



CHIARA FARRONATO

ALAN MACCORMACK

SARAH MEHTA

Innovation at Uber: The Launch of Express POOL

The mood was tense in a conference room at Uber's headquarters in March 2018. For the past several months, the San Francisco-based ride-sharing company had been testing a new product called Express POOL (Express). Express offered a reduced price to riders willing to carpool, walk a short distance to/from their pick-up and drop-off points, and wait for two minutes before being matched to a driver.

The Express product was similar to POOL, except that it offered a cheaper ride in exchange for **walking** and **waiting to be matched**. When a rider requested Express, the Uber app asked her to wait for up to two minutes while the back-end algorithm assessed potential matches given the pick-up and drop-off locations of nearby riders. The app then matched that rider to others heading in the same direction and instructed her to walk to a designated pick-up point to meet her ride. Longer initial wait times enabled the app to make more efficient matches, ensuring that the car was at full seating capacity for as much of the trip as possible and making it financially feasible to sell trips at even lower prices. Riders' tolerance for waiting, however, was finite, and the company recognized the need to balance efficiency with rider experience.

Since concluding pilots in Boston and San Francisco, Uber's data scientists had been running experiments in these two cities to test a number of improvements to Express. Express had been conceived with a maximum two-minute wait time, but by early March, the results of an experiment looking at Boston riders' reactions to waiting five minutes had just come in. Duncan Gilchrist, head of data science for Uber's rider pricing and marketplace experimentation teams, had called a meeting to discuss the results. "The Boston experiment shows mixed results," began Gilchrist. "Longer wait times increase cancellation rates, but reduce our costs per ride."

Gilchrist waited as his colleagues considered how these results might impact the ongoing rollout of Express. Most pressing, just two weeks earlier, Uber had begun a new experiment (launch experiment) in 12 U.S. cities. In this experiment, Uber had launched Express in six "treatment" cities and held constant six additional "control" cities for comparison. The data science team had placed a five-week moratorium on changes to these 12 markets. This freeze was meant to allow one week for the markets to stabilize and then four weeks for data collection, data which Uber used to evaluate the impact of the new Express product on market equilibrium and company profits. Now, armed with data from the Boston experiment about the impact of longer wait times on costs, Uber needed to decide whether to

Professors Chiara Farronato and Alan MacCormack and Case Researcher Sarah Mehta (Case Research & Writing Group) prepared this case. It was reviewed and approved before publication by a company designate. Funding for the development of this case was provided by Harvard Business School and not by the company. Certain details have been disguised. HBS cases are developed solely as the basis for class discussion. Cases are not intended to serve as endorsements, sources of primary data, or illustrations of effective or ineffective management.

Copyright © 2018, 2020, 2023 President and Fellows of Harvard College. To order copies or request permission to reproduce materials, call 1-800-545-7685, write Harvard Business School Publishing, Boston, MA 02163, or go to www.hbsp.harvard.edu. This publication may not be digitized, photocopied, or otherwise reproduced, posted, or transmitted, without the permission of Harvard Business School.

overrule the five-week moratorium and increase wait times from two to five minutes in the six treatment cities midway through the launch experiment.

Ronak Trivedi, senior product manager for shared rides, listened with interest as Gilchrist explained more. Trivedi was focused on cancellation rates. “The fact that cancellation rates increased is a sign that customers *hate* waiting,” he said. “We should keep waiting times to a maximum of two minutes.” At the other end of the table, product manager Miraj Rahematpura had been typing numbers on his laptop. Looking up, he said, “But the reduction in costs per ride is huge. That would more than offset the increase in cancellations. We should increase wait times in our six treatment cities immediately.” “I understand both your points,” replied Gilchrist, “but these results are from just a few neighborhoods in Boston. We have no way to validate them for other cities. Besides, properly collecting baseline data from the 12-city experiment is important to inform future improvements to Express, so we must wait five weeks before changing the product.” Everyone looked to Ethan Stock, director of product management for shared rides, for a decision. “Well Ethan,” said Rahematpura, “can we increase wait times in the treatment group of our 12-city experiment?”

Uber and The Ride-Sharing Industry

Uber was founded in 2009 by serial entrepreneurs Travis Kalanick and Garrett Camp as an on-demand luxury car service targeted at executives in Silicon Valley. The company connected riders to unoccupied private black cars.¹ In mid-2012, to target a larger customer base, Uber introduced a new product, UberX, which was roughly 35% cheaper than the company’s luxury car offerings.² UberX allowed any licensed driver over the age of 21 with a car in good condition to drive for Uber. Potential drivers were required to submit to a background check and a review of their driving record and car registration. Uber provided insurance coverage for its driver-partners (i.e., Uber’s term for drivers) while they drove for the company. Driver-partners could work as much or as little as they wished.

Throughout the 2010s, Uber grew quickly, expanding into international markets and adding new products. Uber tended to launch products as soon as was feasibly possible, reflecting a maxim applied to several disruptive technology companies: “move fast and break things.” As Gilchrist explained, “Uber has this culture of being very experimentally driven. We tend to put a minimum viable product in the market and then iterate based on what we learn.”

By 2018, some 75 million riders and 3 million driver-partners used the Uber platform.³ On any given day, Uber drivers completed 15 million trips in 600 cities across 65 countries.⁴ The company had raised \$21 billion across several funding rounds and was valued at \$62 billion, making it the most valuable startup in the world.^{5,6} Because it was a private company, information about Uber’s financials was relatively sparse. Observers generally believed that Uber operated at a loss, but that its financial performance varied widely between cities. In some of its more mature markets, the company was thought to be profitable.⁷

Uber was a first-mover in the growing ride-sharing industry. By 2018, a number of other U.S.-based ride-share companies had emerged, such as Lyft and Wingz, along with global players, most notably the Chinese company Didi Chuxing Technology Co. (Didi) and Singapore-based Grab. All offered some variation of ride-sharing, which could be broadly defined as an on-demand service that connected idling independent drivers to waiting passengers for a fee. The service was powered by point-to-point software and GPS mapping installed on riders’ and drivers’ smartphones.

Despite the growth of ride-sharing platforms, access to these services remained concentrated in urban areas. Where these services were available, the ride-sharing industry was highly competitive.

Low switching costs between service providers permitted both riders and drivers to utilize multiple platforms with ease.⁸ To gain market share, companies had aggressively lowered rider fares and offered large bonuses for new drivers. Of U.S.-based companies, however, Uber remained the market leader, claiming 77% of the domestic market.⁹

Uber's Platform and Product Offerings

By early 2018, Uber offered eight ride-hailing products: Express and UberPOOL (cheaper, carpool options), UberX, UberXL, and UberSELECT (the company's core economy offerings), and UberBLACK, UberSUV, and UberLUX (premium, more expensive options).¹⁰ (**Exhibit 2** describes each product type.) Product availability varied by city.

To request a ride, users opened the app and entered their desired drop-off point. The app then prompted users to choose their preferred product. In some cities, the app also listed the projected fare and the anticipated arrival time. Once users requested the ride, the app matched them with an available driver and, if a carpool option was selected, with other riders. The app showed riders their driver's name and the license plate number for identification purposes, as well as the driver's customer rating, which ranged from 1 (the lowest rating) to 5 (the highest). Once matched, the rider could monitor the driver's movements in real time. While on their journey, riders could follow the trip's progress on a map. After the ride, the app prompted the rider to rate the driver. Uber charged the rider's credit card. Riders had the option to tip their drivers and leave comments. (**Exhibit 3** shows the rider app.)

On the driver's side, when a ride request came in, the app pinged nearby eligible drivers one at a time. Uber's driver-partners could either accept the request, passively reject it — meaning they failed to accept the ride within a 15-second window — or actively refuse a rider. If the driver accepted the ride, he or she picked up the passenger and then followed the directions in the app to the destination. Upon completing the ride, the app prompted drivers to rate passengers. Additional features on the driver app included a "heat map," indicating areas with high demand, as well as an earnings icon that showed the driver's pay.¹¹ (**Exhibit 4** shows the driver app.)

Riders' fares varied by product type, but were generally based on the trip's length and distance. Drivers were paid a base fare for each ride as well as a set amount per mile and per minute (**Exhibit 5** shows a typical driver's earnings breakdown). Uber kept 25% of a trip's gross fare.¹² While some media outlets reported that Lyft and Uber drivers' take-home pay was exceedingly low,¹³ Uber's own research pegged its driver-partners' median earnings between \$15 and \$30 per hour (see **Exhibit 6** for estimated driver earnings disaggregated by the number of hours driven).¹⁴ The company used dynamic pricing, charging higher "surge" prices when demand outpaced supply. For example, if a surge rate of 1.8x went into effect for a given neighborhood, a normal \$10 fare would be \$18 until demand and supply recalibrated.¹⁵

In addition to its core ride-sharing products, Uber had expanded into other verticals, such as Uber EATS, a food delivery service, and Uber Freight, a service that notified trucking companies of freight awaiting transport. Uber was also working toward advancing autonomous vehicle technology.¹⁶

Organizational Structure

Uber's San Francisco-based product teams were organized into three key verticals: rider, driver, and marketplace. The rider vertical was responsible for the rider-facing app, rider recruitment and customer service. The driver vertical handled analogous tasks for drivers. The marketplace vertical maintained an overarching view of the health of all Uber products, monitored substitution patterns between products, and developed the systems and technologies behind pricing and vehicle matching.

In addition to these verticals, Uber also had a central operations team and individual city teams that ensured smooth product launches and alerted engineers and product teams of issues.

Within the three verticals, staff performed a number of functions, including product management, engineering, data science, product operations, design, and marketing. Product managers shepherded the product development and improvement process, considering the perspectives of all stakeholders. Engineers and data scientists developed the technology behind Uber's products and evaluated improvements to its algorithms. Product operations specialists liaised with technical teams and city operations teams to ensure product-market fit. Designers curated the look and feel of Uber's website and rider- and driver-facing mobile apps. Finally, marketing teams created advertising and marketing campaigns to ensure consistent product messaging. (**Exhibit 7** provides an organizational chart.)

Product operations specialist Bradford Church clarified the role of product operations. "Global teams need local feedback to refine products," he said. "We work closely with city teams to collect information and feed it back to the engineering teams. A good example is how we integrate with airports. Engineering teams might only be familiar with how a couple of airports work, so, without product operations, they might build a product feature that works well at the San Francisco airport but does not work for the vastly different airport setups around the world."

Uber's product development process involved teams staffed with representatives from four functions: a product manager, a data scientist, a designer, and an engineer. "Engineers are practical, designers are aspirational, and data scientists are tactical," Trivedi said. "The product manager takes multiple perspectives into account and helps make decisions." Typically, teams designed a minimum viable product and then invited others within the company to examine it and tell the team what aspects to cut and which to add. "The process is painful," said Trivedi, "but necessary."

Engineering manager Danny Guo believed that for the most part, Uber employees maintained collegial relationships, even across functions. "There is a level of trust between product managers and engineers here," he said. "Engineering tends to be very technology-driven, which means that we do not always consider the way that products will perform in the real world. Product managers and operations specialists give us that perspective, which helps us build pragmatic products."

Innovation at Uber

Innovation at Uber spanned a spectrum with regard to the degree of product change involved. At one end of the spectrum, Uber made continuous incremental improvements to its core products. As Guo said, "We are constantly iterating. Software only lasts for about 18 months here." Next, Uber identified and developed new ride products, such as Express, that were optimized for a different set of customer preferences (e.g., for price-sensitive customers who did not mind waiting and walking). Finally, Uber placed bets on riskier ideas, based upon fundamental changes to its business model and technology (e.g., Uber EATS and autonomous vehicles).

To understand riders' experiences, Uber ran rider surveys and looked at proxy data for rider satisfaction (e.g., ride re-requests, driver ratings) in the app itself. With drivers, the company gathered information more directly, through interviews and other interactions. Uber also followed online forums frequented by its driver-partners to gauge their satisfaction. Hamid Nazerzadeh, a staff data scientist in Uber's marketplace optimization team, said, "We are especially sensitive to drivers' feedback because they are such important partners for us."

A key element of Uber's innovation strategy was its substantial investment in data science. Of the 200 employees staffing the company's marketplace vertical, for instance, 60 were data scientists; these

individuals typically held doctoral degrees in data science and related disciplines from a top-ranked school. (See **Exhibit 8** for the proportion of data scientists employed by Uber as compared to its peer companies.) Uber employees believed that the use of data science at the company was relatively advanced. As Gilchrist noted, “Generally, data scientists at Uber have a fair amount of influence. That’s because the types of problems we’re solving require us to focus on a combination of algorithms, user experience, and scale.” As senior data scientist Connan Snider added, “Uber is unique. A lot of places are engineer-driven because they are dealing with straightforward production problems. But we are sorting through very complex issues that require data scientists to understand and evaluate how users interact with our technology. We start by adjusting algorithms manually, and then build systems to update automatically based on what we learn about demand and supply. For example, we recalibrate many of the parameters in our dynamic pricing algorithms on a weekly basis.”

Uber ran different types of experiments to improve its products, depending on the type of improvements to be tested. Among those most commonly used were user-level A/B experiments, switchbacks, and synthetic controls.

User-level A/B Standard, user-level A/B experiments compared the behavior of app users to test the effects of platform decisions. For example, say Uber wanted to understand how differences in product placement within the app affected riders’ propensity to select one product type over another. In a standard A/B experiment, the platform would randomly allocate riders into either the “treatment” group or the “control” group. When users in the control group opened their app, they would see the standard app design. Users in the treatment group, by contrast, would see the newly proposed product placement on the app. After some time, Uber would compare the frequency with which riders selected one product type over another in the treatment versus the control group to estimate the degree of change in behavior induced by the different app design.

Switchbacks “Switchbacks” were another type of study design used to evaluate the effects of a product tweak on some outcome variable of interest. Say, for example, that Uber wanted to test an improved algorithm for matching riders to drivers. In a switchback experiment, the data science team would expose all riders and drivers in a given market to Uber’s standard matching algorithm for a 160-minute period. During the subsequent 160-minute period, riders and drivers would be exposed to the revised matching algorithm. There would be an odd number of switches per day, ensuring that if a Monday evening rush hour was in the treatment group in the first week, it would be in the control group the following day and the next Monday. Uber would continue to switch back and forth for a two-week period, and compare differences in efficiency and customer satisfaction metrics between the two groups. “We try to keep our switchbacks as clean as possible,” explained Gilchrist. “The problem is that we can only really run one at a time in a city to prevent them from interacting with each other.”

Synthetic control experiments These experiments attempted to create treatment and control cities to study the effects of a product tweak on a set of city-level outcomes. Synthetic control experiments randomized treatment and control at the city level, and despite their name, were not an application of the statistical synthetic control method.¹⁷ For instance, if Uber wanted to understand how a new rider app affected market demand, the company would roll out the new app in a random set of treatment cities and measure total requests and cancellation rates in those cities. Uber would then study the same outcome variables in a group of similar cities that remained unexposed to the new rider app. Because these experiments assigned all riders and drivers within a city to either the treatment or the control group, they allowed the company’s data scientists to evaluate how product tweaks affected city-level outcomes. However, noted Nazerzadeh, “With synthetic controls, in order to detect any changes, they need to be fairly significant – usually 5% or higher – but lots of our experiments generate effects of smaller magnitudes.”

Uber teams ran several experiments simultaneously. As product operations specialist Jane Lee noted, “At any given time, we have teams experimenting with rider-side and driver-side features, and other teams making smaller product tweaks, all within a finite number of cities. So, figuring out how to launch new products and test their effects without interfering with other experiments is challenging,” Nazerzadeh added, “For Uber, the problem is that it is easy to run out of cities in which to experiment.”

Gilchrist had been tasked with keeping track of these experiments and helping teams accurately interpret their results. “In general, it is very hard to measure the effects of our product tweaks because there are a lot of external factors to take into account,” he explained. “There are also network effects to consider. In the past, Uber had a bit of a ‘wild, wild West’ approach to experimentation. Teams were presenting results that overstated effects, so when I came on board in mid-2017, there was a general consensus that we needed to have more trustworthy experimentation.”

To impose more order on the experimentation process, Gilchrist and his team developed a process called the Marketplace Change Protocol (MCP) for scheduling major experiments. (Major experiments were defined as those that could have a significant impact on products other than the target.) During a weekly steering committee meeting, teams hoping to run a major experiment gave presentations explaining the goals and risks of the proposed experiment. Uber’s head of product and head of engineering made the final decision about which experiments would move forward. Smaller experiments, like user-level A/B, were not subject to the same review; hence teams could still implement and run these experiments on their own.

As the importance of data science had grown at Uber, so had its role in driving innovation across the business. Trivedi noted, “It’s not just measuring the results of small experiments. Some big ideas can come out of our own observations around travel patterns supplemented with rapid experimentation. For example, three years ago, we noticed that lots of people were traveling along similar routes at similar times. That’s one of the observations that led to UberPOOL.”

Shared Rides: Launching UberPOOL

In 2014, Uber launched UberPOOL (POOL), its first shared rides product, which offered a discounted fare to riders willing to carpool with other passengers. POOL was still a door-to-door pick-up and drop-off service, with no walking required. Uber hoped that this option would generate higher “seat utilization” — a key metric of interest for the shared rides team — thereby increasing the company’s overall ridership and boosting earnings per ride. POOL drivers were paid based on ride time, distance, and surge rates. In addition, they received a rider pick-up fee of between \$0.50 and \$1.00 for each additional passenger in the car.¹⁸ This meant that if two separate passengers were both going from point A to point B, a driver taking both passengers would be paid between \$0.50 and \$1.00 more than a driver taking only one passenger. Media outlets estimated that passengers’ fares on a POOL trip were roughly half those of an UberX trip.¹⁹

Initially, to generate interest in POOL, the Uber app asked UberX riders to push a button called “I’m Feeling Lucky” if they were willing to share the car with a co-rider in exchange for a price cut. If the app located a suitable co-rider, the UberX became a POOL, and the original passenger’s fare was cut in half. If the app found no compatible co-rider, the original passenger took a normal UberX and paid the UberX fare. “We thought the ‘I’m Feeling Lucky’ option would be popular,” explained Trivedi, “but people gravitated toward paying a bit more for an uninterrupted journey. You just never know what will happen in the market until you release a product. Optimizing toward your assumptions is foolish.” Uber soon switched to modeling the probability of matching. If the app predicted that riders on a given

route would be matched with others, it offered the reduced POOL fare upfront. Riders who chose POOL then paid this discounted fare regardless of whether they matched with a co-rider.

To identify co-riders for POOL trips, Uber used a “greedy algorithm.” Essentially, all POOL rides began as UberX trips. Once on the trip, the algorithm constantly looked for other passengers to add to the trip. If it located a rider whose pick-up and drop-off locations were within certain parameters relating to the original passenger’s route, it directed the driver to pick up the new rider. The algorithm permitted POOL drivers to pick up multiple additional riders, either until the car hit its seating capacity (i.e., three passengers) or the driver had completed 95% of a given rider’s trip.

In many markets, adoption of POOL was growing at a faster rate than UberX. While Uber was pleased with rider uptake, the greedy algorithm system posed challenges for both riders and drivers. POOL riders, for instance, tended to react negatively to the loops and detours necessary to pick-up and drop-off co-riders, especially if they required the driver to backtrack. Drivers worried that POOL might leave them vulnerable to poor reviews due to factors beyond their control. Moreover, the product’s capacity to match co-riders was not as efficient as it could be, with about half of all POOL trips left unmatched. As a result, by the end of 2016, POOL was still unprofitable.

Express POOL Project

In early 2017, Peter Deng, head of Uber’s rider vertical, began advocating for the company to rethink its shared rides strategy. He believed that products like POOL would drive Uber’s future growth, but that they needed to be far more efficient to reach their potential. Several Uber executives bought into Deng’s vision, and the company made improving the efficiency of shared rides an organizational priority. In July 2017, the company’s leadership created a joint task force between the shared rides team and the marketplace team, and asked that they make a series of improvements to shared rides. Deng placed director of product Stock in charge of the task force.

Stock explained, “The merged team comprised product managers and operations specialists, engineers, and data scientists, who all had different perspectives. They had to make hard trade-offs between rider experience and cost efficiency. Products like POOL need to be high-quality, because if they aren’t, people won’t take them, but they also have to increase our earnings. In an ideal world, these objectives would be complementary, but they often conflict.”

Stock noted that his role was primarily to set up the structures and processes to enable the merged team to make good decisions. “I really relied on the team to come up with the details,” he said. “I asked them to create specific metrics to quantify a positive trip experience, rather than relying on vague feelings.” The team identified metrics for measuring rider experience. These included, for example, opt-in rates (i.e., the percentage of total riders who requested a shared ride) and rider cancellation rates. To measure cost efficiency, the team looked at the number of occupied seats per minute and per mile; the ideal scenario was three riders to one driver (i.e., the maximum seating capacity) for as much of the trip as possible. The more riders on a shared trip, the more revenue that Uber earned.²⁰

By August 2017, the team had adopted two key strategies for improving shared rides. They would ask riders to: 1) wait up to two minutes while the algorithm matched them to co-riders, and 2) walk a short distance to/from their pick-up and drop-off points. If done well, this would result in fewer detours and better matches. Uber had already tested walking and waiting in a few key markets. In New York City, for instance, the company in 2016 launched HOP, a product requiring riders to walk to a fixed pick-up point. Custom-built for this market, it accounted for the city’s many one-way streets. Trivedi, who had been deeply involved in the launch of this product, noted, “We put HOP in the

market, and riders loved it. They intuitively knew that it made more sense to meet their driver at the correct side of an intersection. Riders started rating drivers more highly. Drivers' lives got easier too because they no longer had to circle around a city block to pick up passengers. That resulted in better rider ratings. Everyone was happier." Uber had also tested waiting in Chicago and San Francisco by offering riders the option to wait a few minutes to be matched to a ride in exchange for a discount.

August 2017: Adapting the JIT Algorithm to Express

Once the team had settled on adding walking and waiting to shared rides, Uber's leadership asked engineering manager Guo to lead a "tiger team," or a group of ten engineers, to build a new algorithm for rider matching. As Rahematpura summarized, "The customer learnings from HOP carried over, but the software technology did not because it was custom-built for New York City." Whereas POOL's greedy algorithm matched ride requests on a first-come, first-served basis, this new system would use a two-minute window to batch all requests and active rides to jointly find the optimal allocation of passengers to drivers. Stock explained, "By having riders wait to be matched upfront, we exponentially expand the number of possible overall matches because there are hundreds of ride requests coming onto the platform in that timeframe. This allows our system to find highly compatible riders at massive scale, providing both efficiency and quality."

To build this new algorithm, the tiger team adapted an existing system called "Just In Time" (JIT), which had been built six months earlier to improve the driver dispatch process for UberX. Rather than matching rides one by one, the JIT system delayed driver dispatch a few seconds, during which Uber looked at all ride requests across product types and matched a driver to that UberX requester as efficiently as possible. In early 2017, Uber had tested the new JIT driver dispatch system with Uber EATS's food delivery service. As Guo quipped, the company felt comfortable testing the new software on food delivery because "sandwiches don't have feelings." The tiger team would need to extend the batching time window and adapt the JIT system to accommodate co-rider matching and walking.

One early debate was whether pick-up points should be fixed prior to matching, as they were for the HOP product, or whether they should be determined by the new algorithm. "We discussed this issue at length," recalled Trivedi. Dynamism ultimately won over customer preference for fixed pick-up points. As Trivedi said, "In the end, I understood that our efficiency gains will stem from our ability to be as flexible as possible." In lieu of fixed pick-up points, the team decided to select and tag a series of "corners" as possible pick-up points for a given area (see **Exhibit 9**). Once the rider was matched with others, the app directed her to walk to the most convenient corner for pick-up, as determined by the location of her co-riders and driver. Product operations specialists worked with city teams and Uber's engineering teams to refine corners, ensuring that they did not place riders in danger or ask them to complete impossible tasks, such as fording a river to reach the pick-up point.

Stock's team debated whether to simply add walking and waiting to the existing POOL product or to create a new shared rides product. Ultimately, they chose to launch a new product called Express. Stock explained, "We were concerned that if we told people who were used to POOL that they now had to walk and wait, we would be perceived as insensitive to safety and accessibility concerns. We wanted people who were used to POOL to still have that option." In markets without POOL, however, Uber could make Express the default POOL product. For example, UberPOOL in Australia, launched in early 2018, was actually the Express product.

It took Guo and the tiger team two months to develop the additional JIT software required for Express. Their development plan included two milestones: version 0 (v0) which accommodated walking and matched riders and drivers within a fixed waiting time, and v1, which added flexible waiting. In the latter, a rider was guaranteed to wait up to the maximum waiting time, but if the

algorithm found a good match beforehand, she would be immediately notified about her driver and pick-up location. “Throughout this process,” said senior data scientist Lior Seeman, “we had discussions with the shared rides product team, refining the corners and making sure we had all the details right.” Within a few weeks, the team had built a basic, bare-bones prototype of the Express product. They continued to refine the prototype until they had built the Express code base.

September 2017: Simulations and “Trip Parties”

Starting from September 2017, the team began to run simulations mimicking how Express would work in the market as a function of different matching parameters. To do so, the team cloned Uber’s historical POOL requests and simulated the matches that would occur under different Express scenarios, for example with longer or shorter waiting times. “In just two hours per simulation, we could see the details of all Express matches that would have occurred in a two-week period,” said Rahematpura. The team then reviewed these matches together in meetings, called “trip parties.”

These meetings served more than one purpose. Their primary goal was to establish a set of parameters that would dictate Express matching, based on the results of the simulations. The team would focus on simulated trips to evaluate how varying waiting and walking parameters influenced rider experience and efficiency. An equally important goal of the trip parties was to provide a forum for different stakeholders to discuss the experience/efficiency trade-offs. By the end of the trip parties there was broad agreement about the matching parameters for Express. Maximum waiting time was set at two minutes.

Pricing

While they were confident about the waiting and walking parameters they had established, the Express team remained unsure if riders would buy into the service. “This was a completely new product for us,” said Stock. “In meetings, people would say, ‘The magic of Uber is that it’s an on-demand, door-to-door service, and Express is neither.’”

One major uncertainty regarded Express pricing. While the simulations had mimicked actual historical POOL rides, they had not included price points. To assess riders’ willingness to wait and walk at different price points, the company sent a segment of its riders conjoint surveys—a type of questionnaire that measured consumers’ sensitivity to different variables. Based on the surveys’ results, Uber built a calculator that aimed to predict pricing thresholds based on walking and waiting parameters. As Church recalled, “The survey gave us a helpful floor and ceiling for pricing. Most people simply will not wait 15 minutes for a ride, even if prices are far cheaper.”

To help refine its pricing, the company also looked at demand curves previously estimated for the POOL product, which quantified rider sensitivity to price changes. As Church said, “We wanted to launch a product that did not lose money, while still allowing us to offer a lower price.” While pricing would be dynamic, Uber decided that Express would always be at least 20% cheaper than POOL. “But,” said Lee, “if that’s only a \$0.50 difference, that probably isn’t enough to persuade people to take Express and deal with waiting and walking, so in some cases the discounts are deeper.” Thus, Uber also planned to adjust the prices for the POOL product in order to make Express attractive. “We are still determining what the substitution patterns are,” said Snider. “Then we will revise accordingly.”

November 2017: Boston and San Francisco Pilots

By November, the merged team believed that it was ready to release the Express product into test markets. “It was still only half as good as we wanted it to be,” said Church, “but we knew we needed

to get it into the market to see riders' reactions." Thus, in early November, Uber launched Express in select neighborhoods of Boston and San Francisco. The company chose these cities because they were competitive, dense markets. "We had 'home curb advantage' in San Francisco," explained Lee. "If something went wrong, we knew that we could address it fairly quickly." Boston's heavy student population was also a consideration. Launching in just a few neighborhoods would confine the effects of any negative repercussions to smaller geographies, and provide an opportunity to test the algorithm for bugs, since many Uber employees would use the new product in San Francisco.

When riders clicked on the Express product in their Uber app for the first time, a box popped up that read, "Walk a little, Save a lot," and offered a primer on Express. For drivers, the product was quite similar to POOL, so there was less need for driver education. According to Seeman, "The limited launches went well. Nothing was a huge surprise. We learned small things. For example, riders really care about not getting dropped off 100 feet before their door just to make the ride one minute shorter."

In December, Uber expanded the product to all of Boston and San Francisco. Express was well-received in San Francisco. "We saw an increase in volume and cost savings, with no degradation of marketplace metrics around user experience," said Lee. Reception in Boston, however, was more lukewarm. "Weather was certainly a factor influencing adoption," said Seeman. "We launched in the winter, and people likely did not want to walk and wait in the cold."

February 2018: 12-City Synthetic Control Experiment (Launch Experiment)

On February 19, 2018, after seeing that the Express product had not caused any major problems in Boston and San Francisco, Uber launched a synthetic control experiment (launch experiment) in 12 U.S. cities. The company launched Express in six treatment cities and held constant a set of six control cities. In the six treatment cities, riders were made to wait up to two minutes before being matched to a driver.

As was the case in Boston and San Francisco, all of the 12 cities already offered the POOL product prior to the experiment. In preparation for launch, the company ran advertisements and sent emails to all its users in the treatment markets. Following the launch, Uber began monitoring the effects of Express on these markets, as compared with the control cities. As was standard practice, the company placed a five-week freeze on experimental changes to all 12 of these cities to enable data scientists to interpret any market changes as cleanly as possible.

The Wait Time Debate

To evaluate riders' willingness to wait, in mid-February the Express product team launched an experiment in Boston (**Exhibit 10** provides a timeline of all the experiments discussed). By that point, the Express product – with two-minute waiting – had been available in Boston and San Francisco for three months, allowing ample time for the markets to stabilize after the addition of a new product.

To evaluate the effects of longer wait times, the team set up a switchback experiment in Boston. Every 160 minutes, the matching algorithm switched between letting riders wait up to two (control group) and five minutes (treatment group) before being matched to a particular driver. After two weeks, the data scientists analyzed the experiment results. The data indicated that Uber's costs per ride decreased with longer wait times, but there were differences in how rush and non-rush hours were affected. Because the longer wait times resulted in more efficient matches, seating capacity was better utilized, thus making lower prices profitable. (**Exhibit 11** provides a snapshot of the experiment data.)

This was the information that Gilchrist was conveying to his colleagues at the meeting in early March. Gilchrist and his colleagues now needed to decide whether to overrule the 5-week freeze in the launch experiment and increase the wait time from two to five minutes in the six treatment cities. Trivedi emphasized the negative effects of longer wait times on customer experience, while Rahematpura pushed for an immediate increase of waiting times given the economic benefit.

Gilchrist, while sympathetic to the product managers' views, reiterated the importance of continuous data collection after a product launch. "We only get one shot for clean data collection," he said. "The data from the launch experiment will inform all future product improvements to Express." Seeman agreed with Gilchrist, adding, "Say we go ahead and increase wait times to five minutes across all six treatment cities midway through the launch experiment. If, in three months, we find that Express is performing poorly, we won't be able to say with certainty whether this is due to defects with Express or because people in those markets are reacting poorly to the increased wait times." Rahematpura, who had been running numbers, interjected: "According to some back-of-the-envelope calculations, by not increasing wait times now, we stand to lose \$1.6 million in the six treatment cities. That might outweigh the data collection concerns." Everyone looked to Stock for a decision.

Exhibit 1 Selected Data, Uber vs. Competitors, 2018

Company	Founded	Primary Market	# of Trips in 2017	# of Drivers	# of Riders	Last Valuation (in \$bn)	\$ Raised
Uber	2009	U.S.	4 billion	3 million	75 million	62	\$21 billion
Lyft	2012	U.S.	375 million	1.4 million	23 million	15	\$4.1 billion
Didi Chuxing	2012	China	7.4 billion	21 million	450 million	56	~\$19 billion
Ola	2010	India	6 million/wk (2016)	1 million	-----	7	~\$3 billion
Grab*	2012	Singapore	1 billion total (as of 2017)	2 million	68+ million	6+	\$4.1 billion
Go-Jek	2010	Indonesia	-----	900,000	15 million/week	~5	~\$2.1 billion
Taxify	2013	Europe	-----	500,000	10 million	~1	\$177 million
Yandex**	2011	Russia	285 million	-----	-----	3.7 (w/ Uber co-ownership)	
Careem	2012	Middle East	-----	560,000	14 million	1.2	\$572 million

Source: Casewriter, compiled from: Uber, "Company Info," 2018, <https://ubr.to/2xIGJJK>; Jillian D'Onfro and Josh Lipton, "Uber Posts Big Sales Jump in First Quarter and Boosts Valuation to \$62 Billion," CNBC, May 23, 2018, <https://cnb.cx/2LsID7N>; Megan Rose Dickey and Ingrid Lunden, "Uber's Raising up to \$600M in a Secondary Round at \$62B Valuation, Q1 Sales Grew to \$2.5B," TechCrunch, May 23, 2018, <https://tcn.ch/2KJvYwv>; "Lyft Raises New Capital and Continues Momentum," Lyft (blog), June 27, 2018, <https://lft.to/2IKNzCG>; "Our 2017 in Review," Lyft (blog), January 16, 2018, <https://lft.to/2Nk0mzq>; Dara Kerr, Lyft Grows Gangbusters in 2017, Bringing Competition to Uber," Cnet, January 16, 2018, <https://cnet.co/2lQkgpf>; Xiaochun Zhao, "Losing \$300M in 2017, Didi Chuxing Wants to Turn a Profit in 2018 amid Fierce Competition," Kr Asia, April 3, 2018, <https://bit.ly/2z64UWY>; Johana Bhuiyan, "China Ride-Hail Giant Didi Chuxing Has Raised \$4 Billion," Recode, December 20, 2017, <https://bit.ly/2BJddpA>; Xinhua, "Didi Completes 7.43 Bln Rides in 2017," January 8, 2018, <https://bit.ly/2z61UtS>; "Ola," Crunchbase, <https://bit.ly/2MGmuCP>; Arjun Kharpal, "Uber's Biggest Rival in India Just Got \$1.1 Billion from Tencent, SoftBank, Valuing Company around \$7 Billion," CNBC, October 11, 2017, <https://cnb.cx/2hAS9rJ>; Anaya Bhattacharya, "As Uber Sputters, Ola Is Really Stepping on the Gas in India," Quartz, February 15, 2018, <https://bit.ly/2EvFQHq>; Sayan Chakraborty, "Ola, Uber See Rides Rise Fourfold in 2016: Report," LiveMint, February 17, 2017, <https://bit.ly/2Krn4YJ>; Swashwati Shankar, "Undeterred by High Attrition Rate, Ola and Uber Banking on Drivers in Their 20s," The Economic Times, June 1, 2017, <https://bit.ly/2tOGHzn>; "Billion Dollar Unicorns: Grab Becomes the Most Valuable Startup in Southeast Asia," One Million by One Million Blog, December 8, 2017, <https://bit.ly/2z5RDxI>; "You're One in a Billion," Grab (blog), November 6, 2017, <https://bit.ly/2j5Mu0Z>; Jon Russell, "Go-Jek Buys Three Startups to Advance Its Mobile Payment Business," TechCrunch, December 15, 2017, <https://tcn.ch/2BLuAlt>; "GO-JEK," Crunchbase, <https://bit.ly/2rk0yag>; Anshuman Daga, "Indonesia's Go-Jek Raises \$1.5 Billion as Ride-Hailing Market Heats Up: Sources," Reuters, February 26, 2018, <https://reut.rs/2F5QGUC>; Crunchbase, "Taxify," not dated, <https://bit.ly/2NkqiL9>; Jon Russell, "Uber's European Rival Taxify Raises \$175M Led by Daimler at a \$1B Valuation," TechCrunch, May 30, 2018, <https://tcn.ch/2J09giY>; Frank DiPietro, "Yandex.Taxi Is Just One Example of the Sprawling Empire Yandex Is Building in Russia," The Motley Fool, July 26, 2017, <https://bit.ly/2KFaf9a>; Ingrid Lunden, "Uber Rival Careem Closes \$500M Raise at \$1B+Valuation as Daimler Steps In," TechCrunch, June 15, 2017, <https://tcn.ch/2t4kwml>; "Careem," Crunchbase, <https://www.crunchbase.com/organization/careem#section-locked-charts>; Megan Rose Dickey, "Ride-Hailing App Careem Reveals Data Breach Affecting 14 Million People," TechCrunch, April 23, 2018, <https://techcrunch.com/2018/04/23/careem-data-breach/>; all accessed June 2018.


Note: * These numbers are from late 2017, prior to Uber taking an ownership stake in Grab.

** Yandex is now combined with Uber in some regions. These numbers are from June 2017, just prior to the Uber deal.


Exhibit 2 Uber Product Types, 2018

	Type	Launched	Description
Carpool Options	Express	2018	Matches riders going in the same direction; Requires riders to walk to/from their pick-up and drop-off points and wait a few minutes to be matched
	Pool	2014	Matches riders going in the same direction; Offers door-to-door rides with no walking or waiting
Economy Options	UberX	2012	Provides private, affordable rides for 1 to 4 people; Uber's core economy product
	UberXL	2014	Provides private, affordable rides for up to 6 people
	UberSelect	2015	Provides private rides for 1 to 4 people with a driver who has been consistently highly rated
Premium Options	UberBLACK	2010	Uber's original ride option; Provides private rides in high-end black cars with professional drivers for 1 to 4 people
	UberSUV	2015	Provides private rides in luxury SUVs for up to 6 people
	UberLUX	2014	Uber's most luxurious option; Provides private rides in high-end cars for 1 to 4 people

Source: Casewriter, compiled from: Uber, "Ride," 2018, <https://www.uber.com/ride/>, accessed June 2018.

Exhibit 3 Uber's Rider-Facing App, 2016

Source: Company documents.

Exhibit 4 Uber's Driver-Facing App, 2015

Source: Company documents.

Exhibit 5 Driver's Earnings Breakdown in Uber App

In the app

We give you complete clarity on how your earnings are calculated. For more information tap "Fare Details" at the bottom of the receipt.

1. Trip start and end locations.
2. Your trip earnings for the trip.
3. Duration of this trip. This is used to calculate your time-based earnings (see #10).
4. Trip distance. This is used to calculate your distance-based earnings (see #9).
5. Vehicle type. Your earnings may vary by vehicle type.
6. Time you accepted the ride request.
7. Date you accepted the ride request.
8. The base fare you earned for this trip.
9. The distance portion of your earnings for this trip. Your per-mile rate is multiplied by the actual trip mileage.
10. The time portion of your earnings for this trip. Your per-minute rate is multiplied by the actual trip duration.
11. Tap to see further details for this trip.


Source: Uber, "Trip Details," 2018, <https://www.uber.com/drive/resources/earnings-trip-details/>, accessed June 2018.

Exhibit 6 Median Hourly Earnings per UberX Driver, by Number of Hours Worked Weekly, 2014

	1 to 15 hours/week		16 to 34 hours/week		35 to 49 hours/week		Over 50 hours/week	
	% of drivers	Earnings/hr	% of drivers	Earnings/hr	% of drivers	Earnings/hr	% of drivers	Earnings/hr
Bos	58%	\$19.25	30%	\$20.41	9%	\$20.78	4%	\$20.48
Chi	56%	\$15.60	31%	\$16.12	9%	\$16.21	4%	\$16.03
DC	53%	\$16.61	31%	\$17.46	10%	\$17.70	6%	\$17.41
LA	59%	\$16.37	29%	\$17.07	8%	\$17.07	4%	\$16.97
NY	42%	\$26.03	35%	\$28.47	16%	\$29.65	7%	\$29.61
SF	53%	\$23.74	34%	\$25.51	10%	\$25.36	3%	\$25.36

Source: Jonathan Hall and Alan Krueger, "An Analysis of the Labor Market for Uber's Driver-Partners in the United States," *Uber*, January 22, 2015, p. 18, https://s3.amazonaws.com/uber-static/comms/PDF/Uber_Driver-Partners_Hall_Krueger_2015.pdf, accessed June 2018.

Note: Bos = Boston; Chi = Chicago; DC = Washington, DC; LA = Los Angeles; NY = New York City; SF = San Francisco.


Exhibit 7 Uber's Organizational Chart, March 2018

Source: Company documents.

Exhibit 8 Ratio of Data Scientists Employees at Major Tech Companies, July 2018

	Approximate Total Employees	Share of Data Scientists
Facebook	30,000	4.5%
Airbnb	8,000	2.7%
Uber	16,000	2.4%
Netflix	5,000	1.8%
Instacart	2,000	1.0%
Lyft	14,000	0.9%
HomeAway	6,000	0.7%
Google	83,000	0.7%
Amazon	195,000	0.5%
Apple	155,000	0.5%
Upwork	33,000	0.3%
Rover	2,500	0.1%

Source: Casewriter, compiled from LinkedIn, accessed June 2018.

Exhibit 9 Uber's Express POOL User Interface

Source: Company documents.

Note: The circled area in the second picture indicates the “corner” selected for this particular area. Each of the pins is a potential pick-up point.

Exhibit 10 Timeline of Events

Date	Event
2014	Uber launches Uber POOL
2017	Uber begins re-thinking shared rides strategy to increase profitability
September 2017	Uber begins simulations on the Express POOL concept
November 2017	Uber launches pilots of Express POOL in San Francisco and Boston
February 19, 2018	Uber launches a 5-week-long synthetic control experiment to test Express POOL. The experiment has six treatment cities (Denver, Los Angeles, Miami, Philadelphia, San Diego, and Washington DC) with Express POOL and 2-minute wait time, and six control cities.
March 6, 2018	Results of the switchback experiment in Boston are available. The switchback experiment compared 2- versus 5-minute wait times.

Source: Casewriter.

Exhibit 11 Snapshots of the Data Dictionary and Data from the Boston Switchback Experiment

Data Dictionary		
Variable	Type	Definition
city_id	String	Location where the experiment took place. In the data it is always equal to "Boston."
period_start	Date	Start date and time for the 160-minute time period of the current observation.
wait_time	String	This variable takes on two possible values: "2 mins" if the matching algorithm let riders wait up to 2 minutes during the current time period; "5 mins" if the matching algorithm let riders wait up to 5 minutes during the current time period.
treat	Boolean	This variable takes on two possible values: "TRUE" if wait_time equals "5 mins"; "FALSE" if wait_time equals "2 mins".
commute	Boolean	This variable takes on two possible values: "TRUE" if the time period happens during rush hours (7-9:40AM or 3-5:40PM), "FALSE" otherwise.
trips_pool	Numeric	Total number of POOL trips completed in the current time period. Each matched ride request is a separate trip.
trips_express	Numeric	Total number of Express POOL trips completed in the current time period. Each matched ride request is a separate trip.
rider_cancellations	Numeric	Total number of requested trips that were cancelled by the rider in the current time period.
total_driver_payout	Numeric	Total dollars paid to drivers for trips completed in the current time period. This is equal to Uber's total costs for matching trips in the current time period.
total_matches	Numeric	Number of completed trips during the current time period that were paired with at least another rider for part of the trip. Each matched ride request is a separate trip, so two separate riders matched together would count as two matches.
total_double_matches	Numeric	Number of completed trips during the current time period that were paired with at least two other riders for part of the trip. Each matched ride request is a separate trip, so three separate riders matched together would count as three double matches.

Snapshot of Boston Switchback Dataset

city_id	period_start	wait_time	treat	commute	trips_pool	trips_express	rider_cancellations	total_driver_payout	total_matches	total_double_matches
Boston	2/19/2018 7:00	2 mins	FALSE	TRUE	1415	3245	256	34458.41163	3372	1476
Boston	2/19/2018 9:40	5 mins	TRUE	FALSE	1461	2363	203	29764.34982	2288	1275
Boston	2/19/2018 12:20	2 mins	FALSE	FALSE	1362	2184	118	27437.36736	2283	962
Boston	2/19/2018 15:00	5 mins	TRUE	TRUE	1984	3584	355	44995.45299	4035	2021
Boston	2/19/2018 17:40	2 mins	FALSE	FALSE	1371	2580	181	27583.9553	2200	979
Boston	2/19/2018 20:20	5 mins	TRUE	FALSE	1401	2022	135	23888.11085	2066	1062
Boston	2/19/2018 23:00	2 mins	FALSE	FALSE	1216	2543	103	29642.90567	2600	1406
Boston	2/20/2018 1:40	5 mins	TRUE	FALSE	1691	2018	150	25794.86992	1918	1281
Boston	2/20/2018 4:20	2 mins	FALSE	FALSE	1248	2481	131	23238.94629	2623	1059
Boston	2/20/2018 7:00	5 mins	TRUE	TRUE	1815	2539	284	34047.4739	2624	1565
Boston	2/20/2018 9:40	2 mins	FALSE	FALSE	1594	2773	166	28053.3648	2723	855
Boston	2/20/2018 12:20	5 mins	TRUE	FALSE	1629	2380	185	25964.4189	2477	1322
Boston	2/20/2018 15:00	2 mins	FALSE	TRUE	1640	3290	236	39912.40151	3777	1969
Boston	2/20/2018 17:40	5 mins	TRUE	FALSE	1173	1891	144	22029.73431	1893	998
Boston	2/20/2018 20:20	2 mins	FALSE	FALSE	1853	2305	159	33299.09273	2234	1145
Boston	2/20/2018 23:00	5 mins	TRUE	FALSE	1145	2419	182	25062.18271	2214	1104
Boston	2/21/2018 1:40	2 mins	FALSE	FALSE	1490	1949	125	25419.06091	1904	1120
Boston	2/21/2018 4:20	5 mins	TRUE	FALSE	1560	2309	180	25104.44488	2125	990
Boston	2/21/2018 7:00	2 mins	FALSE	TRUE	1916	3469	245	44871.10585	3575	1809
Boston	2/21/2018 9:40	5 mins	TRUE	FALSE	1664	2320	183	27626.13599	2344	1309
Boston	2/21/2018 12:20	2 mins	FALSE	FALSE	1039	3188	168	29259.93547	3045	1607

Source: Casewriters.

Note: This snapshot is provided as a case supplement. Note that the instructor may choose to assign the case without the data supplement.

Endnotes

¹ Moon, "Uber: Changing the Way the World Moves," p. 2.

² Alexia Tsotsis, "Uber Opens Up Platform to Non-Limo Vehicles with 'Uber X,' Service Will Be 35% Less Expensive," *TechCrunch*, July 2, 2012, <https://techcrunch.com/2012/07/01/uber-opens-up-platform-to-non-limo-vehicles-with-uber-x-service-will-be-35-less-expensive/>, accessed June 2018.

³ Uber, "Company Info," *Uber*, 2018, <https://www.uber.com/newsroom/company-info/>, accessed June 2018.

⁴ Uber, "Company Info."

⁵ Jillian D'Onfro and Josh Lipton, "Uber Posts Big Sales Jump in First Quarter and Boosts Valuation to \$62 Billion," *CNBC*, May 23, 2018, <https://www.cnbc.com/2018/05/23/uber-q1-financial-data-increased-sales-valuation-with-new-tender-offer.html>, accessed June 2018.

⁶ As reported by: Megan Rose Dickey and Ingrid Lunden, "Uber's Raising up to \$600M in a Secondary Round at \$62B Valuation, Q1 Sales Grew to \$2.5B," *TechCrunch*, May 23, 2018, <https://techcrunch.com/2018/05/23/uber-q1-2018/>, accessed June 2018.

⁷ Moon, "Uber: Changing the Way the World Moves," p. 3.

⁸ Miranda Katz, "This App Lets Drivers Juggle Competing Uber and Lyft Rides," *Wired*, February 15, 2018, <https://www.wired.com/story/this-app-lets-drivers-juggle-competing-uber-and-lyft-rides/>, accessed June 2018.

⁹ Leslie Hook, "Can Uber Ever Make Money?" *Financial Times*, June 23, 2017, <https://www.ft.com/content/09278d4e-579a-11e7-80b6-9bfa4c1f83d2>, accessed June 2018.

¹⁰ Uber, "Ride," 2018, <https://www.uber.com/ride/>, accessed June 2018.

¹¹ Moon, "Uber: Changing the Way the World Moves," p. 4.

¹² Stephen Antczak, "11 Things to Know before Becoming an Uber or Lyft Driver," *Forbes*, April 23, 2017, <https://www.forbes.com/sites/nextavenue/2017/04/23/11-things-to-know-before-becoming-an-uber-or-lyft-driver/#503d71fd6579>, accessed June 2018.

¹³ Noah Smith, "Uber Better Not Be the Future of Work," *Bloomberg*, March 8, 2018, <https://www.bloomberg.com/view/articles/2018-03-08/uber-drivers-earn-pay-that-s-just-above-the-poverty-line>, accessed June 2018.

¹⁴ Jonathan Hall and Alan Krueger, "An Analysis of the Labor Market for Uber's Driver-Partners in the United States," *Uber*, January 22, 2015, p. 18, https://s3.amazonaws.com/uber-static/comms/PDF/Uber_Driver-Partners_Hall_Kreuger_2015.pdf, accessed 2018.

¹⁵ Uber, "How Surge Pricing Works," 2018, <https://www.uber.com/drive/partner-app/how-surge-works/>, accessed June 2018.

¹⁶ Eric Newcomer, "Uber, Daimler Strike Partnership for Self-Driving Vehicles," *Bloomberg*, January 31, 2017, <https://www.bloomberg.com/news/articles/2017-01-31/uber-daimler-strike-partnership-for-self-driving-vehicles>, accessed 2018.

¹⁷ Wikipedia, "Synthetic Control Method," https://en.wikipedia.org/wiki/Synthetic_control_method, accessed June 2018.

¹⁸ Johana Bhuiyan, "Uber Drivers Will Get a Flat Fee for Every New Pick Up on Rides," *Recode*, September 26, 2017, <https://www.recode.net/2017/9/26/16365904/uber-drivers-pool-driver-improvement>, accessed June 2018.

¹⁹ "From Zero to Seventy (Billion)," *The Economist*, September 3, 2016, <https://www.economist.com/briefing/2016/09/03/from-zero-to-seventy-billion>, accessed June 2018.

²⁰ Johana Bhuiyan, "Uber's new 'Express Pool' is all about getting more riders to share rides," *Recode*, February 21, 2018, <https://www.vox.com/2018/2/21/17032598/uber-express-pool-transit-bus-cheaper>, accessed February 2020.