

Metadata

#Author	sachin.kansal@uber.com	SEMANTIC
#Date Modified	07/19/2022	SEMANTIC
#DateCreated	07/18/2022	SEMANTIC
#Title	Safety @BoD July '22 - Product FAQ	SEMANTIC
Account	mike.akamine@uber.com;	SEMANTIC
All Custodians	Akamine, Mike;Kansal, Sachin;	SEMANTIC
All Paths	Akamine, Mike: \Akamine, Mike\EDISCO-23471_mike.akamine@uber.com_16.zip; Akamine, Mike: \Akamine, Mike\EDISCO-23471_mike.akamine@uber.com_16.zip; Akamine, Mike: \JCCP_Drive001\JCCP_Drive001_14.zip; Akamine, Mike: \JCCP_Drive001\JCCP_Drive001_14.zip; Kansal, Sachin: \EDISCO- 25937_sachin.kansal@uber.com\EDISCO-25937_sachin.kansal@uber.com2_37.zip; Kansal, Sachin: \EDISCO-25937_sachin.kansal@uber.com\EDISCO- 25937_sachin.kansal@uber.com2_37.zip	SEMANTIC
Application	Microsoft 2007 Word Document	SEMANTIC
Begin Family	UBER_JCCP_MDL_002059720	SEMANTIC
Collaborators	dharmin@uber.com; mike.akamine@uber.com; aarti@uber.com; jennifera@uber.com; nadiam@uber.com; natalia.galvez@uber.com; rebecca.payne@uber.com	SEMANTIC
Confidentiality	Confidential	SEMANTIC
Date Created	07/18/2022 6:59 am	SEMANTIC
Date Modified	07/19/2022 8:42 pm	SEMANTIC
DocID	1xLHGvWTHYYvWZ2fCDUoTz87SCgANTTGP1UQVt_gVspQ	SEMANTIC
Document Type	Electronic File	SEMANTIC
End Family	UBER_JCCP_MDL_002059724	SEMANTIC
File Path	\Akamine, Mike\EDISCO-23471_mike.akamine@uber.com_16.zip	SEMANTIC
File Size	15122	SEMANTIC
Filename	Safety @BoD July -22 - Product FAQ_1xLHGvWTHYYvWZ2fCDUoTz87SCgANTTGP1UQVt_gVspQ.docx	SEMANTIC
GoogleDocumentType	DOCUMENT	SEMANTIC
Hash Value	f2753c91b34416ab2eaae934ab4f0b4d	SEMANTIC
Hidden Content	Yes;	SEMANTIC
ILS All Bates	UBER_JCCP_MDL_002059720;UBER_JCCP_MDL_002059721;UBER_JCCP_MDL_0 02059722;UBER_JCCP_MDL_002059723;UBER_JCCP_MDL_002059724	SEMANTIC
ILS Document Date	07/19/2022	SEMANTIC
ILS Prod Date	1/3/2025	SEMANTIC
ILS Prod Vol	JCCP_MDL057	SEMANTIC
LINKSOURCEBEGBATES	UBER_JCCP_MDL_002169036; UBER_JCCP_MDL_002169040; UBER_JCCP_MDL_002337349; UBER_JCCP_MDL_004918354; UBER_JCCP_MDL_004979197; UBER_JCCP_MDL_005313782	SEMANTIC
Other Custodians	Akamine, Mike;Kansal, Sachin;	SEMANTIC
Primary Date	07/18/2022 6:59 am	DOC_TYP E_ALIAS
Production Volume	JCCP_MDL057;	SEMANTIC
Redacted	Yes	SEMANTIC
Sort Date	07/19/2022 8:42 pm	SEMANTIC
SourceHash	53706652263c33b8e0ffcea5fb16d039	SEMANTIC

Trial Exhibit No.
P-00798

July 19 2022
Uber Board meeting
Safety Product section ([narratives/slides](#))

Safety Product FAQs: Being prepared for questions that could come up in the board meeting.

- Can you pls add links to additional documents
 - Latest audio deck - [US expansion](#)
 - Latest Dashcam deck - [BYOD analysis](#)
 - Latest IAVR deck - [Checkpoint 2](#)
 - Latest teens deck [\[UberTeens\] ELT Review Q2](#)
 - Latest Left turns deck [MTR Maps + Safety collab](#)
 - Latest Driving insights deck [Driving Insights Review - May 2022](#), [Master Deck](#)
 - Latest [selfie deck](#) as well as other [Rider Identity deck](#) (multiple is ok)
 - Latest SRAD deck [\[BOD\] SRAD Overview July 2022 | A/C Privilege](#)

- What kind of impact have you seen with audio recording?
 - Media reaction and qualitative survey feedback from users has been largely positive.
 - We have not seen incident rate impact, but believe this is largely due to how the product has been positioned and used - awareness is low and it is used mostly during safety incidents, as opposed to being used as an 'insurance policy' where users record all of their trips, just in case. We believe using the product like this will have the most potential impact on incident rate (because folks will know that recording is happening before they get in the vehicle), and will also be most useful for support (as the recording will capture the whole development of the incident, not just the part after the user started feeling unsafe). We are working on comms campaigns and new feature development to allow for auto-recording of trips, so that users can set it and forget it.

- Video recording! How are you getting around privacy issues and single/dual consent markets? Will you be limited in where you can roll out?
 - We've built all of our recording products with privacy as a key requirement, and have consulted with various 3rd party privacy advocate organizations. As much as possible, we try to limit access to recordings, unless it is being used for safety purposes. For example, with our audio recording product, the recording is stored encrypted on a user's phone - so they can't access it and we can't access it, unless they choose to release that recording to Uber. This prevents the public sharing of recordings and privacy violations we often see with personal dashcams.

Commented [1]: @nat@uber.com @rebecca.payne@uber.com @mike.akamine@uber.com @dharmin@uber.com @aarti@uber.com pls see some questions and requests below for the BoD meeting prep.

Commented [2]: @rebecca.payne@uber.com this one for you

Commented [3]: Added.

- Consent is the other key factor. In one-party consent markets, only one person needs to know that the recording is happening - so the person recording obviously knows this is happening and we fit the requirement. In all-party consent markets, everyone in the vehicle needs to consent to being recorded. This is more difficult to manage as we only have direct communication with the person who called the ride, but not necessarily the folks they are bringing with them on the trip. So today, we rely on notifications to the rider to let them know that their trip will be recorded, and this will suffice for all of the one-party consent markets. But in order to expand to all-party consent markets, we will need to do more work on physical signage - helping drivers get access to and confirming that they are using stickers and signs in their car to alert all riders that they will be recorded. So this is not off the table, but requires further investment from an operations and tech perspective.
- Is reducing left turns increasing ETA and making routes less efficient?
 - As part of the experiment we ran to reduce left turns, we closely monitored metrics like ETAs, rider cancellations and ATAs. We did not see changes to any of these metrics as part of the XP. Part of the XP design was to look at offline analysis including a larger left turn penalty, which caused problems in guardrail metrics, so we are confident we landed on the right configuration. Our next step in impact improvement is to look at harsher penalties for specific sub-types of left turns.
- Speeding: Is driving above the speed limit considered speeding? If so, that would mean everyone is speeding.
 - For most of our speeding products, including driving insights, we're looking for signals of high levels of consistent speeding across a trip. To identify these instances, we're looking to identify whether a driver was going 30% or more over the speeding for 30% or more of the trip.
- **Selfie:** Do you validate the selfie with their ID? Why does just taking a selfie increase safety?
 - Checks: No - we validate that the photo was taken with a front facing camera, that there is only one person in the image, and that the person is not wearing glasses or hats.
 - We believe that selfies improve safety because they reduce anonymity - we deter bad actors from committing safety incidents as they know they can be held accountable for their actions.
 - Experiment: We tested asking riders for a Selfie in Panama, Honduras, Paraguay, Peru, Ecuador, and South Africa. We ran an experiment against Social Connect, a proven (2017) method for incident reduction.

- We found that in small, safer countries, Selfies work great - we saw a 20% increase in the verification pass rate and sign-up to first trip rate with no change in the incident rate.
 - But in bigger, more unsafe countries, our current selfies product is not enough of a deterrent - in Peru, Ecuador, and South Africa we saw a 50-100% increase in the incident rate. We rolled back the feature in these countries.
 - H2 plans: we are building a deduplication check for our Selfies feature as our analysis showed that incidents were committed by the same person (same face) creating different accounts. We are also tightening up our fraud checks by using a new photo-of-a-photo model.
- Teens: What age group are you targeting? Are they going to take an Uber regardless?
 - 13-17 y.o.
 - Yes, based on recent user research we estimate that around 10-20% of trips happening in our platform today are taken by unaccompanied minors (% trip share varies by market).
- S-RAD: When you determine a pairing is more risky, isn't that just a sign that one of those people is not safe enough? Should we just take them off the platform?
 - We have a separate deactivation program which identifies users with a pattern of risky behavior and takes them off the platform. For example, when we receive a report of a serious sexual misconduct or a pattern of low level sexual misconduct incidents, we take the user off the platform.
 - SRAD, on the other hand, evaluates the risk of a pair in a given situation, and not an individual. Consider this situation - It is Friday night at midnight. Dan has been in a bar with his friends. He is getting ready to go home, so he requests a trip with Uber. Dan is generally an okay rider, but has received low ratings from drivers in similar trip situations. Dan's request can potentially be paired with any of the 5 drivers we see here. Chloe is closest to Dan, but is a relatively new driver. Joe is a little farther away than Chloe is. But Joe is a seasoned driver with 3x more lifetime trips than Chloe. Our hypothesis is that if Joe is dispatched for this request, we reduce the probability of a safety incident on this trip.
 - The hypothesis is that some of these incidents are crimes of opportunity. So if we reduce the opportunity, we reduce the probability of an incident from happening. Now because safety incidents are rare, SRAD is a very low precision machine learning model. The model is right about 50 times out of a million predictions that a serious sexual assault or sexual misconduct incident may happen. Hence, despite a high SRAD score, the likelihood of an sa/sm on the trip is still extremely low. So blocking drivers or riders seems like a blunt tool to prevent incidents. Instead, we deprioritize such high score pairs in order to reduce likelihood of such a dispatch. While prediction has low precision, the cumulative effect of millions of soft interventions has a significant impact on incident rate.

- Which is the most predominant factor in the model?

Every trip has a different relative feature importance that drives the risk. However, on the whole, here are the top factors

- Trip hour and day of week
- Rider & Driver incident rates
- Rider & Driver star ratings
- City & local area incident rate
- Rider's late night request rate

- Example of preventing a **rider on driver** incident

- It is Friday night at midnight. Dan has been in a bar with his friends. He is getting ready to go home, so he requests a trip with Uber. Dan is generally an okay rider, but has received low ratings from drivers in similar trip situations. Dan's request can potentially be paired with any of the 5 drivers we see here. Chloe is closest to Dan, but is a relatively new driver. Joe is a little farther away than Chloe is. But Joe is a seasoned driver with 3x more lifetime trips than Chloe. Our hypothesis is that if Joe is dispatched for this request, we reduce the probability of a safety incident on this trip.

- Example of preventing a **driver on rider** incident

- It is late Friday evening. Chloe is in a bar with her friends. She's getting ready to go home, so she requests a trip with Uber. Her request can potentially be paired with any of the 5 drivers we see here. Bob has the lowest ETA, but has received some negative feedback on similar trips. Joe is a little farther away than Bob is. But Joe hasn't had any issues with late night trips, has 3x more lifetime trips and 3x fewer 1-star ratings than Bob. Our hypothesis is that if Joe is dispatched for this request, we reduce the probability of a safety incident on this trip.

- S-RAD: Are you using gender at all in this algorithm?

- We are not using gender in the US Model, but we do use it in Brazil. We saw a 15% lift in the model's ability to predict incidents if we used gender information.

REDACTED - PRIVILEGED

- How do you conduct fairness testing?

- We knew that this product would raise a lot of concerns, and some of the biggest would be around discrimination. Even with a safety benefit, we had to make an affirmative case that this product wasn't harming historically marginalized groups.

Commented [4]: @sachin.kansal@uber.com
Assigned to Sachin Kansal

- We've conducted three rounds of fairness testing in the US and Brazil. We've looked to see whether, for example, riders in HUD-qualified census tracts have to wait longer for a trip, or whether female drivers see a drop in earnings. And we've consistently found no disparities of any statistical or practical significance. In the US, we've looked by gender, ethnicity, geography, and for drivers who self-identify as non-binary or transgender where we have that data.
 - Now, this is reassuring, but we recognize it's not good enough for Uber to just say – there's no problem here, nothing to worry about. And that's where the process comes in. To that end, we worked with a well-respected civil rights auditor to help shape and then vet our fairness testing analysis and safety results. We presented the work to a host of civil rights groups to hear and incorporate their feedback. We worked with two NGOs in this space – Upturn and NYU Law's Center on Race, Inequality and the Law – to write a public-facing paper on best practices for this type of testing when you don't have self-reported ethnicity data. And in the long-term, we plan to replicate our work with an external academic as well.
- S-RAD: Why not announce this?
 - While we are confident in the safety, marketplace and fairness of SRAD, the product is complex to explain and can be easily misinterpreted by the general population. So we are taking a more cautious approach on where we fit the message to the audience. We have updated our privacy notice to explicitly state that data may be used in driver-rider matching to improve safety. We have a "What moves us" page on Uber's website where we disclose that Uber may modify pairings in certain instances to help maintain a safer platform. In the long-run, we will develop a deeper dive on the results of the algorithm, targeted at academic audiences.