the world for some time and how to disaggregate those two.

A: I should mention that even though my diagram had speed on the horizontal axis, it's really speed slash distraction? Cause I find when I drive, my Level 1 cars, they're not as fast as my non-Level 1 cars, so I don't drive them faster, but I do drive them a lot more distracted.

Q: Have we accounted for increasing vehicle size and weight? Why are collisions decreasing in Europe? This is a related set of questions. One of our senior fellows at MIT, David Zipper, has been writing a lot about car bloat and vehicles just getting larger and larger and with higher surfaces at the at the front end of the vehicles resulting in greater risk and indeed fatalities to pedestrians. What are your thoughts on car bloat and then if that is a specific to the US conversation?

A I think car bloat is a very serious issue and electric vehicles tend to be much heavier. I should give you a little anecdotal story, just earlier this week, I bought a 1976 triumph TR6, which is a 2,500 pound little British sports car. But the car that was on the trailer that was backed out before it was some BMW I6 or something and I looked at the weight of that car is like over 6,000 pounds and the weight of my little roadster is 2,500 pounds. With car bloat I think we are seeing perhaps a greater variance in the weight distribution on cars on the road, perhaps more than any time in the history of the vehicle mix because now we have cars from like 2,500 pounds up to 6,000 pounds and over. These are just regular cars and you know not including heavy trucks and so on so that weight distribution I think imposes an externality. People with heavy cars or imposing an externality on those with lighter cars.

Now it's not just weight because you have crumple zones and everything. You can have a car that is lighter and still be pretty safe if you have very innovative crumple zones. But the car weight is again, It seemed like in the 1990s we were designing cars to be highly efficient from an air perspective. The coefficient of drag was very good and the consequences of a good coefficient of drag is you have a nice slope on the hood. It probably was safer for pedestrians as well. But since then, we've become less, concerned about coefficient of drag as we have more efficient engines. As you pointed out, a lot of the larger vehicles have these vertical grills and you know when you think of the injury and even a low kilometer per hour collisions, it can be pretty significant.

Now, with regard to Europe, it's interesting that almost all of the developed countries, except for the United States, have noticed significant improvements over time in, fatalities per kilometer driven. I'm not sure if it's a uniquely American phenomenon, what makes maybe people in the US more likely to risk compensate. I haven't followed the literature of the behavior of US drivers versus European drivers.

There's also an adaptation that goes on, let me give you an example. How people drive in Boston is much different than how they might drive in Tampa. So you can see spatial disparities within the US. Now let's say you had an autonomous vehicle that is actually learning from the vehicles around it which you have to be really cautious having it learn driving behaviors from vehicles around it. If you have an AI autonomous vehicle that was trained in Boston and then you started driving it in, Indianapolis, probably in the first 100 miles, you'd have like three wrecks. When we think of the US compared to other developed countries there's fundamental differences in behavior and that tends to be reflected in how effective the safety features are from one country to the next.

Q: Do these flaws mean physical interventions, for example, traffic calming, support public safety more successfully than safety features in individual vehicles? How do we achieve safe designs in the US transportation network given not just flaws in individual thinking, but what Roger Ruddock has called an institutional flaw with how we educate our transportation engineers. How do you think about the impact of road network design on safety?

A: At some point, maybe 60 years from now, we'll have all autonomous vehicles and then the problem will resolve itself. Then it just becomes a software problem. Until then, you're probably correct and I'll have to make sure I do another edition of my textbook Principles of Highway Engineering and Traffic Analysis, but when you think of what could be the most effective safety policy, it's probably not looking at controlling the individual behavior in specific vehicles, although artificial intelligence assisted driving once we get safety features enough we can start looking at those risk compensating elements. As pointed out in the question, I do think probably before we get to that level things like traffic calming are probably the best ways to address safety in the very short-term.

Now, It's interesting that there is a demand for safety features in vehicles. Certain people want to be safe and they tend to pay significant amounts of money for additional safety features and the market is responding to that. There will be people that are seeking those safety features. but if you want to make an immediate impact designing safer roads is probably more effective than adding more safety features to cars. Unless you add the safety features and then develop the AI to adjust the for the compensating behavior, then you might be on the same playing field. Without that, we probably overestimate the effectiveness of technology. We're saying, “we'll put this safety feature in it decreases crashes,” and as I mentioned before these safety features are often justified, “a 30% decrease in the probability that crash”. That number comes from the safest drivers because those are the ones in those cars. It's really not 30%. We have to think about that self-selectivity when we get into looking at individual safety features. I think designing safer roads is probably the best way to go. But there's a cost of that too; roads aren't cheap.

Q: We have trained transportation engineers for generations at MIT and other institutions to focus on this concept of level of service. And level of service is a concept that just says you want to move more cars more quickly through the road network. So it's interesting that that's the objective rather than you know, protecting people's lives. And I think that's a little bit of where at least in my mind it creates this cognitive dissonance and we need to sort of maybe rethink what the larger objectives are. Do you have thoughts on that?

A: When we design roads it's not all for a level of service. A lot of it is the actual site distance. There is a lot of accident prevention in there, for example stopping site distance. This is like chapter 3 in my text where we design roads for an anticipated eye-height, expected stopping distance of the vehicle, etc. It's not just the number of lanes. There is a significant safety element being considered in design of highways. For example, interstates are designed for 70 mile an hour driving speeds, but that's based on a standard vehicle with specified braking characteristics. If you have a high performance car, you could probably drive that road with the same level of risk at maybe 80 or 90 miles an hour. But for the standard designed vehicle, we design them for a certain speed to allow the vehicle to stop in a safe distance. We do consider that, but as you can imagine as was pointed out earlier, there's an incredible variance in vehicles on the road. You have heavy trucks that have terrible braking performance are not really the design vehicle because the drivers eye-height is so high you know when you do vertical and horizontal alignment of design that's you know that's much more effective than minimizing the braking distance. So we do consider elements of road design, but the high variance in the vehicles, that drive on the roads makes it a serious challenge. There's only so much we can do and designing the Newtonian physics of avoiding crashes.

When we look at it though, most crashes are caused by human error. That's the great promise of autonomous vehicles we get the humans out of the decision-making process, then it becomes much easier. If I give this talk 50 years from now, it's like, well, we did it. We have the machine learning and autonomous vehicles but to the point is a lot of people try to correlate traffic congestion with safety. Some people say if you have high congestion, it's inherently safer because you can't get to the speeds that cause fatalities. So there's an interesting interaction between level of service and safety. There's a lot of elements involved like the speed variance and the vehicles. There's a lot of things that have to be considered there.

Part IV. Summary of Memos.

Themes from Other Memos

1. Concerns over the absence of discussion on multi-modal transit systems and their influence on highway safety. Interest particularly in context of comparison to European cities with different infrastructure and lower pedestrian/bicycle death rates.
2. Concern over the focus on technology to address roadway safety and the importance of not only advancing AI and safety technologies but also incorporating urban planning and policy interventions. The example of Hoboken, NJ, illustrates that political commitment and the implementation of current technologies can lead to remarkable safety achievements, challenging the narrative that AI is the sole solution for the future.
3. Questions on how to appropriately measure the safety of a vehicle. In the US, vehicles are considered safe insofar as they protect the driver and passengers rather than pedestrians. Tension over whether we should be framing the question of highway safety around vehicles/drivers or those identified as increasingly at risk (i.e., vulnerable road users who are not driving).
4. Concerns over the paradoxical effects of increased car safety features fostering more distracted driving behaviors and a decrease in the long-term effectiveness of these technologies. Previous San Francisco Bay area Tesla crashes demonstrate the potential that AVs might inadvertently promote negligence among drivers who rely on machines for decision-making. Interest in exploring the problem from a behavioral science or psychology perspective.

My Reflection

This week’s forum, Perspectives on Highway Safety: Contemporary issues and the forthcoming age of autonomous vehicles, was a presentation by Professor Fred Mannering of the University of South Florida. Fred motivated his talk by reviewing the latest statistics showing disturbing trends of rising traffic, bicycle, and pedestrian fatalities in the United States in spite of improvements in vehicle safety technologies. He then spent the majority of his time review the four factors he identified that contribute to this trends: 1. Not accounting for the effects of safety innovations on human behavior, 2. Unconsidered externalities, 3. Not accounting for the evolution of human behavior, and 4. relying on flawed data analytics.

I appreciated the framework and explanatory graph Fred provided in his presentation to describe the first factor. It was a clean and elegant way to show he potential outcomes of implementing a new safety technology such as ideal effectiveness, reduced effectiveness due to compensatory behavior, or over-compensatory behavior that results in more danger due to overestimating a new feature's abilities. During the question and answer period, Fred went on to revise that the x-axis should be a combination of either increased speed or level of distraction by the driver. His comments included mention of the parallel, but often separate developments of in-car safety technologies and entertainment systems, but I would have appreciated more discussion on the potential gains in safety to be had by removing the sources of driver distraction.

Fred concluded his talk by discussing the potential for artificial intelligence assisted driving to counteract the compensatory driver behavior particularly during the adoption period of AVs when a mix of automated and human driven vehicle intermix on roadways. For instance, he envisions an AV that detects tailgating by a following human driver so the AV begins strategically flashing its brake lights to cause the human driver to back off. I'm skeptical of this example though as it seems that once the countermeasures of AVs like this became known and experienced by human drivers, then the humans would just adjust as his third factor states and develop new compensatory behaviors to improve their personally realized performance.

Part V. Other Information

Other questions: How should we measure car safety? How do car safety advancements from AI fit into the improvement of safety across the highway transportation system?

Other literature (referenced during the Q & A):

Festinger, L. (1962). Cognitive Dissonance. Scientific American, 207(4), 93–106. https://www.jstor.org/stable/24936719

Mannering, F. (2023, December 13). *Data analytics and the promise of Artificial Intelligence to improve highway safety*. https://www.youtube.com/watch?v=_6MnWjpd8l8

Mannering, F., Washburn, S. S., & Kilareski, W. P. (2009). *Principles of highway engineering and traffic analysis* (4th ed.). John Wiley.

Rudick, R. (2024, March 5). *American-Dutch Engineer Calls Out Root Problem with American DOTs—Streetsblog San Francisco*. https://sf.streetsblog.org/2024/03/05/its-not-just-politics-american-dutch-engineer-calls-out-root-problem-with-american-dots