

# The persistence of driver bias on ride-sharing platforms



Ride-sharing platforms, such as Didi Chuxing, Uber, Lyft, and Via, are a manifestation of the sharing economy that has been disrupting traditional taxi industries worldwide. These platforms employ a simple-to-use mobile app to enable customers to request a ride. On the other side of the platforms, the app connects the rider to available local drivers who will fulfil the ride request. Drivers are independent contractors who work on a commission based on the price of the ride they provide, which depends on the location, ride length, and time of day.

These platforms attempt to optimise the supply (number of drivers) and demand (number of riders) by increasing the price of rides when demand outstrips supply and vice versa. This policy is called dynamic (or surge) pricing and provides strong incentives to both the drivers, who receive a larger payment for an identical ride than they would for the same ride at another time of day, and to riders, who may choose an alternative form of transportation (i.e., a bus or waiting until prices return to normal levels) in lieu of high prices. The net effect of this efficient reallocation of labour and resources is a large increase in consumer welfare.

With the advent and success of ride-sharing platforms, there has been hope that discrimination against under-represented minorities may be reduced relative to the levels documented in traditional taxi systems. However, early evidence suggests that the bias persists in ride-sharing platforms. Several platforms responded by removing information about the rider's gender and race from the ride request presented to drivers. In doing so, drivers do not have the information that enables a biased decision when a ride is requested, hopefully leading to reduced bias in the initial ride request stage. However, following this change, bias may still manifest itself through driver cancellation after a request is accepted, when the rider's picture is displayed.

In a recent article, we examine whether the alterations removed bias from the platform. The study aims to understand how rider race, gender and support for a social cause (visible, for instance, when they place a rainbow filter on their profile picture to show support for the lesbian, gay, bisexual, and transsexual community) affect the quality of the service they receive on ride-sharing platforms.

We investigate this by requesting 1,600 rides on a major ride-sharing platform in a large North American city. By randomly manipulating rider names and profile pictures, we observe drivers' patterns of behaviour in accepting and canceling rides. We measure the bias by the time it takes to have a ride confirmed, the quoted time to wait for the ride, and post-confirmation driver cancellation rates.

Our results confirm that bias at the ride request stage has been eliminated. However, after a ride has been confirmed, underrepresented minorities are nearly twice as likely to have a ride canceled as Caucasians. Riders that show support for the lesbian, gay, bisexual, and transsexual community, regardless of race or gender, also experience significantly higher cancellation rates. Biases related to gender appear to have been eliminated

We also explore the role of dynamic pricing on the extent to which bias exists. It is possible that a higher price for the same ride alleviates the bias because the financial incentive overcomes the (perceived) utility loss exhibited in the bias. On the other hand, a higher price can signal that there are many alternative potential riders so the driver can be selective about who they pick up. This would amplify the bias as a result of high prices. We find that the financial incentive of higher prices can alleviate the extent to which bias exists; bias is reduced or completely eliminated during times of peak demand as opposed to non-peak demand.

Our research leads to four main takeaways. First, the timing of communicating information to service providers must be carefully considered. In particular, while we all aim to minimise and eliminate bias, if it does exist they may be better off if the bias occurs earlier in the process (at the ride request stage) rather than later (post-acceptance). This is because cancelled riders also incur costs related to waiting before the cancellation occurs, and the inconvenience of re-requesting a ride. Rider characteristics either need to be fully hidden until the last possible moment for rider and driver to be safely connected, or fully visible from the ride request stage.

Second, while platforms are trying to make changes that remove bias from the system, they have not achieved a completely successful outcome yet. It is possible that a data-driven solution exists wherein rider characteristics are captured when a driver cancels and is penalised by the platform for biased behaviour. One possible way to punish drivers is to move them down the priority list when they exhibit biased cancellation behaviour so they have fewer ride requests.

Third, the economic incentive delivered via dynamic pricing is an important social benefit that has not previously been included in the estimates of ride-sharing platforms' benefits generally and dynamic pricing in particular. To further these benefits, it may be necessary for ride-sharing platforms to adopt a dynamic fee structure that incentivises drivers to accept rides from social groups or in specific areas of cities that are traditionally underserved.

Finally, despite public opinion polls indicating that support for the LGBT community is strong, there are still negative associations that can impact supporters.

Our results also raises some important questions. We document bias against LGBT supporters, which begs the question of whether/what other biases exist? Do signals of religious, political, or cultural affiliations result in drivers' biased behaviour? We regrettably expect that it will, but it is important to quantify the manner in which this bias may influence the quality of the service.

Additionally, if a ridesharing company has drivers who consistently favour a certain demographic, to what extent is the platform complicit and potentially legally responsible for the biased behaviour? If society's intent is to reduce discrimination, it may be necessary to increase firms' cost of discrimination. One would expect that, like the drivers with whom they contract, firms would respond to increased costs with better policies and monitoring of biased behaviour.



#### Notes:

- This blog post is based on the authors' paper [When Transparency Fails: Bias and Financial Incentives in Ridesharing Platforms](#), 2018
- The post gives the views of its author, not the position of LSE Business Review or the London School of Economics.
- Featured image credit: [Photo](#) by [iphonedigital](#), under a [CC-BY-SA-2.0](#) licence
- When you leave a comment, you're agreeing to our [Comment Policy](#).



**Jorge Mejia** is an assistant professor at Indiana University's Kelley School of Business, department of operations & decision technologies. ([jimmejia@iu.edu](mailto:jimmejia@iu.edu))



**Chris Parker** is an assistant professor of supply chain management at Pennsylvania State University's Smeal College of Business, department of supply chain & information systems. ([chris.parker@psu.edu](mailto:chris.parker@psu.edu))