



INFORMS Transactions on Education

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Case—Calyber: A Ridesharing Game

Yifan Shen, Chiwei Yan, Julia Yan

To cite this article:

Yifan Shen, Chiwei Yan, Julia Yan (2026) Case—Calyber: A Ridesharing Game. *INFORMS Transactions on Education* 26(3):226-232. <https://doi.org/10.1287/ited.2025.0163cs>

This work is licensed under a Creative Commons Attribution 4.0 International License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as “*INFORMS Transactions on Education*.” Copyright © 2026 The Author(s). <https://doi.org/10.1287/ited.2025.0163cs>, used under a Creative Commons Attribution License: <https://creativecommons.org/licenses/by/4.0/>.”

Copyright © 2026 The Author(s)

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes. For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>




Case

Calyber: A Ridesharing Game

 Yifan Shen,^a Chiwei Yan,^{b,*} Julia Yan^c

^aDepartment of Industrial and Systems Engineering, University of Washington, Seattle, Washington 98195; ^bDepartment of Industrial Engineering and Operations Research, University of California, Berkeley, California 94720; ^cSauder School of Business, University of British Columbia, Vancouver, British Columbia V6T 1Z2, Canada

*Corresponding author

Contact: syf1996@uw.edu,  <https://orcid.org/0009-0003-0758-4707> (YS); chiwei@berkeley.edu,  <https://orcid.org/0000-0001-5770-2079> (CY); julia.yan@sauder.ubc.ca,  <https://orcid.org/0000-0001-6263-2589> (JY)

Received: August 8, 2025


Revised: November 26, 2025;
January 13, 2026

Accepted: January 28, 2026

Published Online in Articles in Advance:
March 18, 2026

<https://doi.org/10.1287/ited.2025.0163cs>

Copyright: © 2026 The Author(s)

 **Open Access Statement:** This work is licensed under a Creative Commons Attribution 4.0 International License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as “*INFORMS Transactions on Education*. Copyright © 2026 The Author(s). <https://doi.org/10.1287/ited.2025.0163cs>, used under a Creative Commons Attribution License: <https://creativecommons.org/licenses/by/4.0/>.”

Funding: This work was supported by the National Science Foundation [Grant 2517861] and the Natural Sciences and Engineering Research Council of Canada [Grant RGPIN-2022-03524].

Keywords: ridesharing • pricing • matching • dynamic stochastic optimization

1. Introduction

Eliot George is the CEO of Calyber, a ridesharing company based in Chicago, Illinois. When he founded Calyber several years ago, Eliot was driven by a firm belief: to truly succeed, Calyber must offer shared rides—where each driver transports multiple compatible passengers in a carpool-like service. Having grown up in Los Angeles, Eliot was all too familiar with the frustrations of traffic congestion. This experience deeply shaped his vision for Calyber. He was fond of saying that if the company failed to achieve a high level of ride sharing, it would be no different from a traditional taxi service—just one with a slick mobile app.

But offering a profitable shared rides service had turned out to be more difficult than he had once imagined. Downtown, Calyber kept prices low to attract customers. But prices were so low that even with most riders matching with compatible corridors, Calyber was unable to recoup its dispatching costs. Outside of downtown, prices were higher, but this caused separate issues. Most riders who booked a shared ride were never matched with compatible corridors; meanwhile, those who did share were often sent on long, costly routes picking up corridors all over the city. This had turned Calyber into the butt of jokes on social media, as vocal customers complained about irritating corridors and long detours. (Customers with positive experiences rarely posted on social media, Eliot noted with some chagrin.)

Faced with these challenges, Eliot decided to discuss a way forward with Calyber’s chief scientist, Thea Bruckner. He hoped that together, they could design a shared rides product that could achieve high ridership *and* profitability at the same time.

2. Background

A ridesharing platform is a transportation service provider that connects customers (*riders*) with drivers who can take them from their origins to their desired destinations, usually via a mobile app.

Lyft was founded in 2007 by Logan Green and John Zimmer as an intercity carpool service between college campuses before shifting in 2013 to short-distance rides (Sherman et al. 2019). Uber was founded in 2009 by Travis Kalanick and Garrett Camp after they had difficulty hailing a taxi in Paris and sought to make the experience as simple as pushing a button on a smartphone (Uber 2025).

Initially, both companies focused on solo rides, similar to a taxi service. This changed in 2014 when both companies (within a day of each other!) launched *shared* rides in San Francisco (Anand 2017). The products, dubbed UberPool and Lyft Line, offered riders lower prices in exchange for potentially sharing the ride with other riders in a carpool-like service. The idea was that multiple riders with compatible origins and destinations could be served concurrently (*matched*) with the same driver, reducing vehicle miles and improving efficiency. As Kalanick put it, “We saw

a lot of people pushing the same button at the same time going essentially to the same place. And so we started thinking about, well, how do we make those two trips and turn them into one?" (Kalanick 2016). The platform would benefit by serving more riders at lower cost, riders would benefit if lower costs translated into prices low enough to make them willing to share, and Zimmer even claimed that vehicle miles would be reduced so much as to enable the "end of traffic" (Zimmer and Green 2017).

Within a few years, the new shared rides products became core offerings, comprising 20%–30% of all rides (Lunden 2016, Lyft 2018). In 2018, Uber doubled down on shared rides and launched Uber Express Pool, a low-cost variant of UberPool that charged riders low fares in exchange for walking toward meeting points with other riders (Uber 2018). At the same time, Lyft announced its ambition for 50% of its rides to be shared by the end of 2020 (Lyft 2018).

However, operational challenges have prevented shared rides from living up to their lofty aspirations. After the launch in 2014, both companies sought to "win" riders in San Francisco, chiefly by keeping prices low—Uber even offered a \$5 flat fare.¹ But at the same time, Uber's algorithms struggled to match riders: the *match rate*, or the fraction of riders who get matched with another rider, was reportedly only 8%. The combined effect was that Uber was essentially offering ultracheap solo rides, a practice that lost \$1 million a week and that was subsequently described as "bleeding cash subsidizing rides [without] a plan for tomorrow." Over the next few months, Uber improved its matching algorithms and achieved a match rate of almost 25%.² But this improvement was still not enough to break even. Attempting to cut its losses, Uber raised prices, which then tanked ridership by 22% (Anand 2017).

Compounding these difficulties, the COVID-19 pandemic compelled both companies to shut down shared rides entirely to aid physical distancing. Recognizing the difficulty of bringing customers back to shared rides, Lyft's and Uber's relaunches in 2021 and 2022 came with a strong emphasis on affordability for inflation-conscious customers. Lyft included an option to reserve shared rides early for an extra discount, promising "the most affordable pricing available" (Lyft 2021). Meanwhile, Uber launched a pricing policy that would offer further discounts to riders if they were matched with coriders. As Uber CEO Dara Khosrowshahi explained, "You get a deeper discount if you get matched up with somebody, so it's in your interest to get matched up... taking that sharing to the next level and the discounts to the next level" (Uber 2022, Morning Brew Daily 2024). However, these launches struggled to find traction

(Desai 2022), and in the face of financial difficulties, Lyft stopped offering shared rides in 2023 (Davalos 2023). Announcing the decision, Lyft CEO David Risher explained, "The problem with shared trips is that they take people out of their way. At some point you have to pay attention to what your customers want."

The up-and-down history of shared rides highlights the following operational challenges. First, matching requires multiple riders with compatible origins, destinations, and request times, which can be rare unless demand density is sufficiently high. Second, low prices are needed to attract riders and achieve enough density for efficient matching, but can create losses for the platform if riders cannot be matched effectively. In trying to strike a balance, platforms may set prices that are both too high to attract riders *and* too low to comfortably sustain. These operational challenges fuel skepticism that shared rides can be operated viably.

At the same time, these challenges present an opportunity for any ridesharing company that can solve the puzzle of a viable shared rides product. In 2025, Uber promoted new features called route share and price lock, aimed at making regular trips (e.g., to and from work or major events) more affordable to combat continued inflation and improve accessibility (Jackson 2024, Kansal 2025, The Verge 2025). Lyft also started experimenting with shared rides again, reintroducing them at airports with high demand density (Lung 2025). Moreover, beyond the high-density corridors formed by commutes, major events, and airport trips, low-density neighborhoods are an underserved market primed for innovation.

3. Calyber's Operations

Calyber is based in Chicago, Illinois, and focuses entirely on offering shared rides (see Figure 3 for its service area). Its platform is used by about 35,000 customers per day. The service is *on demand*, which means that riders can expect to be served quickly. Calyber might let riders wait on the platform to give time for compatible coriders to arrive, but this waiting time is kept short (e.g., at most a few minutes) to ensure customer satisfaction. Calyber also guarantees that each rider will be matched with at most one other corider.

To request a ride, a rider can open Calyber's mobile app, enter the origin and destination, and receive a price quote (this is called a rider *session*). If the rider is willing to pay the quoted price (the rider decides chiefly based on the price per unit distance in dollars per mile), then the rider *requests* a ride. Otherwise, the rider leaves the platform. The rider's willingness to pay is unknown to Calyber.

After a rider makes a request, Calyber makes a dispatch decision. If there is another compatible rider

currently waiting on the platform, Calyber can match the two riders and dispatch them immediately on a shared ride. Otherwise, Calyber can let the rider wait on the platform for a short period. The rider has limited patience, so she is willing to wait only for a limited sojourn time, which is also unknown to Calyber. If the rider's patience expires, then, as the rider is about to exit the platform, Calyber can intervene and dispatch the rider immediately in a solo ride. See Figure 4 for example routes corresponding to shared and solo requests. Whether the dispatch is shared or solo, Calyber pays the driver \$0.70/mile to perform the dispatch. As Calyber is known to treat drivers well, there are sufficient drivers to ensure that a driver is available for immediate dispatch.

A critical feature of shared rides is that the average cost per request depends on demand density. When demand density is high (i.e., more riders request similar trips at similar times), it is easier to find matches with low detours. Then, a single driver can serve multiple requests with only slightly longer routes, reducing the average cost per request. By contrast, when demand is sparse, it is more difficult to find matches. Consequently, the platform dispatches more requests along routes with longer detours or even solo, increasing the average cost per request.

4. A New Engine for Pricing and Matching

Calyber's pricing and matching engine had been designed by one of its founding engineers, Nick Bailey. Unfortunately, Nick was no longer at Calyber, and the engine was a black box to the rest of the engineering team (hastily written code from the early years of a startup is not known for tidiness).

Eliot and Thea agreed that Calyber could not address its operating issues without overhauling its pricing and matching engine, and they decided to discuss the best way to move forward.

Eliot: The shared rides product isn't working for us—we're bleeding cash.³ I think we need to raise prices. And it seems that the engine is pretty complicated and has caused your team a lot of trouble. What if we simplify and charge everyone a flat price, like \$0.80/mile?

Thea: I agree that we should raise prices, although we should be careful not to drive away riders. As for the flat price, I see the appeal of simplicity, but I don't think charging different riders different prices is the problem. Actually, charging different riders different prices can be a *good* thing. Downtown riders are matched more than the riders in the outer neighborhoods, so they cost us less per mile. Shouldn't they get lower prices?⁴

Eliot: What do you think is the problem with the engine, then?

Thea: To me, the problem with the engine is that its pricing decisions are completely black box. It's

impossible to troubleshoot. Remember a couple weeks ago, when a bug made it into production, and we tanked the price of airport rides for an entire weekend? It took my team days to track down the issue.

Eliot: Yeah, we need to make sure that this kind of mistake never happens again. That's one reason why I proposed a flat price for everyone. But you made a good point just now—different riders cost us different amounts per mile. So, what about pricing at cost? It seems like we could at least differentiate downtown and outer riders.

Thea: We can model costs at an even more granular level than downtown versus outer. We record every route that our drivers take, including all rider pickups and drop-offs. We can easily use the routes to calculate the average cost by origin and destination.⁵

Eliot: Great. What do you think riders will think of our pricing strategy?

Thea: That's another piece we should figure out. Every time a rider uses the mobile app, we record the origin, destination, price, and whether the rider makes a request.⁶ We could use this historical data to learn how riders respond to our prices and build a demand model. Combining the demand model with the cost model, we can set prices that maximize profit.

Eliot: This sounds good to me. So, pricing is resolved then. What will we do about matching, then?

Thea: There's so much to do on the pricing side that I haven't thought about matching yet. I'm hoping that we'll have enough customers that we won't have to worry so much about matching. Then we could do something simple. Say a rider has just made a request on the platform. If there are no compatible riders waiting for dispatch, then let the arriving rider wait. But if there are compatible riders waiting on the platform, then match the arriving rider with the waiting rider that generates the highest level of cost saving in their shared route.

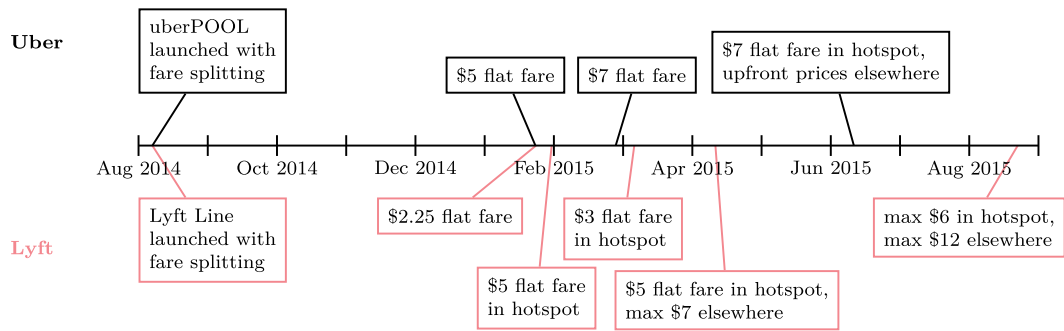
Eliot: That sounds reasonable to me. Why would we ever match any other way?

Thea: Well, it's possible that if we let the arriving rider wait even longer before we dispatch, then other compatible riders might arrive, and we might be able to find better matches. But by waiting, we also risk riders trying to cancel their rides if we make them wait too long.

Eliot: Waiting can be valuable, then. You talked about using historical data to understand how riders respond to prices. Can we also use it to understand how they feel about waiting?

Thea: Maybe ... We do record riders' waiting time from arrival to dispatch in our database.⁷ But we're missing some information—for riders who are matched with other riders, we don't know how long they would

Figure 1. Timeline of Pricing Policies Tested by Uber and Lyft in the First Year



Source. Figure reproduced from Anand (2017).

have waited otherwise. I think this would be difficult for us to do right—we already have a lot to do on the pricing side, and my team is stretched pretty thin. Why don't we focus on getting a simpler prototype up and running and work on this waiting idea later?

Eliot: I see. I do like the idea of giving customers fast service, so let's go with your first idea.

Thea: Sounds great to me. We'll get started right away on implementing the new engine, and we'll kick off an experiment to track profit, match rate, and throughput.

Eliot: Sounds like a plan.

Thea's team implemented the new engine and piloted it in an experiment.⁸ Thanks to an increase in prices, profit increased—instead of losing \$3.76 per minute, Calyber was now making \$2.47 per minute in profit. Thea was happy that the new engine was profitable but suspected that there was still room for improvement. Throughput had dropped by 76%, from 16.82 to 4.02 riders per minute. As a result, the match rate dropped from 80% to 59%.

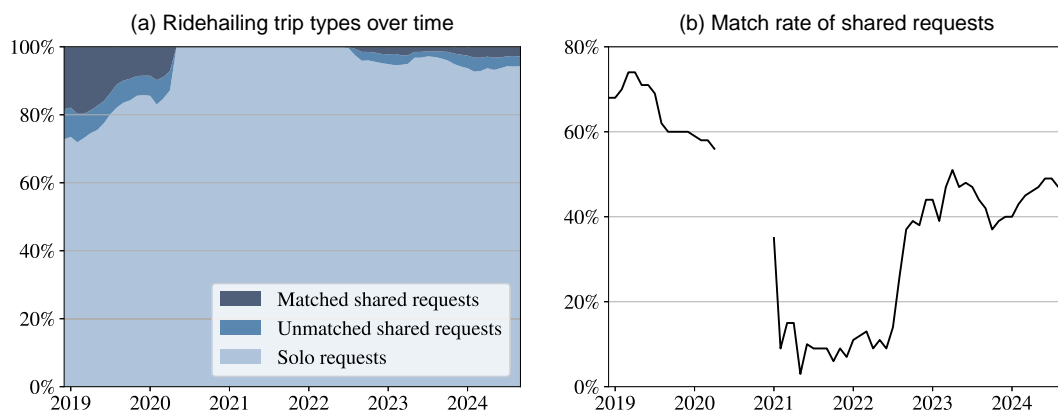
Thea wondered if they had raised prices too quickly. Should they back off the higher prices? If so, were the

demand and cost models accurate enough to quantify the impact of such a pricing change? And could they do anything else to increase demand beyond simply adjusting prices up and down? Finally, if they did have to keep prices at the higher levels, could they improve match rates with the existing throughput? She called a team meeting to discuss these questions and chart a path forward.

Endnotes

- ¹ See Figure 1 for a timeline of the pricing policies that Uber and Lyft experimented with over the first year launching shared rides.
- ² See Figure 2 for the match rates of Uber and Lyft's shared rides in Chicago over recent years.
- ³ See Table 1 ("Nick's engine") for Calyber's key performance indicators (KPIs) under the current engine.
- ⁴ See Figure 5 for demand, cost, and price (as priced by Nick's engine) across different neighborhoods of Chicago.
- ⁵ See Table 2 for a snapshot of Calyber's routing database.
- ⁶ See Table 3 for a snapshot of Calyber's rider database.
- ⁷ See Table 3 for a snapshot of Calyber's rider database, "Waiting Time" column.
- ⁸ See Table 1 ("Thea's engine") for Calyber's KPIs under the new engine.

Figure 2. Ridesharing Trip Types and Match Rates in Chicago over Time



Source. Figure reproduced from data in Schneider (2025).


Notes. From April 2020 to April 2021, shared rides were suspended across major ride-hailing platforms because of COVID-19. Lyft relaunched shared rides in May 2021; Uber did the same in June 2022.

Table 1. KPIs for Nick’s and Thea’s Pricing and Matching Engines by Neighborhood Type

Key performance indicator	Nick’s engine			Thea’s engine		
	Downtown	Outer	Overall	Downtown	Outer	Overall
Profit (\$/min)	−1.96	−1.79	−3.76	0.76	1.72	2.47
Throughput (#/min)	6.92	9.90	16.82	1.71	2.31	4.02
Match rate	0.86	0.76	0.80	0.67	0.53	0.59
Detour rate	0.14	0.14	0.14	0.14	0.14	0.14
Average quoted price (\$/mile)	0.48	0.60	0.56	0.80	0.84	0.83
Average cost per rider (\$/mile)	0.59	0.63	0.62	0.61	0.64	0.63

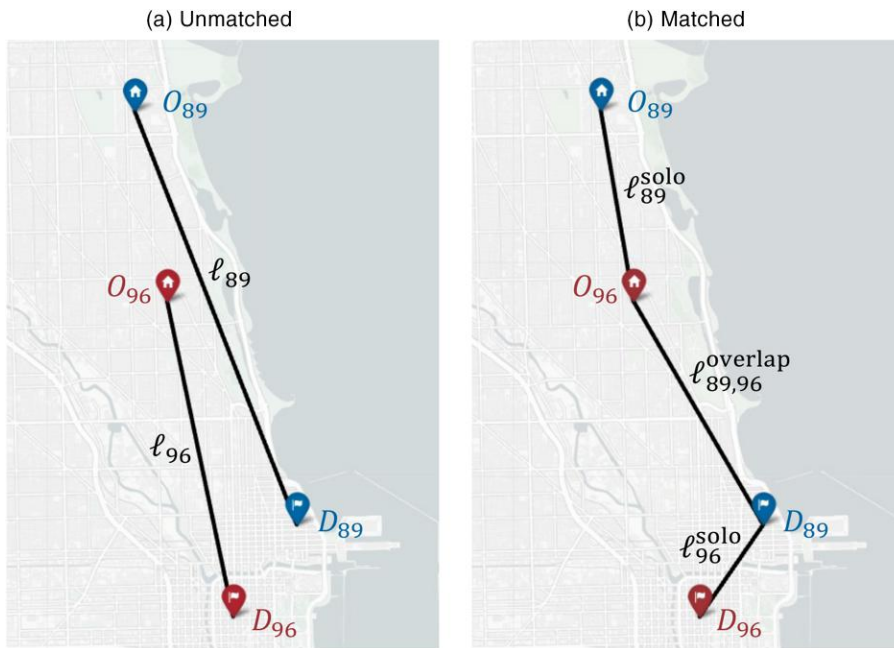
Notes. Throughput is the number of requests served per minute. Match rate is the fraction of requests that are matched with another request. For the calculation of Detour rate and Average cost per rider, see the caption of Table 2.

Figure 3. Downtown and Outer Neighborhoods of Chicago



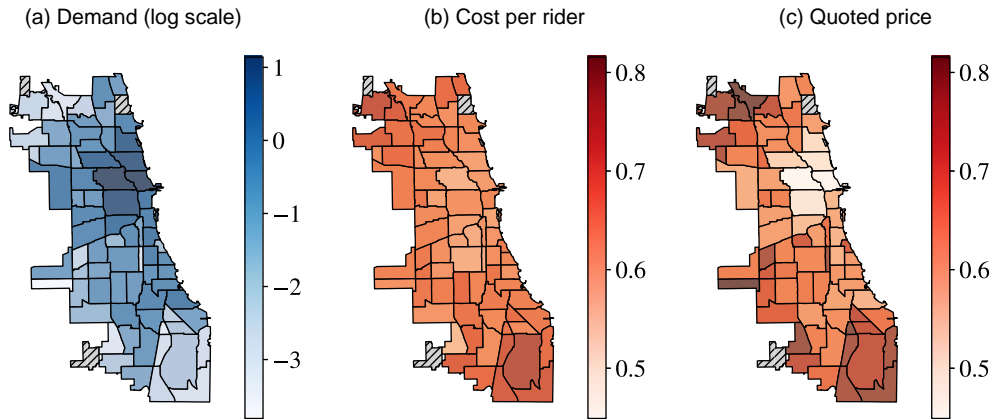
Notes. The downtown neighborhoods cover 20.36 square miles and have demand of 9.74 riders per minute. The outer neighborhoods cover 197.08 square miles and have demand of 18.73 riders per minute.

Figure 4. Illustration of the Components of Routes Serving Riders 89 and 96



Notes. Calyber currently cannot afford either a commercial mapping service or the development of an in-house alternative, so it relies on Haversine distance for routing. The labels O_{89} , O_{96} , D_{89} , and D_{96} denote the origins and destinations of riders 89 and 96, respectively. (a) The routes followed if both riders are dispatched unmatched. In this case, their dispatching costs are based on the origin-destination (OD) lengths l_{89} and l_{96} . (b) The route followed when the riders are matched. In this case, the route is divided into solo components of length l_{89}^{solo} and l_{96}^{solo} and an overlap component of length $l_{89,96}^{overlap}$. The dispatching cost attributed to riders 89 and 96 are $l_{89}^{solo} + \frac{1}{2}l_{89,96}^{overlap}$ and $\frac{1}{2}l_{89,96}^{overlap} + l_{96}^{solo}$, respectively.

Figure 5. Demand (Riders per Minute), Cost per Rider (Dollars per Mile, Under Nick’s Engine), and Quoted Price (Dollars per Mile, Under Nick’s Engine) Across Different Neighborhoods of Chicago for a Typical Weekday from 7:30–8:30 a.m



Notes. Cost per rider can be above \$0.70/mile because of detours. Neighborhoods with limited data because of low demand are shaded gray and excluded.

Table 2. Sample of Calyber’s Data on Routes Taken to Fulfill Rider Requests

Route ID	Rider ID	OD length (miles)	Component length			Detour rate	Cost (\$/mile)
			Overlap	Solo	Total		
1	76	5.28	0.00	5.28	5.28	0.00	0.70
2	89	4.69	2.72	2.04	5.94	0.01	0.51
2	96	3.39	2.72	1.18	5.94	0.15	0.52

Notes. Details of route 2 are given in Figure 4. “OD length” is the direct distance from the rider’s origin to destination. “Overlap length” is the distance of the route component when both riders are in the vehicle. “Solo length” is the distance of the route component when only the given rider is in the vehicle. “Total length” is the length of the entire route, including shared and solo components for all riders (e.g., $5.94 = 2.72 + 2.04 + 1.18$). “Detour rate” is the fraction of the OD length that is added for sharing (e.g., $0.01 = (2.72 + 2.04 - 4.69)/4.69$). “Cost” is the cost per mile of fulfilling the given rider, allocating the cost along the shared component equally between the two riders (e.g., $0.51 = 0.7(2.72/2 + 2.04)/4.69$).

Table 3. Sample of Calyber’s Data on Rider Sessions

Rider ID	Route ID	Arrival date	Arrival time	Pickup latitude	Pickup longitude	Dropoff latitude	Dropoff longitude	Pickup area	Drop-off area	Solo length (miles)	Quoted price (\$/mile)	Converted	Waiting time (seconds)	Matched with
61	NaN	September 9	7:31:55	41.892	-87.665	41.881	-87.633	24	32	1.83	0.43	0	NaN	NaN
76	1	September 9	7:32:20	41.961	-87.752	42.009	-87.673	15	1	5.28	0.74	1	473.59	NaN
89	2	September 9	7:32:40	41.958	-87.653	41.895	-87.620	3	8	4.69	0.61	1	15.00	96
96	2	September 9	7:32:55	41.929	-87.646	41.881	-87.633	7	32	3.39	0.57	1	0.00	89

Notes. Riders who were quoted a price but did not request a ride have “Converted” = 0 and “Waiting Time” = NaN. Riders who requested a ride but were not matched with another rider have “Converted” = 1 and “Matched With” = NaN.

References

- Anand P (2017) Uber burned more than \$1 million a week on UberPool to win San Francisco. *Buzzfeed News* (May 31), <https://www.buzzfeednews.com/article/priya/uber-pool-burn-rate-frisco>.
- Davalos J (2023) Lyft will discontinue pooled rides. *Bloomberg* (May 11), <https://www.bloomberg.com/news/articles/2023-05-11/lyft-will-discontinue-pooled-rides-roll-out-new-features>.
- Desai S (2022) UberPool is a Zombie. *Atlantic* (July 5), <https://www.theatlantic.com/technology/archive/2022/07/uberx-share-carpooling-ride-app-cost/661483/>.
- Jackson AL (2024) Why Uber CEO Dara Khosrowshahi wants ride sharing to be “cool” again. *Fast Company* (March 13), <https://www.fastcompany.com/91056936/sxsw-uber-ceo-dara-khosrowshahi-sharing-rides-cool>.
- Kalanick T (2016) Uber’s plan to get more people into fewer cars. Accessed July 8, 2025, https://www.ted.com/talks/travis_kalanick_uber_s_plan_to_get_more_people_into_fewer_cars/.
- Kansal S (2025) For every day. *Uber* (May 14), <https://www.uber.com/newsroom/go-get-2025/>.
- Lunden I (2016) Uber says that 20% of its rides globally are now on UberPool. *TechCrunch* (May 10), <https://techcrunch.com/2016/05/10/uber-says-that-20-of-its-rides-globally-are-now-on-uber-pool>.
- Lung N (2025) Lyft revives pooled rides at airports in push for cheaper trips. *Fortune* (May 19), <https://fortune.com/2025/05/19/lyft-pooled-rides-at-airports-cheaper-trips/>.
- Lyft (2018) Lyft’s new app creates positive change for passengers and cities. *Lyft* (June 7), <https://www.lyft.com/blog/posts/new-app>.
- Lyft (2021) Shared rides return to more cities. *Lyft* (May 4), <https://www.lyft.com/blog/posts/shared-rides-return-to-more-cities>.
- Morning Brew Daily (2024) Uber CEO unveils plan to make rides cheaper, reveals own Uber rating, and launches “party bus.” Accessed July 25, 2025, <https://www.youtube.com/watch?v=Pw29jyXCihA>.
- Schneider T (2025) Taxi and ridehailing usage in Chicago. Accessed July 24, 2025, <https://toddschneider.com/dashboards/chicago-taxi-ridehailing-data/>.
- Sherman I, Greiner A, McFarland M, Tse J (2019) A history of Lyft, from fuzzy pink mustaches to global ride share giant. *CNN* (March 28), <https://www.cnn.com/interactive/2019/03/business/lyft-history/index.html>.
- The Verge (2025) Uber CEO Dara Khosrowshahi on the company’s new Route Share feature. Accessed July 24, 2025, <https://www.youtube.com/shorts/ERSAJRtrc2U>.
- Uber (2018) Introducing Express POOL: Walk a little to save a lot. Accessed June 29, 2023, <https://www.uber.com/newsroom/expresspool/>.
- Uber (2022) UberX Share. Accessed May 14, 2024, <https://www.uber.com/us/en/ride/uberx-share/>.
- Uber (2025) The history of Uber. Accessed June 19, 2025, <https://www.uber.com/en-CA/newsroom/history/>.
- Zimmer J, Green L (2017) The end of traffic. *Medium* (January 17), <https://medium.com/@johnzimmer/the-end-of-traffic-6d255c03207d>.