Riders, drivers and the price: evolutionary dynamics in ride-hailing pricing

Evidence from 301 million New York trips and a natural price shock

Rodrigo Hermont Ozon

2026-11-25

Abstract

A ride-hailing price acts on two populations that adapt with delay: riders, who choose between apps, and drivers, who choose where to work. We study both with the public trip records of New York City, 301.6 million Uber and Lyft trips. On the rider side, Lyft's January 2025 credit of the congestion fee is a natural price shock. Nested choice models pick a logit dynamic with inertia: the daily persistence of Lyft's log-odds share is 0.61 (95% interval 0.58 to 0.64), a half-life of 1.4 days, yet the dynamic model forecasts the post-credit share no better than a static one, and the credit's effect cannot be told apart from the same weeks a year earlier. On the driver side, a simulator calibrated on the data shows that a market-clearing surge amplifies drivers' reaction to earnings and makes prices cycle at imitation speeds where a fixed price is stable; smoothing delays the cycles. We also report a fare model, a week-ahead demand forecast, the size of a switchback experiment and a drift monitor.

1 Introduction

Surge pricing is usually analysed in equilibrium (Castillo et al. 2017; Garg and Nazerzadeh 2019). Equilibrium is where the market settles; pricing decisions are made on the way. Riders learn which app is cheaper and faster, drivers learn where the money is, and both learn with delay. Evolutionary game theory describes such populations: strategies spread in proportion to how well they pay relative to the alternatives (Sandholm 2010; Fudenberg and Levine 1998). The logit dynamic with inertia, in which a population moves part of the way toward its logit best response each period (Hofbauer and Sandholm 2002), nests two extremes that matter for pricing: static choice, where shares follow prices at once, and replicator-like drift, where a temporary advantage leaves a permanent mark. This paper asks which one the data support for riders, and what adaptive drivers imply for the choice of a pricing policy.

2 Data

The NYC Taxi and Limousine Commission publishes every high-volume for-hire trip (NYC Taxi and Limousine Commission 2026). We read fifteen monthly files (December 2023 to April 2024 and September 2024 to June 2025) with DuckDB, one pass per month, into zone-hour aggregates by app, a 0.5% trip sample and a monthly quality report. The report found that the time the car arrived is missing for 99.7% of Lyft trips until February 2025 and filled from March 2025; the wait is therefore measured as pickup minus request.

The congestion zone is identified from the data: pickup zones where at least 90% of trips paid

the congestion fee after February 2025 (38 zones). From 5 to 31 January 2025 Lyft credited riders the fee on rides in the zone, each credit valid for a week; we treat 5 January to 7 February as the credit period.

3 Methods

Rider dynamics. For each congestion-zone pickup zone and day, L_{zt} is the log-odds of Lyft's share among Uber and Lyft trips. The static model is $L_{zt} = a_z + d_{\text{dow}} + b\,C_t + c\,\Delta \log p_{zt} + e\,\Delta w_{zt} + u_{zt}$, with C_t the credit, $\Delta \log p$ the log ratio of Lyft's to Uber's fare per mile and Δw the difference in waits. The dynamic model adds $\rho L_{z,t-1}$: $\rho = 0$ is static choice, $0 < \rho < 1$ a logit dynamic with inertia, $\rho = 1$ replicator-like drift. Both are estimated through 7 February 2025 and forecast 8 February to 30 April, the dynamic model recursively from its own predictions. They are compared by root mean squared error of the share and a one-sided Diebold-Mariano test. The credit is also measured by difference in differences against zones outside the congestion zone, with zone and date fixed effects, and the same calendar a year earlier as placebo.

Drivers and policy. Six areas are calibrated on Lyft's weekday 18:00 profile from February to June 2025. Demand for the platform is $q = Q\,m^{\varepsilon}\,s/s_0$ with $\varepsilon = -0.4$ (Cohen et al. 2016); the wait is $w = w_0\,[(q/Q)(x_0/x)]^{\alpha}$ with α estimated within zone and hour of week; the rider share follows the dynamic choice model; drivers move by an exponential replicator on earnings relative to what each area usually pays, with imitation speed η. Policies: a fixed price, the market-clearing multiplier that restores the usual wait given today's drivers, and that multiplier smoothed by a factor λ. Under the clearing policy a driver's relative earnings equal the multiplier, which moves as $(x/x_0)^{1/\varepsilon}$: the policy multiplies drivers' reaction by $1/|\varepsilon| = 2.5$.

Forecast, fares, experiments, drift. Trips by zone and hour are forecast a week ahead with XGBoost on calendar and lags of one to four weeks, by rolling origin. A switchback planner uses the variability of log block margin. Drift is the population stability index of trip features each month.

4 Results

Riders adapt with inertia, but fast. The daily persistence is 0.61 (95% interval 0.58 to 0.64), a half-life of 1.4 days. A higher relative price per mile and a longer relative wait lower Lyft's share, by 0.30 in log-odds per unit of log price ratio and 0.08 per minute. Forecasting 8 February to 30 April 2025, the share error is 2.83 percentage points for the dynamic model and 2.80 for the static one (Diebold-Mariano p = 0.94): at horizons of weeks the inertia does not help.

The credit left no mark that the calendar cannot explain. Inside the zone, Lyft's log-odds share rose 0.021 during the credit (interval -0.002 to 0.045); the same weeks of the previous year, with no credit, show 0.029 (interval 0.013 to 0.045).

A clearing surge invites cycles. With slow drivers ($\eta = 0.5$) every policy settles, and the clearing surge earns 4.6% more margin than a fixed price. Prices start to cycle at $\eta = 0.75$ under the clearing surge, at 1.75 when it is smoothed to a tenth of the gap per day, and between 1 and 2 under a fixed price. At $\eta = 1$ the cycling clearing surge earns more margin than the smoothed one (US$ 72,156 against US$ 65,011 per evening peak hour across the six areas), because short-run demand is inelastic, while the mean wait rises from 5.08 to 5.59 minutes.

Forecast, experiment and drift. A week ahead, XGBoost's error (WAPE) averages 13.1%, against 13.0% for the mean of the last four weeks and 16.1% for the same hour last week. With

four-hour switchback blocks, two weeks detect a 15% change in margin, eight weeks 7% (Bojinov et al. 2023). Every feature's monthly stability index stays below 0.05.

5 Discussion

The rider side looks closer to static choice than the evolutionary framing suggests: inertia is real but short, and a five-week credit did not shift shares for good. The driver side is where dynamics bite. A policy that clears the market every period turns drivers' earnings into a function of the multiplier itself and makes the system oscillate once drivers imitate quickly; smoothing trades a little margin for stability. For a platform the practical rules are to smooth the multiplier, to judge policies on waits and supply stability as well as margin, and to test them with switchbacks long enough for drivers to settle.

6 Limits

Prices are chosen by the platforms, so the price coefficients of the choice model are associations. Credits and promotions are not in the fares. The daily persistence may mix habit with persistent local shocks. The simulator borrows the demand elasticity, assumes the observed split of drivers is a rest point and cannot identify the imitation speed. In 2025 Uber began passing some requests to yellow cabs, which leave this source.

Code and data

Code, tests and every table: https://github.com/rhozon/mobility-dynamic-pricing. Data: NYC TLC trip records, downloaded by `python -m mdp.cli download`.

References

Bojinov, Iavor, David Simchi-Levi, and Jinglong Zhao. 2023. "Design and Analysis of Switchback Experiments." *Management Science* 69 (7): 3759–77.

Castillo, Juan Camilo, Dan Knoepfle, and Glen Weyl. 2017. "Surge Pricing Solves the Wild Goose Chase." *Proceedings of the 2017 ACM Conference on Economics and Computation*, 241–42.

Cohen, Peter, Robert Hahn, Jonathan Hall, Steven Levitt, and Robert Metcalfe. 2016. *Using Big Data to Estimate Consumer Surplus: The Case of Uber*. Working Paper No. 22627. National Bureau of Economic Research.

Fudenberg, Drew, and David K. Levine. 1998. *The Theory of Learning in Games*. MIT Press.

Garg, Nikhil, and Hamid Nazerzadeh. 2019. *Driver Surge Pricing*. arXiv:1905.07544.

Hofbauer, Josef, and William H. Sandholm. 2002. "On the Global Convergence of Stochastic Fictitious Play." *Econometrica* 70 (6): 2265–94.

NYC Taxi and Limousine Commission. 2026. *TLC Trip Record Data*. Https://www.nyc.gov/site/tlc/about/tlc-trip-record-data.page.

Sandholm, William H. 2010. *Population Games and Evolutionary Dynamics*. MIT Press.