

**PREPARATION SHEET**

# **Uber: Applying Machine Learning to Improve the Customer Experience**

1. Where does the pickup experience actually break down?  
Draw on your own experience as a rider or manager of a service operation. Identify two or three moments in the Uber pickup sequence where friction is most costly - not merely inconvenient, but strategically damaging. Who bears the cost, and why does it matter to the business?
2. Does Uber have one customer or two?  
The pickup experience must satisfy both the rider and the driver simultaneously. Where do their interests align, and where do they conflict? What are the implications for how Uber defines "quality"?
3. What does Uber actually know - and what can it only infer?  
Uber cannot directly observe whether a pickup was good or bad. It can only observe signals. List the active, passive, and third-party signals available to the company. Which signals are most reliable? Which are most easily distorted? What does this tell you about the limits of data-driven decision-making?
4. Build a pickup quality metric.  
Using the signals available to Uber, propose a quantitative metric for pickup quality. Specify which variables you would include, what weights you would assign, and why. Be prepared to defend your weighting choices against alternatives.
5. When should the algorithm defer to a human?  
Uber's ML model works well for recurring, predictable situations. But some contexts - a post-concert surge, a road closure, a local festival - require judgment that no algorithm has been trained to handle. Where exactly should the boundary be drawn between automated and human decision-making, and who should draw it?