

Metadata

#Author	gus@uber.com	SEMANTIC
#Date Modified	12/05/2016	SEMANTIC
#DateCreated	12/05/2016	SEMANTIC
#Title	Issues with Safety	SEMANTIC
Account	travis@uber.com	SEMANTIC
All Custodians	Fuldner, Gus;Kalanick, Travis	SEMANTIC
All Paths	Fuldner, Gus: \JCCP-EDISCO-23800_BEFORE2018\JCCP-EDISCO-23800_BEFORE2018_gus@uber.com\JCCP-EDISCO-23800_BEFORE2018_gus@uber.com_24.zip; Fuldner, Gus: \JCCP-EDISCO-23800_BEFORE2018\JCCP-EDISCO-23800_BEFORE2018_gus@uber.com\JCCP-EDISCO-23800_BEFORE2018_gus@uber.com_24.zip; Kalanick, Travis: \JCCP_EDISCO-23992_DriveTravis\JCCP_EDISCO-23992_DriveTravis_3.zip; Kalanick, Travis: \JCCP_EDISCO-23992_DriveTravis\JCCP_EDISCO-23992_DriveTravis_3.zip	SEMANTIC
Application	Microsoft 2007 Word Document	SEMANTIC
Begin Family	UBER_JCCP_MDL_000172204	SEMANTIC
Collaborators	travis@uber.com	SEMANTIC
Confidentiality	Confidential	SEMANTIC
Date Created	12/05/2016 9:03 am	SEMANTIC
DocID	1pxz1u5Dj26CN4eeZaLQcj6TARrpJeDEUwqk1EfCx-TU	SEMANTIC
End Family	UBER_JCCP_MDL_000172207	SEMANTIC
File Path	\JCCP_EDISCO-23992_DriveTravis\JCCP_EDISCO-23992_DriveTravis_3.zip	SEMANTIC
File Size	11434	SEMANTIC
Filename	Issues with Safety_1pxz1u5Dj26CN4eeZaLQcj6TARrpJeDEUwqk1EfCx-TU.docx	SEMANTIC
GoogleDocumentType	DOCUMENT	SEMANTIC
Hash Value	69eeea4a576d889717ce3d9c6f643f2	SEMANTIC
Hidden Content	No;	SEMANTIC
ILS All Bates	UBER_JCCP_MDL_000172204;UBER_JCCP_MDL_000172205;UBER_JCCP_MDL_000172206;UBER_JCCP_MDL_000172207	SEMANTIC
ILS Document Date	12/05/2016	SEMANTIC
ILS Prod Date	8/19/2024	SEMANTIC
ILS Prod Vol	Vol039	SEMANTIC
Other Custodians	Fuldner, Gus;Kalanick, Travis	SEMANTIC
Primary Date	12/05/2016 9:03 am	DOC_TYP E_ALIAS
Production Volume	Vol039;	SEMANTIC
Redacted	No	SEMANTIC
Sort Date	12/05/2016 12:29 pm	SEMANTIC
SourceHash	2b11513752dc350dce81074ed0c9514b	SEMANTIC

Trial Exhibit No.
P-00694

ANALYSIS OF SAFETY AT UBER

Summary:

Issue 1: High fragmentation of groups that manage or influence safety decisions

Sub Issue 1A: Diffuse decision making:

Sub Issue 1B: Low Ops Integration:

Sub Issue 1C: Core Safety Team does not play a clear coordinating role:

Sub Issue 1D: Limited domain expertise in core team:

Sub Issue 1E: Cost of incidents separated from prevention / response

Issue 2: Use of data in decision making

Sub Issue 2A: Metrics and definitions:

Sub Issue 2B: Analyzing small data sets:

Sub Issue 2C: Product priorities not sufficiently data driven:

Sub Issue 2D: Insufficient focus on severe incidents:

Issue 3: Onboarding / Deactivation and Safety not coordinated

Sub Issue 3A: Onboarding Conflict of Interest:

Sub Issued 3b: Little use of safety data in Onboarding / Deactivation:

Details:

Issue 1: High fragmentation of groups that manage or influence safety decisions

Major groups:

- a. Safety Product (Nirveek De)
- b. Safety Research (Heather Rothenburg)
- c. Safety Intitatives (Kate Parker)
- d. Data Science - Trust & Safety (Qi Dong)
- e. Regional Safety Managers (Transferred from T&S to Ops as of September, scattered throughout ops at Regional or Mega Region level)
- f. Driver Product - Compliance (Chris Weber)
- g. Driver Product - Driver Recognition [Ratings] (Mike Truong)
- h. Legal - Safety (Candace Kelly)
- i. Legal - Insurance Litigation (Katie Waitzman)
- j. Comm Ops - IRT (various - Roger Kaiser [coordinator])
- k. Policy Information & Research - Safety (Dorothy Chou)
- l. Corporate Communications - Safety Incidents (Brooke Anderson)
- m. Insurance (Gus Fuldner)
- n. Insurance Claims (Chris Boedeker)
- o. Insurance Actuarial (Frank Chang)

- p. Policy Information & Research - Economics (Jonathan Hall)
- q. Engineering - Trust and Safety
- r. Engineering - Driver Compliance

Sub Issue 1A: Diffuse decision making: Decision making is very diffuse and has become more diffuse over time. Shedding IRT (2015), BGCs (Q1 2016) and along with RSMs (Q3 2016) from the safety org has reduced coordination and effectiveness

Sub Issue 1B: Low Ops Integration: Safety team (especially leadership) is mostly hired externally and has a low share of former ops folks or people with deep internal organizational connections to effectively implement changes.

Sub Issue 1C: Core Safety Team does not play a clear coordinating role: A distributed team might be viable if the core team played a strong coordinating role, but there is not a clear plan or process for soliciting stakeholder input or keeping stakeholders informed. As a result topics frequently bounce around from group to group until all relevant stakeholders are found: Example topics: Helix Ratings, Global Vehicle quality review, IRT review. Lots of decision making and issue escalation is lead by policy and comms team

Sub Issue 1D: Limited domain expertise in core team: Other than Heather Rothenburg, there are few people with a background in traffic or road safety. Thus there is a lot of “reinventing the wheel” going on as well-meaning people work on topics from first principles. The safety data science team in particular uses generalist analyses vs. traffic safety or insurance specific methods.

Sub Issue 1E: Cost of incidents separated from prevention / response and ops: The cost of safety incidents comes mainly in the form of insurance premiums, self-insurance reserves, and settlements. The cost sits in city P/Ls under the management of the insurance team and main prevention and response levers sit in the safety team. Cities do not bear the cost of individual incidents in their P/Ls and instead get the blended cost of incidents through insurance. Better alignment of interests may be achieved by passing costs of city specific performance to city teams.

Issue 2: Use of data in decision making

The collection of reliable data and the effective use of data in decision making has been a consistent issue.

Sub Issue 2A: Metrics and definitions: Metrics are not well defined and keep changing. For example OKR for H1 was based on safety incident data that was highly influenced by settings of “Arnold” text classifier and baseline rate of incidents changed with each release of the ML model. There are many different definitions of “incident” or “hard break” in different products. L1-L4 are very coarse definitions that make actionable analysis difficult(e.g. minor incidents = L1, vs measuring left turn accidents with property damage)

Sub Issue 2B: Analyzing small data sets: Safety Data Science and Product teams repeatedly doing analysis on specific cities (e.g., Houston) over small periods of time and projecting it elsewhere. For example the “Bouncer” model is based on handful of cities as was most telematics “Delphina” research. Safety incidents are rare events and need very large data sets to draw useful conclusions. We have a large global business and analysis should be going on at the country level to achieve stronger statistical significance. These are Hadoop questions not Excel / Python on desktop questions.

Sub Issue 2C: Product priorities not sufficiently data driven: There is a lot of experimentation on the product side without a clear analytical framework to define success. Example: Delphina (in app trip safety scorecard) developed and rolled out without a clear connection to accident rates. Which factor is most important? Why these specific factors? What is the right threshold for a “hard break” or “hard acceleration”. Projected outcomes such as “20% reduction” in driving incidents are unachievable from the beginning because they are not grounded in analysis.

Sub Issue 2D: Insufficient focus on severe incidents: Metrics and efforts have been focused on all incidents (L1-L4) vs. the severe ones that matter (L3&L4) and should be the primary focus. This is what led to the lack of focus on sexual misconduct matters for most of 2016 until Q3

Issue 3: Onboarding / Deactivation and Safety not coordinated

Onboarding (BGCs, MVRs) and Deactivation are huge drivers of cost (both direct cost of BGCs and indirect cost of Deactivations) and are obviously a major driver of safety (perhaps the largest lever). However BGCs operations are largely housed in product and ops and not tightly coordinated with safety policy. Safety team involvement is focused on emerging markets (e.g., psychometric tests).

Sub Issue 3A: Onboarding Conflict of Interest: Onboarding team sits within driver growth. There is an inherent conflict between being in charge of growing new supply and managing compliance and a need for checks and balances. This is like the “fox guarding the hen house” in terms of setting priorities. Chris Weber is really strong, but his broader team has different objectives than compliance (cost and first trips) which affects resource allocation.

Sub Issued 3b: Little use of safety data in Onboarding / Deactivation: There is also little use of safety data in establishing the thresholds for driver screening (should it be 3 MVRs, which minor offenses should be excluded?) or deactivations (thresholds set on rider ticket #s vs safety incidents our outcomes). Insurance team has been doing some work in this area and managed to get insurers to relax vehicle year criteria in the US after demonstrating low correlation to safety incidents leading to ~6% of incremental US supply in October. We shouldn't be limiting new drivers in the name of safety and should know which criteria matter and which don't in great detail.

Outcomes:

- A. Safety incidents have been increasing on a per-trip/per-mile basis, but there is disorganized data to support change or analyze drivers.
- B. Uber is viewed as unsafe by many. Safety is a major barrier to new rider growth (top survey response for why not to use Uber), U4B, and driver acquisition (especially women drivers) as a result of media coverage of safety incidents
- C. Safety response (IRT) is not performing effectively (though significant strides made in last few weeks) and is not well instrumented to provide data for improvement.
- D. We are probably sub-optimizing the tradeoffs between onboarding criteria and safety