

When Transparency Fails: Bias and Financial Incentives in Ridesharing Platforms

Jorge Mejia

Kelley School of Business, Indiana University, Bloomington, IN 47405, jmmejia@iu.edu, <http://go.iu.edu/jorge>

Chris Parker

Kogod School of Business, American University, Washington, DC 20016, chris.parker@american.edu,
<https://www.american.edu/kogod/faculty/christop.cfm>

Providing transparency into operational processes can change consumer and worker behavior. However, it is unclear whether operational transparency is beneficial with potentially biased service providers. We explore this in the context of ridesharing platforms where early evidence documents bias similar to what has been observed in traditional transportation systems. Platforms responded by reducing operational transparency through removing information about riders' gender and race from the ride request presented to drivers. However, following this change, bias may still manifest through driver cancellation after a request is accepted, at which point the rider's picture is displayed. Our primary research question is to what extent a rider's gender, race, and perception of support for lesbian, gay, bisexual, and transgender (LGBT) rights impact cancellation rates. We investigate this through a large field experiment on a major ridesharing platform in Washington, DC. By manipulating rider names and profile pictures, we observe drivers' behavior patterns in accepting and canceling rides. Our results confirm that bias at the ride request stage has been eliminated. However, after acceptance, racial and LGBT biases are persistent, while we find no evidence of gender biases. We also explore whether peak times moderate (through increased pay to drivers) or exacerbate (by signaling that there are many riders, allowing drivers to be more selective) these biases. We find a moderating effect of peak timing, with lower cancellation rates for non-caucasian riders. We do not find a similar moderating effect for riders that signal support for the LGBT community.

History: September 20, 2019

1. Introduction

Providing transparency into operational processes can change consumer and worker behavior resulting in improvements in terms of perceived quality (Buell and Norton 2011), worker efficiency (Buell et al. 2016), and likelihood of citizens to re-engage with a government (Buell et al. 2019). However, common practice in the sharing economy is to connect customers (demand) with suppliers (supply) via a two-sided platform that actively tries to optimize matches based on market conditions. In such cases, it is unclear how operational transparency may affect service outcomes. Moreover, while the algorithmic matching of these two-sided platforms may be efficient based on specific objective

criteria, the behavior from the humans executing the service delivery after matching may not. In particular, humans have been known to exhibit biases towards specific demographics that may shape service delivery in the presence of customer transparency, or providing a worker with a view to the customer for whom they are providing the service.

Examining operational transparency while accounting for biased behavior in humans is important because it may alter the conclusions reached so far in the literature. Indeed, we believe this is a crucial expansion for service operations research because despite the well-documented prevalence of bias, its effects on service delivery remain mostly unknown. In the U.S., extant research has identified bias in a myriad of contexts, such as hiring (Goldin and Rouse 2000), performance evaluations (Greenwood et al. 2017), housing options (US Department of Housing and Urban Development 2013), and professors' responses to student requests (Milkman et al. 2012). Indeed, a recent NPR report found that 92% of black Americans surveyed believe that discrimination exists in some form today (Blendon et al. 2017). More than half of the respondents have personally been discriminated against in three categories: promotion or pay, applying for jobs, and interacting with police. Nineteen percent said they had been discriminated against when trying to vote or participate in politics. The negative outcomes of the discrimination can be devastating; recent data from 52 U.S. cities shows that 48 of them are more likely to make an arrest for a Caucasian homicide victim case rather than a minority victim case (Rich et al. 2018). Based on this evidence, it would thus appear that providing information about customer demographics may enable these biases and provide sub-optimal results. We reach then a potentially serious conflict: while providing operational transparency may increase service quality, it may also enable biased behavior, which may in turn lowers service quality. This paper aims to address this gap in the literature.

In particular, this study aims to better understand bias and operational transparency in the context of ridesharing platforms—such as Didi Chuxing, Lyft, Uber, and Via—that have been disrupting traditional taxi industries worldwide. These platforms employ a simple-to-use smartphone application where customers can request rides. The app matches the rider to available drivers who are independent contractors working for a commission based on the price of the ride; origin, duration, distance, and time of day determine the price. Platforms attempt to match supply (number of drivers) and demand (number of riders) through peak (or dynamic or surge) pricing, i.e., increasing the price of rides when demand outstrips supply and vice versa. Peak pricing provides strong incentives to both drivers, who receive a larger payment for the ride than they would for the same ride at another time of day, and to riders, who may choose an alternative form or time of transportation (i.e., driving, public transportation, or waiting until prices return to normal levels) in lieu of high prices. The net effect of this efficient reallocation of labor and resources is a large increase in consumer welfare (Cohen et al. 2016).

Despite the positive welfare gains and popular press reports on the benefits of the sharing economy, early work has found that bias persists in the sharing economy (e.g., Cui et al. 2016, Ge et al. 2016, Edelman et al. 2017). Indeed, in these situations where the biased decision is made before physical interaction between people occurs, it is difficult (if not impossible) for those biased against to be aware that the bias even exists. As Blendon et al. (2017) note: “The full scope of the disparate treatment often becomes clear only in the aggregate, once the camera zooms out.”

Platforms responded to drivers’ biased behavior by removing information about a rider’s gender and race from the ride request sent to drivers, i.e., altering the timing of customer transparency. In doing so, drivers do not have the information that enables a biased decision when a ride is requested, ideally removing bias in the initial ride request stage. If this occurs, either the platform change in the ride request stage reduced discrimination on the whole or has simply shifted the biased behavior until after the ride is confirmed resulting in higher cancelation rates rather than longer service times associated with waiting for a ride to be confirmed or a driver to arrive.

Our primary research question is whether this change in the timing of customer transparency removed bias from the platform. Understanding whether the bias has been removed is important for ridesharing companies as they continue to compete not only against each other but also with traditional transportation options. Building trust in the consistency and quality of the platform is key but requires carefully balancing operational transparency and responsiveness. In the rideshare context, increased operational transparency makes the bias exhibited by service providers more readily observable by those discriminated against and may reduce the perceived value of the service and ultimately their overall usage rates.

To investigate this, we perform a field experiment on a ridesharing platform from early October to mid November 2018 in Washington, DC. We randomly vary three dimensions of rider characteristics: 1) Gender: Female vs. Male using the name and profile picture, 2) Ethnicity: Caucasian vs. Under-Represented Minority (URM) using the name and profile picture, and 3) Support for LGBT rights: rainbow profile picture filter vs. no rainbow profile picture filter. In addition, we vary the time at which we request rides to determine how peak and non-peak price periods impact bias.

We make three contributions to the literature. First, we add to the operational transparency literature by demonstrating that transparency does not operate in isolation but must be considered in the context of potentially biased actors. In particular, bias remains an unintended consequence of ridesharing platforms despite recent changes to the timing of customer transparency. While the change has, as expected, eliminated ride-request stage bias observed in previous research (Ge et al. 2016), bias persists at a later stage in the process. In the aggregate, URMs are more than twice as likely to have a ride canceled as Caucasians (3.1% vs. 7.5%). In contrast to previous research (e.g., Ge et al. 2016), we find no differences in cancelation rates across genders.

Second, we show that signaling support for a social cause (in our case, the lesbian, gay, bisexual, and transgender (LGBT) community) can impact service provision. Riders that show support for the LGBT community, regardless of race or gender, also experience significantly higher cancellation rates. We believe we are the first to use social support as a bias-enabling characteristic.

Third, we examine the role of peak time 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, peak times can signal that there are many other potential riders so the driver can be selective about who they pick up. This would amplify the bias. Our results show that the financial mechanism of surge pricing during peak times can alleviate the extent to which bias exists. However, this reduction in bias appears to only occur for URMs and not for those that support the LGBT community. While we cannot observe characteristics of the drivers, analysis of available data and a subsequent survey using Amazon Mechanical Turk show that the populations of peak and non-peak drivers are not fundamentally different providing evidence that the reduction observed for URMs in peak times is driven by pricing and not other factors.

In addition, we contribute managerially by highlighting how “small” pieces of information affect service quality. Our study shows that certain nudges (i.e., delays of information) for the sake of reducing bias can backfire on the very users it sought to protect and also decrease the overall economic efficiency of the platform for everyone. Our analysis should allow platforms and policymakers to reflect on the type and timing of information they give to service providers and whether financial mechanisms can impact the extent to which bias exists on their platforms. We hope to start a conversation about the rider-driver relationship and emphasize the important role of platform governance. We conclude the paper by discussing managerial implications, highlighting limitations of the current analysis, and identifying opportunities for future research.

2. Literature Review

Our research is related to five literature streams: 1) operational transparency, 2) behavioral operations management, 3) bias and discrimination, 4) information technology’s (IT’s) impact on markets and society, and 5) dynamic pricing and economic incentives. Below we briefly summarize these and highlight where our paper fits into the broader research.

2.1. Operational Transparency

Operational transparency is the idea that providing customers and employees with information about inner operations of a firm or government can impact outcomes. Traditionally, operational transparency has been thought of as what is now termed “process transparency,” where customers are allowed to observe operational processes. Providing process transparency can increase quality

ratings (Kruger et al. 2004, Chinander and Schweitzer 2003) as well as customer satisfaction and gratitude (Mohr and Bitner 1995, Morales 2005). In the end, with process transparency, customers may perceive results to be of a such higher quality that they prefer to wait longer for a service (Buell and Norton 2011). However, Buell et al. (2019) demonstrate that operational transparency increases trust in the government and willingness to re-engage in the future only if citizens are shown that their service request was adequately addressed. Increasing process transparency is detrimental when the government fails to act on the service request.

More recently the idea of operational transparency has expanded beyond process transparency in two important ways. First is distinguishing cost transparency where customers receive information about the costs the firm faces. Providing this information to customers has been found to increase sales by more than 20% (Mohan et al. 2018). Second, and more closely related to this paper, is customer transparency where employees are given information about the customers they are serving. Research has hinted at the possibility that customer transparency can create value for employees (Oldham and Hackman 1978), lead to a situation where employees identify with customers resulting in improved job performance (Gino et al. 2010, Small and Loewenstein 2005), and focus an employee on an individual customer leading to improved service levels (Friedrich et al. 1999). Building on this work, Buell et al. (2016) show that allowing employees to see the customers they are serving results in a 22% increase in customer-reported quality as well as increased job satisfaction for employees. This paper expands the knowledge of customer transparency by examining how it interacts with bias in the ridesharing economy, a previously underexplored terrain.

2.2. Behavioral Operations Management

Previous research has shown that behavioral biases impact operations management outcomes. Gino and Pisano (2008) provide a clear definition of behavioral operations management, while Bendoly et al. (2006) provide a thorough review of this field. Their note (on p. 737) that “the success of operations management tools and techniques, and the accuracy of its theories, relies heavily on our understanding of human behavior” resonates in the work that has followed. Behavioral decision making has been shown to have several detrimental effects including low cognitive reflection exacerbating the bullwhip effect (Narayanan and Moritz 2015, Moritz et al. 2019) and increased workloads leading to significant differences in midwife services (Freeman et al. 2017).

Behavioral decision making and its impact on services employees provide can have drastic implications for their firm. In some cases, poor service quality can impact firm performance so much that firms are forced to close (Mejia et al. 2015). While previous work has looked at traditional ways in which employee behavior impacts a firm, we explore how personal biases may impact the service a firm provides.

Behavioral operations management has also identified that seemingly simple policy fixes may have unintended consequences. For example, Lurie and Swaminathan (2009) show that providing timely information to inventory managers can actually induce demand chasing behavior and ultimately lead to worse performance, while Choudhary et al. (2018) find that real-time feedback induces worse driving behavior leading to a 1.65% increase in accident probability. Timely data is not the only culprit in these seemingly simple fixes. Chen and Savva (2018) find that the introduction of the Hospital Readmissions Reduction Program reduced hospital readmissions as designed but simultaneously increased the number of patients under observation. Similar issues are at play in the sharing economy. Indeed, after recent reports of bias in the sharing economy (e.g., Cui et al. 2016, Ge et al. 2016, Edelman et al. 2017), some platforms have to delay or even limit the amount of information provided to users to reduce instances of bias. These actions fit a pattern of unintended, and often paradoxical, consequences from transparency decisions well-established in the literature, such as stifling innovation (Brown and Martinsson 2018), increasing mistrust (Strathern 2000, Hansen and Flyverbom 2015), and weakening accountability (Roberts 2009). However, unlike most of this work, which focuses on unintended consequences from increasing transparency, our study follows the unintended consequences from reducing transparency. Moreover, by reducing the amount of transparency, these platforms are going against the long-held view in economics that reducing information asymmetry increases market efficiency and lower transaction costs (Stiglitz 2002, 2003, Kono 2006, Libich et al. 2006).

2.3. Bias and Discrimination

Biases have been broadly documented in the academic literature. Below we briefly introduce some of these papers in an effort to provide context for this paper. An interested reader should read Bertrand and Duflo (2017) for a comprehensive review. Much of the work in this area has been in the context of hiring, where experiments have been created by submitting resumes or creating job postings that differed in the applicant's name but were otherwise identical. Differences in callback rates are attributable to the gender and/or race of the applicant. Bertrand and Mullainathan (2004) is one of the earlier papers that explores bias in the context of help wanted ads. White-sounding names received 50 percent more callbacks than African-American-sounding names with minimal differences across the job type, industry, or employer size. Furthermore, there is little evidence that the results are due to differences in perceived social class. The names we use for riders come from this paper. Other work has found that removing gender information from hiring decisions can eliminate the gender gap previously observed in hiring for orchestras (Goldin and Rouse 2000) and that performance evaluations also tend to be lower for women (Greenwood et al. 2017).

While much of the evidence surrounding bias has focused on hiring in traditional labor markets, there is substantial racial and gender bias in other markets as well. In a broad study of the U.S. housing market, US Department of Housing and Urban Development (2013) finds that URM are both told about and shown 6.6% to 18.8% fewer properties than Caucasians. The effect is consistent across renting and buying and are not the result of potentially-biased credit rating differences. Disappointingly, academia has also been found to be susceptible to these biases. Milkman et al. (2012) sent meeting requests for potential doctoral students to more than 6,500 faculty members and altered both gender and race as well as the timing of the request (that day or in one week). The authors find no differences when the request is to meet the same day but that Caucasian males were more likely to get a meeting with the faculty member and received faster responses.

Bias against people who are LGBT is also well documented (see Ozeren (2014) for a recent review). For example, a recent report on bias in the workplace “finds consistent evidence of sexual orientation and gender identity discrimination” (Badgett et al. 2007). This discrimination has been shown to affect hiring decisions (e.g., Weichselbaumer 2003, Tilcsik 2011) as well as wages and income (Berg and Lien 2002). URM in the LGBT community have been shown to be especially economically disadvantaged because of persistent discrimination on multiple levels. In fact, a 2011 report found that “32 percent of children being raised by Black same-sex couples live in poverty, compared to 13 percent of children being raised by heterosexual Black parents and just 7 percent being raised by married heterosexual white parents” (Movement Advancement Project, Family Equality Council, and Center for American Progress 2011).

Our paper differs from previous work on biases in two important ways. First, while we are aware of discrimination based on sexual orientation and gender identity, it is unclear whether support for this specific community may also give rise to similar biases. This type of implicit bias against the community has not been as extensively explored to our knowledge. This paper extends our understanding of bias beyond race and gender and includes signaling support for the LGBT community. Second, while the majority of the literature focuses on biases from the employers’ perspective in hiring decisions, we focus on the bias from the supply side (the drivers), which a priori should not see any difference in the service provision (i.e., cost) for any individual or group.

2.4. IT Impact on Markets and Society

IT has the potential to make significant societal improvements. Indeed, research has shown that IT and IT-enabled applications can improve educational outcomes (Aker et al. 2012), increase economic efficiency in India’s agriculture markets (Parker et al. 2016), and reduce violence against women (Cunningham et al. 2017). The benefits of IT have also been noted in the sharing economy. The introduction of Uber has been shown to reduce motor vehicle fatalities (Greenwood and Watal 2017) and decrease low-quality entrepreneurial activity (Burtch et al. 2017).

Table 1 Recommendations Regarding Transparency in the Literature

Study	Context	Recommendations
Cui et al. (2016)	Airbnb	Increase transparency
Edelman and Luca (2014)	Airbnb	Reduce transparency
Edelman et al. (2017)	Airbnb	Reduce transparency
Ge et al. (2016)	Ridesharing	Reduce transparency
Greenwood et al. (2017)	Ridesharing	Maintain transparency; focus on debiasing intervention
Pope and Sydnor (2011b)	Crowdfunding	Maintain transparency; focus on debiasing intervention
Pope and Sydnor (2011a)	Crowdfunding	Maintain transparency; focus on debiasing intervention
Younkin and Kuppaswamy (2017)	Crowdfunding	Reduce transparency

Recent research makes different recommendations as to the way firms should handle transparency.

Despite the general promise of IT, evidence is mounting that bias persists in the IT-enabled sharing economy. Ge et al. (2016) show that black passengers on a ridesharing platform that shares profile names and pictures at the time of a ride request must wait for rides about 35% longer than an otherwise identical Caucasian passenger. Edelman et al. (2017) document similar issues with Airbnb. However, few studies explore potential mechanisms by which a bias can be reduced. One notable exception is Cui et al. (2016) which uses a field experiment to show that Airbnb hosts exhibit biases when deciding on which guests to host. Guests with predominantly African American names are 19 percentage points less likely to be accepted than guests with a Caucasian name. The authors manipulate guest ratings in their experiment and find that one online review can remove this bias in acceptance rates, leading to the conclusion that platforms should carefully consider *what* information the service providers can access. Greenwood et al. (2017) find that ratings of ridesharing drivers are also susceptible to riders' biased behavior. We extend this literature by testing whether showing support for the LGBT community also results in discriminatory behavior within the sharing economy and testing whether a simple economic pricing mechanism amplifies or alleviates biases.

Furthermore, our experiment comes after a major platform change. In Ge et al. (2016), race and gender information was available to drivers at the request stage. As a result, the authors conclude that it is important to decide *when* information is provided to the driver. Table 1 shows several other studies and the recommendations they make about operational transparency. Observing these results, many ridesharing platforms made changes to the information available to drivers at the request stage, pushing a large portion of information that enables a driver to make a bias-motivated decision until after the ride is confirmed. By running our experiment after this change, we can determine whether the alteration eliminated the biased behavior or simply shifted it from the ride request to the post-accepted stage.

2.5. Dynamic Pricing and Economic Incentives

The proliferation of cheap IT has resulted in a near elimination of menu costs. As such, dynamic pricing models are being applied to transportation, parking, online retail, and even beer. Empirical research in this area finds that dynamic pricing has a double-edged sword. Zhang et al. (2017) find that dynamic pricing for items abandoned in an online shopping cart can increase short-term purchase probability and revenues. However, customers quickly learn that adding items to their cart and waiting may get them a discount, eliminating the longer-term benefits of dynamic pricing. In another online retail setting, Bandi et al. (2017) find that dynamic pricing induces strategic consumer behavior that leads to an increase in operational costs associated with returns.

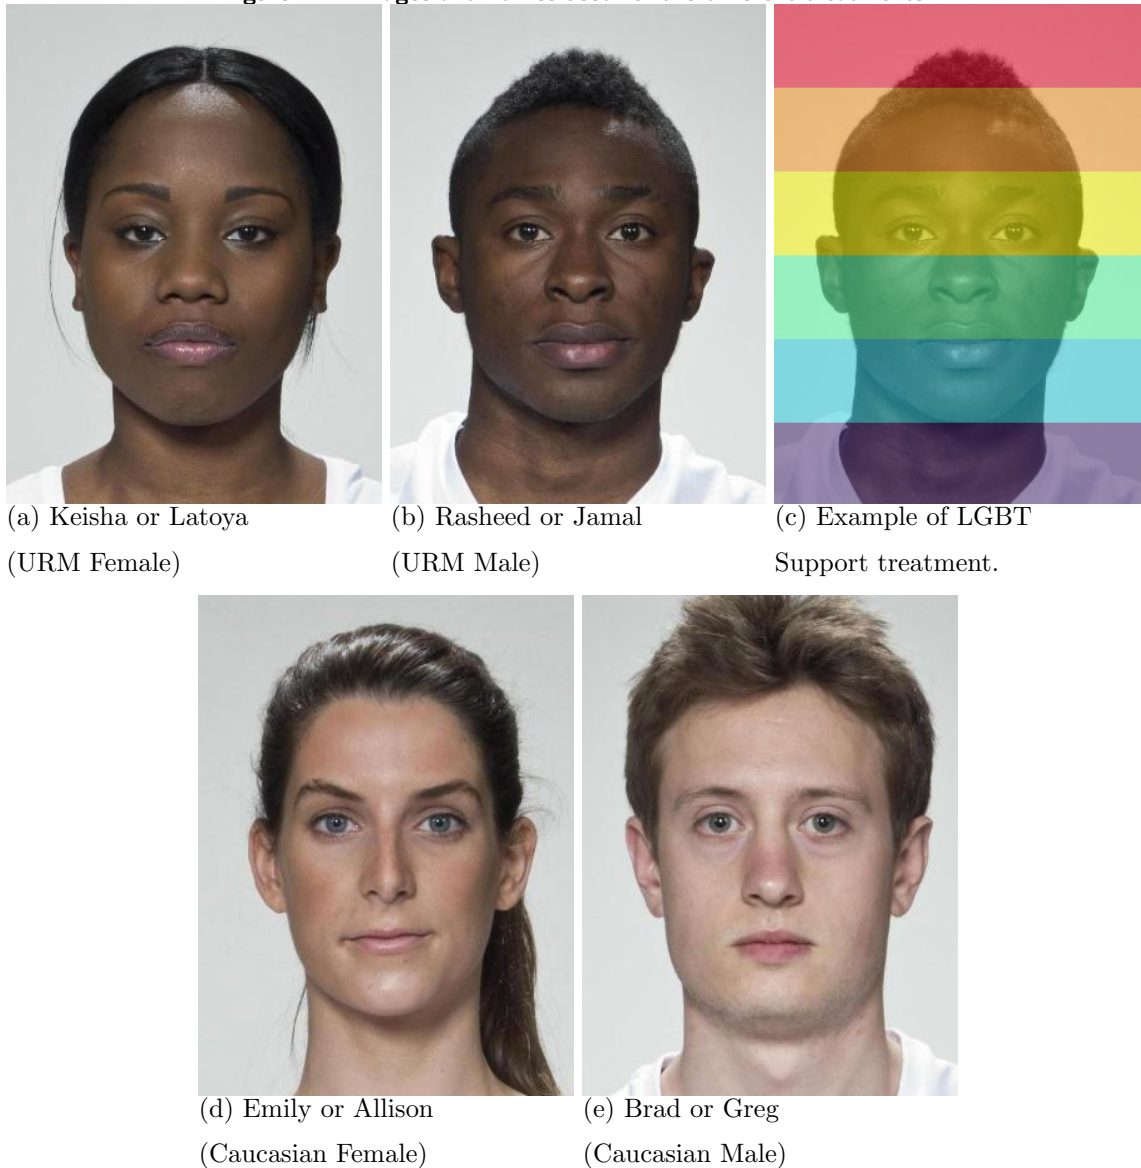
Labor markets also exhibit a similar double-edged sword. On the one hand, introductory economics suggests that a higher wage will encourage more people to apply for the same job. Alternatively, Yu (2018) demonstrate that a thick market can lead to increased selectivity in two-sided markets. In the context of ridesharing and discrimination, there are two related effects that seem feasible. First, higher prices can increase the utility of drivers, resulting in fewer cancelations for minority riders as a result of raising the cost of discrimination for the driver. Second, high prices can signal strong demand. In doing so, drivers may be more selective in who they drive which could increase cancelation rates for URMs. It is an empirical question as to which effect dominates; answering this question is part of our contribution.

3. Experimental Design

In this section we describe the experiment performed to test whether the operational transparency change eliminated bias and the effects of dynamic pricing on that bias. We also detail the three measures of bias we use.

3.1. Treatment Groups

Our experiment contains four treatments each with two possible scenarios for a total of sixteen distinct treatment cells. Three of the treatments alter rider characteristics. Two of these treatments vary names and profile pictures to show that the rider is likely to be 1) Female or Male (F/M) and 2) Caucasian or URM (Cau/URM). The third treatment uses a rainbow-colored overlay on the profile picture to signal that the rider is an ally of the LGBT community (LGBT/Non-LGBT). Rainbows are widely known to be a signal of support for the LGBT community (e.g., Dewey 2015). Figure 1 shows the names and pictures used for the different ethnicity and gender treatments as well as an example of the rainbow filter applied for the LGBT support treatment. At least 92% of people perceive Emily, Allison, Brad, and Greg as Caucasian names and Keisha, Latoya, Rasheed, and Jamal to be black names (Bertrand and Mullainathan 2004). No last names were used as this information is never available to the driver. The pictures come from The Face Database (Minear

Figure 1 Images and names used for the different treatments

and Park 2004). Prior research has shown a connection between attractiveness and social and economic outcomes (e.g., Joseph 1982, Benson et al. 1976, Marlowe et al. 1996, Cryder et al. 2017). To ensure differences in attractiveness are not driving our results, we pre-tested the pictures using two different online tests. All four pictures had an attractiveness of seven or eight out of ten on www.anaface.com, which relies on users selecting a gender and manually marking critical points on the subject's face to create the score. Additionally, all four pictures had a six or seven out of ten on www.hotness.ai, which relies on face recognition algorithms and deep learning to calculate an attractiveness score.

The final treatment is used to examine whether pricing amplifies or alleviates biases on ridesharing platforms. In this treatment, we vary the time of day to be when there is typically not much demand such that prices are likely to be low or at a time when there is high demand such that prices are likely to be high (Non-Peak/Peak).¹

The experiment is run using a major ridesharing platform in Washington, DC from early October to mid November in 2018. Throughout the experiment, the pickup is fixed at a central Metro stop, and the drop-off location is fixed at a major airport to reduce any geographic variation that may exist in the bias. The rider's rating is 4.8 throughout.

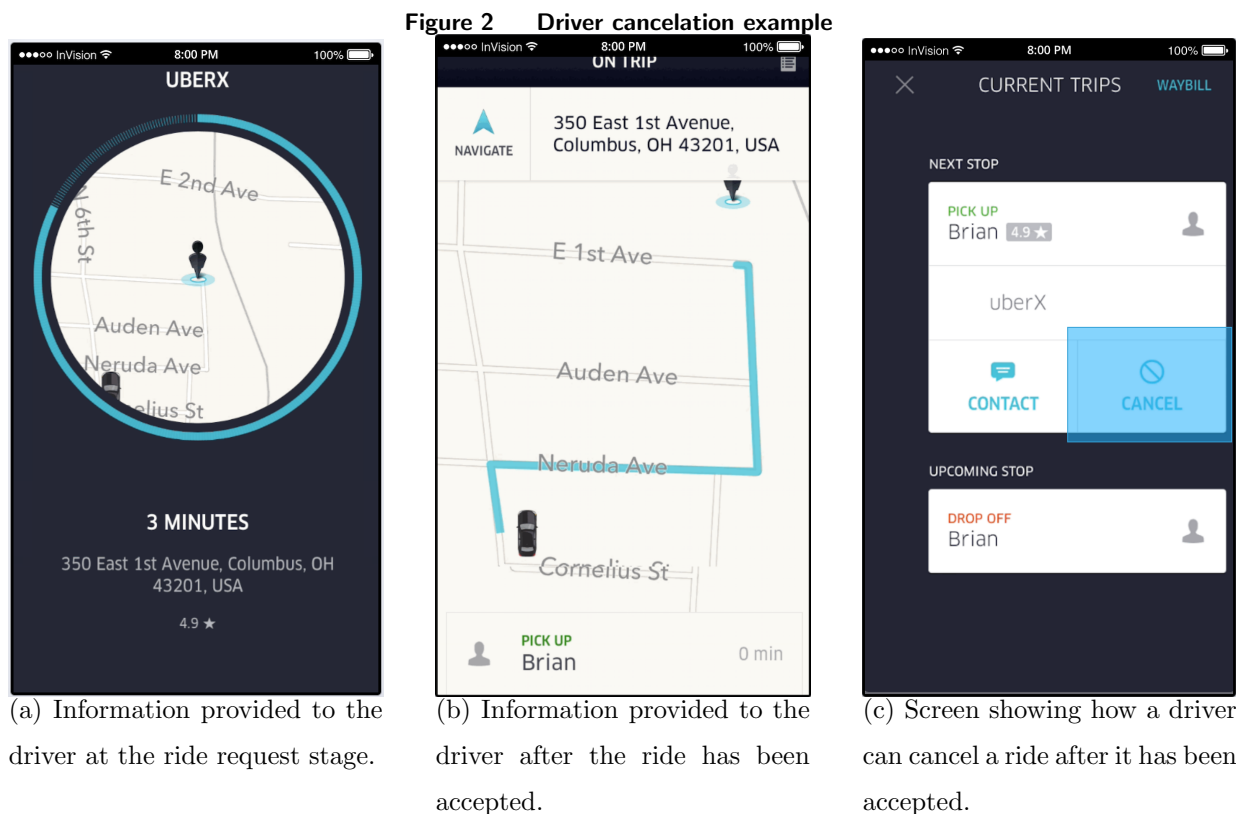
Throughout the experiment we alter gender, ethnicity, LGBT support, and timing of ride requests. Specifically, the researchers' computer program randomly selects a treatment cell from the possibilities described above. Then a ride request is generated through the platform's API with the manipulated rider characteristics corresponding to the treatment cell. The platform offers a driver (i.e., the subject in this study) the ride at which point the driver decides to accept or reject the ride. As the experiment is implemented after the platform change, the driver can only observe the general location of the rider at this stage. If the driver rejects the ride, then the platform offers the ride to a new driver. This continues until a driver accepts the ride at which point the platform notifies the API that the driver is on the way.

Once the ride is confirmed, the rider's name and profile picture become visible to the driver as well as the exact pickup location. We wait three minutes to allow the driver to cancel if they no longer want to accept the rider. If the driver cancels the ride, they receive no compensation. If the driver has not canceled after three minutes, we cancel the ride, and the driver receives compensation in the form of a cancellation fee, which ranges from \$5 to \$10. They are paid at the end of the interaction through the platform as with any other ride and in line with the cancellation policies established for both riders and drivers on the platform. Figure 2 shows an example of the rider information available to Uber drivers at different points in time and how those drivers cancel a ride. The platform on which the experiment was performed has a similar interface.


We request an equal number of rides during Peak and Non-Peak times across our experiment so that there are 1,600 in each (for a total of 3,200 observations). However, there is not an equal number of Peak and Non-Peak hours each day; AM Peak is 7-10am and PM Peak is 4-7pm for a total of six Peak and eighteen Non-Peak hours each day. During Peak hours we call once every ten minutes for a total of 36 peak observations per day. In order to get 36 Non-Peak observations per day, we need to make 36 ride requests during the 18 Non-Peak hours, which means we request a ride every 30 minutes.²

¹ Peak and Non-Peak times are defined by the ridesharing platform itself and not by the researchers.

² Technically, we call a ride randomly in minutes 1-7 of each ten minute period during Peak and in minutes 1-27 during Non-Peak times. The additional three minutes allows us to cancel the ride before the next 10 minute or 30 minute period begins.

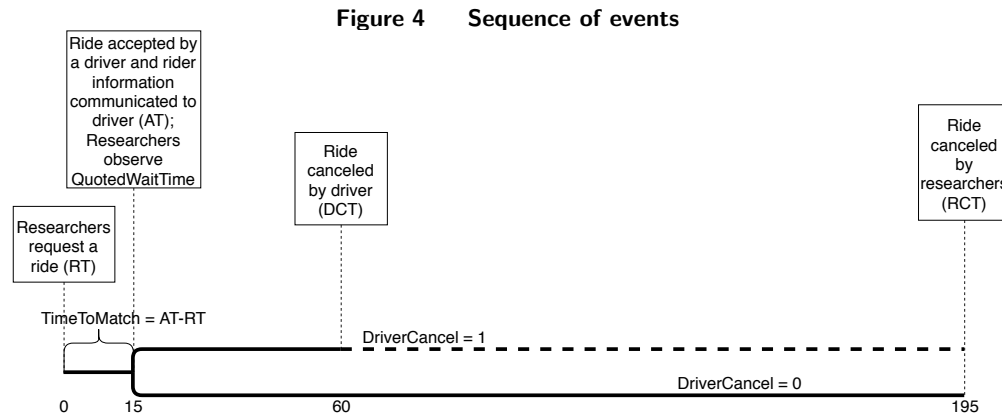


Note. Example of what a typical ridesharing platform driver would see during the request process (a), after accepting a ride (b), and when canceling a ride (c). We note that these are not actual requests from the experiment but are from early tests in an emulator. The city shown in the figures is Columbus, OH, but the experiment was performed in Washington, DC.



Note. Non-Peak and Peak calls are spaced at different intervals to account for the fact that Peak times are a smaller part of the day than Non-Peak times.

Institutional Review Board (IRB) approval required that we minimize the chance that we were contacting the same riders repeatedly. In particular, IRB approval was contingent on: 1) performing the experiment in a large city with a large number of ridesharing platform drivers, and 2) requesting at most 5,000 rides throughout the experiment. These two conditions combined make



Note. Sequence of making a ride request. The top line is when the driver cancels the ride and the bottom line is when the researchers cancel the ride.

it unlikely that we would be matched to the same driver two times in a row. Publicly-available estimates suggest that there are approximately 40,000 active Uber drivers in the Washington, DC area (Di Caro 2017). The platform on which the experiment was performed has a similar number of drivers in the area.

IRB also mandated that we could not record any driver information outside of the experiment. We were allowed only to put driver details through a hash function that enabled us to record an anonymized identifier of the driver that accepted the ride request. Of the 3,200 rides we requested, there were a total of 3,189 drivers that accepted the ride requests. The same driver was not observed twice in the same day. We could not observe any information about the drivers who were offered but did not accept the rides we requested, so we cannot explicitly measure how often a driver who was just matched to our rider is offered the next ride we request.

Figure 3 shows how ride requests are timed throughout the day. The middle row shows actual time, where midnight is both the left and the right side of the row. A call is made in each of the 72 blocks in the day. During the AM and PM Peak times, rides are requested at a faster rate as indicated by the skinnier blocks. However, when we zoom into both the Peak (top row) and Non-Peak (bottom row) hours, both have 36 blocks each that are equally spread out across time. Rides were requested for a period of just over 44 days from early October to mid November in 2018. In particular, for each ride request we randomly selected one of the eight rider characteristic combinations above. Any requests during Peak hours were coded as such. This randomization scheme results in an equal number of Peak and Non-Peak rides, but an unequal number of ride requests for each of the eight rider characteristics combinations.

3.2. Measures of bias

Bias in ridesharing can be exhibited in several ways (Ge et al. 2016). Our study is particularly concerned with drivers either (1) refusing to accept ride requests from certain types of riders, or (2) canceling a ride once rider characteristics are shared.³ We measure these biased behaviors through three operationalizations of the service quality, each of which is connected to one of the biased behaviors above: 1a) the time between when a ride is requested and when a ride is confirmed $TimeToMatch = AcceptTime - RequestTime$, 1b) the initial expected wait time quoted to the rider upon receiving a confirmed ride $QuotedWaitTime$, and 2) whether or not the driver canceled the ride $DriverCancel$, a dummy variable taking the value of 1 if a driver canceled within three minutes of accepting the ride and zero otherwise. Differences in these measures between the different cells are an estimate of the implicit bias towards that group. Figure 4 shows how the sequence of making a ride request occurs and how the metrics above are calculated.

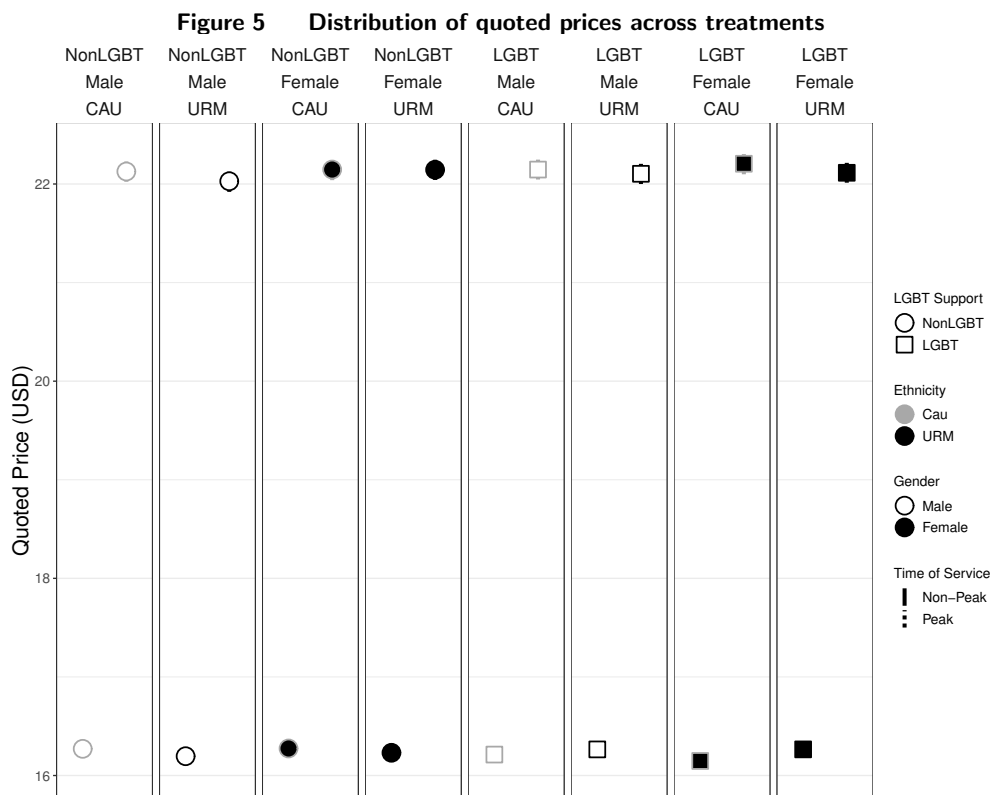
4. Experimental Results

In this section we estimate the cancellation rates for our different treatments (e.g., race, gender, and LGBT support) after the platform decided that drivers should not be shown the race and gender information at the ride request stage. Our objective is to estimate possible biases that remain on the platform after the transparency change. We first show model-free evidence of bias, then provide formal regression analysis to determine whether the differences in cancellation rates are significant and explore whether Peak period cancellation rates are lower for riders with different characteristics. We end the section with additional robustness checks where we (1) explore whether drivers are different during Peak and Non-Peak hours and (2) whether drivers view the rainbow filter as a signal of being in the LGBT community and/or supporting the community.

4.1. Model-Free Evidence of Bias

Before examining driver-related bias, we first demonstrate that the pricing algorithm does not exhibit any price differences across the different rider-related treatments. Figure 5 shows that quoted prices are similar across the different treatments and are indeed higher during Peak times than in Non-Peak times. We note that the price we are shown is actually a range of prices and not the upfront pricing. For example, we observe a price of \$16 to \$20. The plot shows the average of these two prices. The price ranges we observe do not vary in any meaningful way from day to day or across treatments but vary substantially between Peak and Non-Peak times.

³ We also measure the time to cancel conditional on the driver canceling. However, there are only a few observations for each of the treatment cells making an analysis of this measure of bias largely uninformative.

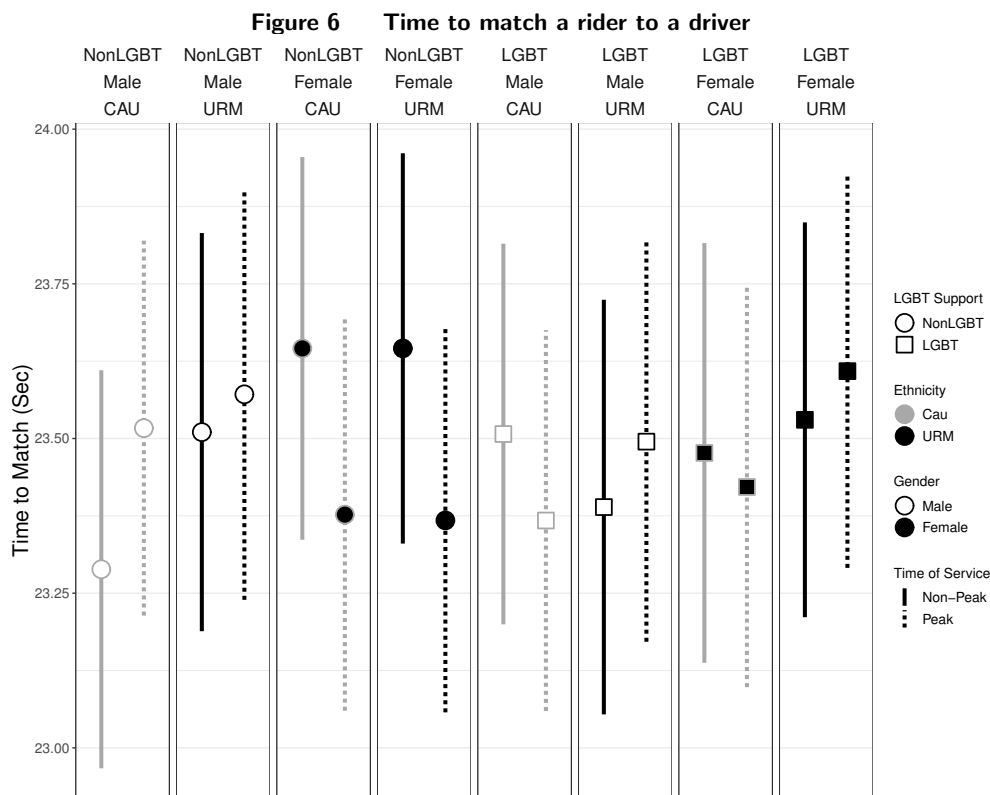


Note. There are no statistical differences in quoted prices across the different treatments. Bars show 95% confidence intervals.

Having found no evidence of algorithmic price differences, we now turn to examining the drivers’ bias. The field experiment, conducted after the platform change, demonstrates that the bias previously observed through higher *TimeToMatch* and *QuotedWaitTime* metrics appears to have been eliminated. Figure 6 shows that *TimeToMatch* is close to 23.5 seconds for all of the treatments, while Figure 7 shows that *QuotedWaitTime* is close to six minutes for all of the treatments. None of the differences are statistically significant.

While removing bias from the ride request stage is an achievement in and of itself, it does not mean that bias has been removed from the platform. Indeed, we find evidence of bias against URMs and LGBT supporters once a rider is confirmed as shown in Figure 8.

First, we explore the gender dimension. Although previous research did find biases against female riders (e.g., Ge et al. 2016), we find no significant differences across the gender treatment. It is unclear why the biases previously found against women do not appear in our experiment. One possibility is that the platform change removed bias against women. Alternatively, the timing (during/post the #MeToo movement) may have played a role. Also, the specific location of the experiment could be such that these biases did not previously exist.



Note. *TimeToMatch* is statistically indistinguishable across treatments. Bars show 95% confidence intervals.

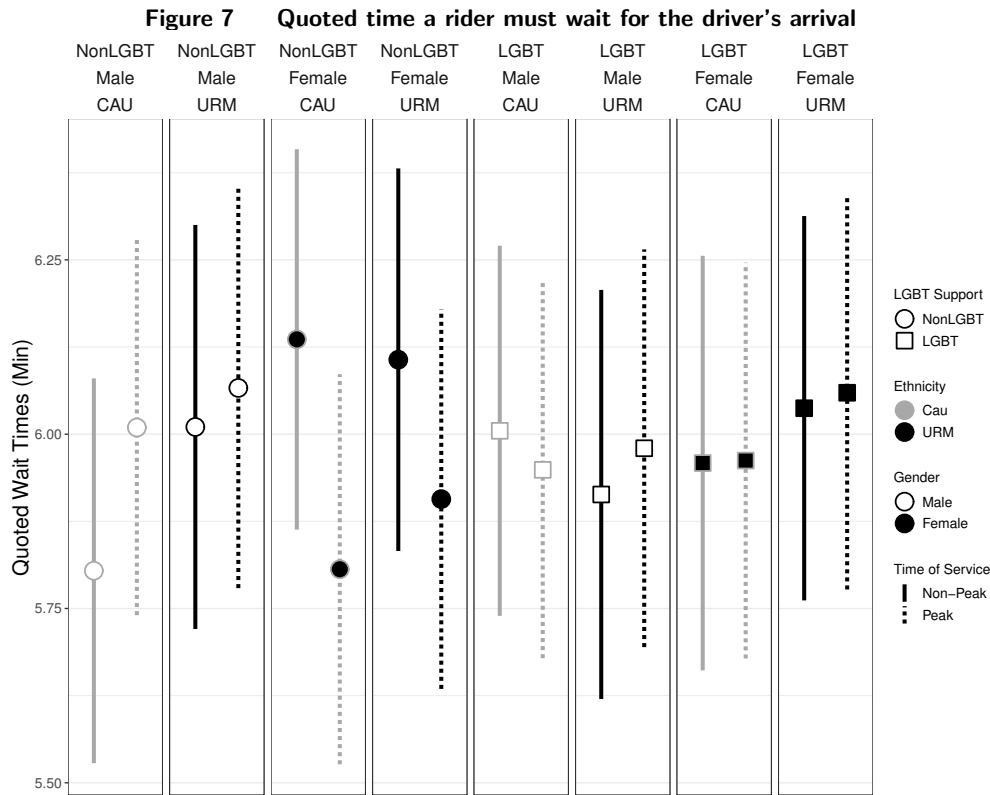
Although there is no evidence of bias against women riders, we do observe that cancellation rates are higher for URM than Caucasians during both Peak and Non-Peak times. Comparing across Peak and Non-Peak treatments, it appears that the pricing effect dominates, i.e., cancellation rates are lower during Peak times for URM. While there are still differences in cancellation rates, it appears that the pricing effect dominates the selectivity of increased demand effect.

The results also demonstrate that there is a bias for supporting the LGBT community as exhibited by higher cancellation rates during Non-Peak times. For these riders the apparent difference between Peak and Non-Peak times is not as strong as the differences observed for URM.

4.2. Regression Analysis

Having shown model-free evidence of bias against URM and LGBT-supportive riders, we now turn to a regression analysis. This analysis has two key benefits. First, it allows us to perform significance tests jointly rather than independently for each of the rider characteristics. Second, we can include day fixed effects that control for any day-to-day differences in cancellation rates. While day-to-day variation is orthogonal to the treatments, it may still introduce additional noise. Specifically, we run two logistic regressions of the form:

$$DriverCancel_i = f(\alpha + \beta_1 Female + \beta_2 URM + \beta_3 LGBTSupport + \beta_4 Peak + \delta_t + \epsilon_i) \quad (1)$$



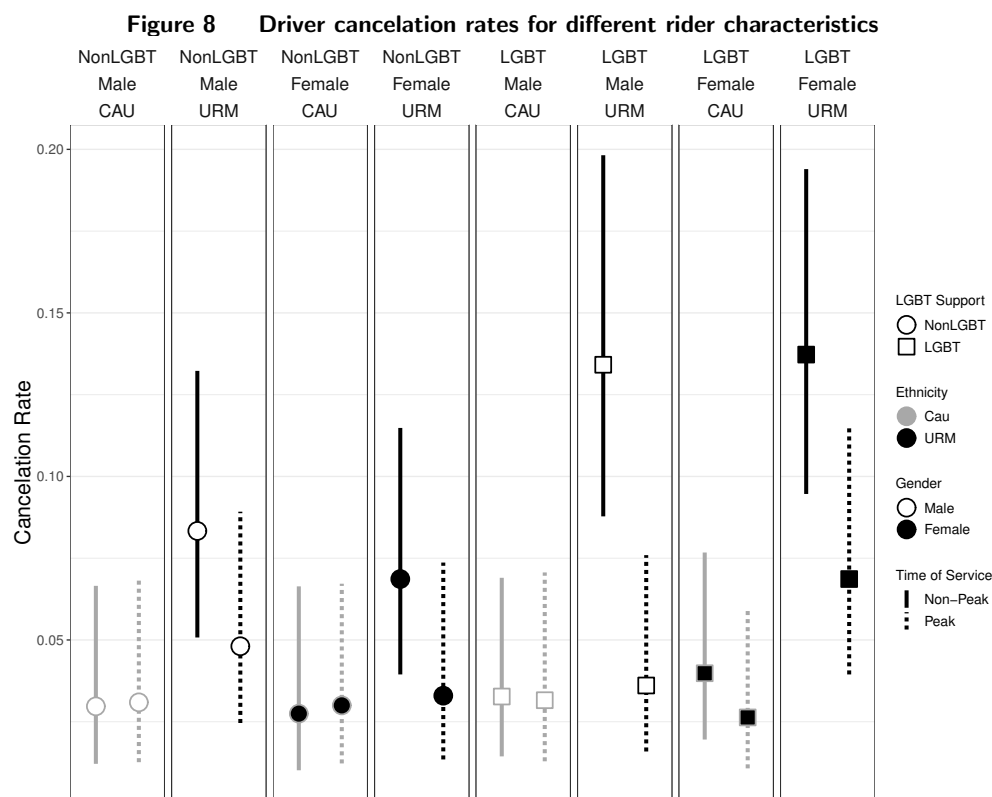
Note. *QuotedWaitTime* is statistically indistinguishable across treatments. Bars show 95% confidence intervals.

and:

$$\begin{aligned}
 DriverCancel_i = f(\alpha + \beta_1 Female + \beta_2 URM + \beta_3 LGBTSupport + \beta_4 Peak & \quad (2) \\
 + \beta_5 Female \times Peak + \beta_6 URM \times Peak + \beta_7 LGBTSupport \times Peak + \delta_t + \epsilon_i)
 \end{aligned}$$

where f is the logistic function, $Female$ is a dummy variable that is one when the rider profile is female and zero otherwise, URM is a dummy variable that is one when the rider profile is a URM and zero otherwise, $LGBTSupport$ is a dummy variable that is one when the rider profile has a rainbow filter signaling support for the LGBT community and zero otherwise, $Peak$ is a dummy variable that is one when the ride was requested during Peak hours and zero during Non-Peak hours, and δ_t are a set of date fixed effects. We focus solely on driver cancellation as the dependent variable for brevity. As would be expected, there are no significant differences on the other measures of bias presented in the previous subsection.

The results of the regressions are in Table 2. Column 1 is consistent with the results observed in the previous section: we find no evidence of bias against Females, URMs are more likely to be canceled on, and signaling support for the LGBT community leads to higher cancellation rates. There are also fewer cancellations during Peak times, which demonstrates that the financial mechanism Peak times outweighs the demand signaling mechanism. Converting the coefficients to odds ratios



Note. Cancellation rates for riders of different ethnicities in different peak times. Bars show 95% confidence intervals based on a proportions test. As such, they are not of equal length for the upper and lower bounds of the 95% confidence interval.

allows us to more meaningfully compare the magnitude of the effects. The odds ratios show that URMs are 2.60 times more likely to be canceled on and signaling support for the LGBT community results in being 1.46 times more likely to be canceled on. During Peak times, cancellations are 0.54 times as likely as during Non-Peak times.

Column 2 shows that Peak times can actually reduce the bias, but only for URMs. Signaling support for the LGBT community results in similar cancellation rates during Peak and Non-Peak times. Odds ratios for this regression show that URMs (LGBT supporters) are 3.57 (1.72) times as likely to be canceled on during Non-Peak times. During Peak times URM odds are 0.45 relative to URM during the Non-Peak times. People who signal support for the LGBT community experience an odds ratio of 0.66 during Peak times, but this is not a significant difference.⁴

Regressions interacting LGBT Support and Ethnicity treatments showed that there were no significant additive effects of being a URM and signaling support for the LGBT community. The results are omitted for brevity.

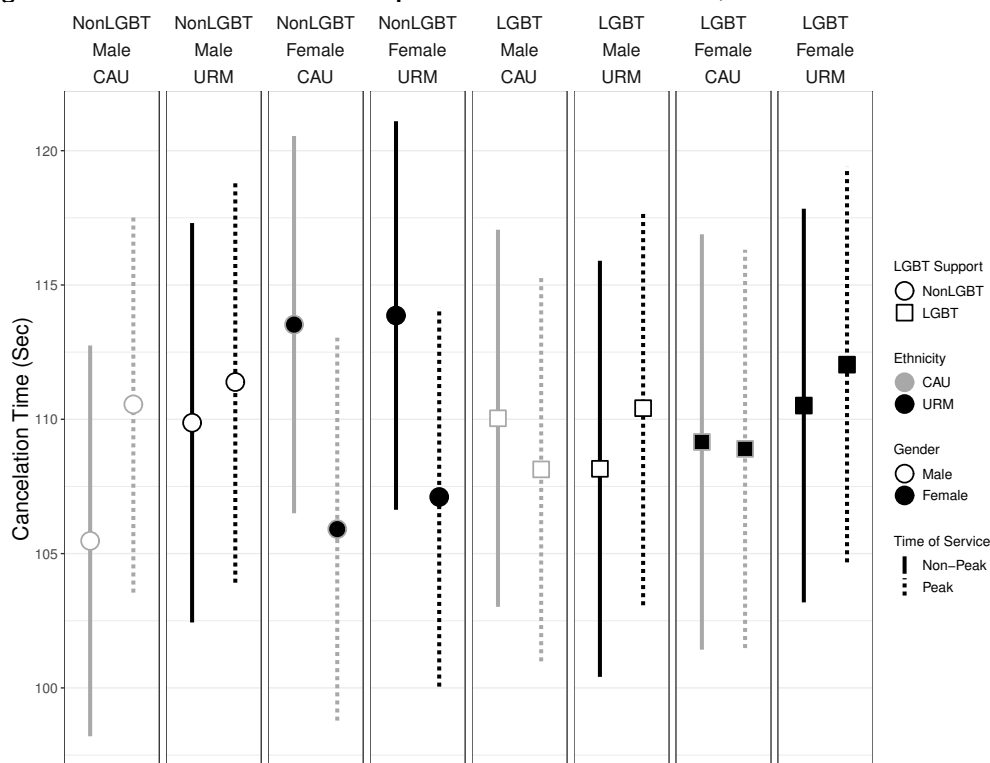
⁴ Another way of looking at this analysis would be to test whether $LGBTSupport + LGBTSupport \times Peak$ is significantly different from zero. Despite the fact that $LGBTSupport$ is significantly different from zero and $LGBTSupport \times Peak$ is not significantly different from $LGBTSupport$, the sum of the coefficients is not different from zero.

Table 2 Logistic regression analysis

	(1)	(2)
<i>Female</i>	0.025 (0.160)	-0.026 (0.202)
<i>URM</i>	0.956*** (0.172)	1.273*** (0.229)
<i>LGBTSupport</i>	0.379** (0.162)	0.545*** (0.207)
<i>Peak</i>	-0.622*** (0.165)	0.094 (0.384)
<i>Female</i> × <i>Peak</i>		0.104 (0.331)
<i>URM</i> × <i>Peak</i>		-0.788** (0.352)
<i>LGBTSupport</i> × <i>Peak</i>		-0.414 (0.334)
Constant	-3.402*** (0.199)	-3.697*** (0.256)
Date FEs	YES	YES
Observations	3,200	3,200
Log Likelihood	-635.4	-632.2
Akaike Inf. Crit.	1,280.8	1,280.4
McFadden Pseudo R ²	0.074	0.078

*** p<0.01, ** p<0.05, * p<0.1, standard errors reported in parentheses.

There are two potential concerns with this analysis. First, errors may be autocorrelated. To rule this out, we performed a Durbin Watson Test and find no evidence of autocorrelation in the error terms. The p-value for Table 2, Column 1 (Column 2) is 0.65 (0.66). Second, errors may be correlated across days but within the same randomization block. We estimated models with errors clustered at the randomization block and obtain similar results.

Figure 9 Time between driver acceptance and driver cancelation, conditional on cancelation

Note. Drivers are canceling on riders in approximately the same amount of time. Bars show 95% confidence intervals.

4.3. Are drivers different during Peak and Non-Peak hours?

One potential concern could be that the population of drivers is fundamentally different during Peak and Non-Peak hours. Unfortunately, IRB restrictions prevent us from recording any driver characteristics.⁵ One thing we can observe is the time between when a driver accepts a ride and the time at which they cancel a ride, conditional on canceling. Figure 9 shows that there are no differences across any treatments, indicating that drivers are acting in a similar manner during Peak and Non-Peak times.

To further alleviate concerns that the driver population is different, we surveyed 200 ridesharing drivers on Amazon Mechanical Turk (AMT). We attempted to limit AMT respondents to only workers who are currently or have been in the past ridesharing drivers for a specific ridesharing platform in the past 12 months in the U.S. by pre-screening participants with a questionnaire. The ten questions were selected from the on-boarding materials from one of the major ridesharing platforms in the U.S. and were reviewed by a ridesharing driver who oversees the on-boarding process for drivers for the platform in the D.C. area. We discarded 51 AMT workers who did not answer at least 9 of the 10 questions correctly, and they did not participate in our survey.

⁵ The only information we were allowed to record was a hashed driver ID which is impossible to link to any individual driver. This was solely allowed to be kept so that we could know if we were repeatedly being matched to the same driver.

Since our focus is to better understand the driver population in the ridesharing platform, we asked respondents when they drive for Uber, their attitudes towards URMs and the LGBT community, and their demographic information. LGBT attitudes are measured using the revised Attitudes Towards Lesbians and Gay Men (ATLG-R) seven-item Likert scale with high levels of internal consistency and questions that include: “Male homosexuality is a natural expression of sexuality in men” (Herek and McLemore 1998, 2011). For measuring attitudes towards URMs, we use the General Social Survey (GSS-Race) single-item question (Davis and Smith 1991), which was found to be the most reliable item to measure racial bias in a recent review of the literature (Axt 2018). We asked: “If you could find the housing that you would want and like, would you rather live in a neighborhood that is all Black; mostly Black; half Black; half White; mostly White; or all White?” Additionally, participants responded to an explicit racial preference item as recommended by Axt (2018). The questionnaire items used to measure both attitudes towards LGBT and URMs can be found in Table 3 in the Appendix.

We found that (1) 60% of respondents drive during peak times while 61% drive during non-peak, (2) 38.5% of respondents drive only during peak times while 39.5% only drive during non-peak, and (3) approximately 22% of drivers drive during both peak and non-peak times. In addition, Peak-only drivers are similar to Non-Peak-only and drivers that drive both during Peak and Non-Peak in terms of both their attitudes towards URMs and the LGBT community. In particular, there are no statistically significant differences using pairwise t-tests or ANOVA. See Figures 10 and 11 in Appendix B. In addition, we compared drivers along their demographics (i.e., gender, age, income) and found no evidence of differences across drivers. Despite finding no differences in these observable characteristics of drivers, we note that there may be other differences that lead to the cancellation patterns we observe. This is a limitation of our current analysis that should be addressed in future research.

4.4. Does the rainbow filter signal being LGBT or support for LGBT causes?

A logical question is whether the rainbow filter signals support for or membership in the LGBT community. To determine which association prevails in the broader population, we performed an additional AMT survey. We showed respondents the pictures from the different treatments with the rainbow filter (i.e., Caucasian male and female, URM male and female), and asked subjects: 1) Which of the following best describes the person in the picture? (NYC Administration for Children’s Services 2016) and 2) The person in the picture would _____ laws that would protect gay, lesbian, bisexual, and transgender people against discrimination in jobs, public accommodations, and housing (Greenberg et al. 2019). The order of the two questions was randomized. After removing respondents who do not pass the manipulation check by correctly identifying the gender or race of the person in the picture, we obtained 412 responses.

The first question is meant to determine whether respondents believe the person in the picture is a member of the LGBT community. Respondents describe sexual orientation of the person in the picture as: 42.5% LGBT, 34.0% heterosexual, and 23.5% as not sure. On the other hand, the second question determines whether the respondents believe the person in the picture would support the LGBT community. We found that subjects believe the person in the picture would 40.5% strongly favor, 36.5% favor, 12.5% be indifferent to, 6.1% oppose, and 3.9% strongly oppose these policies.

To summarize, we find that 42.5% of subjects would describe a person with a rainbow filter as LGBT. 87% of subjects described the person with a rainbow filter as expressing ‘strong support’ or ‘support’ for the LGBT community. Thus, while our field experiment cannot disentangle the reasons behind the bias in cancellation rates, it is possible that it may be due bias against the LGBT community, against the expression of support for the LGBT community, or both. We believe more work in this area is needed to better understand the mechanisms behind this type discrimination.

5. Conclusion

The efficient reallocation of resources by ridesharing platforms is leading to tremendous consumer welfare gains. Unfortunately, recent academic research and investigative journalism indicate that not all users of the platform reap the benefits in the same manner.

In an attempt remove observed driver bias, some platforms have altered their systems to reduce operational transparency. Our field experiment shows that racial biases are persistent after the change, but finds no evidence of biases related to gender. Furthermore, signaling support for the LGBT community results in higher cancellation rates. We complement this by showing that higher prices associated with Peak times can alleviate some of the bias for URM. However, this effect does not apply to the bias observed against those who signal support for the LGBT community.

We make three primary contributions. First, we extend the operational transparency literature by examining how biased actors respond to customer transparency changes, finding that altering the timing of operational transparency does not remove the bias observed in ridesharing platforms. Decisions to provide customer transparency must consider responses by service providers that may alienate minority customers or risk the unintended consequence of a seemingly simple policy change. Data-driven solutions may exist wherein rider characteristics are captured when a driver cancels and the platform penalizes the driver for the biased behavior. One possible way to punish drivers is to move them down the priority list when they exhibit biased cancellation behavior so they have fewer ride requests. Alternatively, less-punitive measures may provide “badges” for drivers that exhibit especially low cancellation rates for minority riders.

Second, we find a novel way to explore how signaling support for a social cause can impact service provision. The results suggest that, despite public opinion polls indicating support for the LGBT community is strong, there are still negative associations that can impact supporters.

Third, we demonstrate that the economic incentive delivered via dynamic pricing has a net positive effect in terms of reducing bias against URM riders on ridesharing platforms during peak times. This is an important social benefit that has not previously been included in the estimates of ridesharing platforms' benefits generally and dynamic pricing in particular.

We note that investments in reducing bias may not occur organically as ridesharing platforms are trying to maximize the number of participants in the platform (i.e., they want to attract both riders and drivers). As a result, it may be necessary for policymakers to mandate what information can be provided to a driver to ensure an unbiased experience while maintaining the safety of everyone involved or to create regulation that requires ridesharing platforms to monitor and remove drivers based on biased behavior. Careful attention should be paid to these policies both before and after implementation as unintended consequences are almost surely to follow any simple fix.

5.1. Limitations and Future Research

Our results raise many questions that should be addressed in future research. First, a limitation of our experiment is that we do not observe the full experimental design before the platform change to remove rider information at the request stage. Without pre-change data, we cannot conclusively comment on whether the change increased, decreased, or had no impact on driver bias behavior at the ride request and/or post-acceptance stages. Analyzing this change would require working with a ridesharing company to examine pre- and post-change data. This is a particularly interesting problem as while minority riders clearly prefer for no bias to exist, 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 they incur costs related to waiting for the cancellation to occur and the inconvenience of re-requesting a ride. Minority riders may be better off on the whole if rider characteristics are fully hidden until the last possible moment for rider and driver to be connected (no operational transparency), or fully visible from the ride request stage (full operational transparency). We note this with some restraint as not providing riders or drivers with any information about their counterpart may induce anxiety, which has been shown to lead to increased balking in queueing situations (Buell 2019). Nevertheless, it is a clear research question to be addressed in the future.

Second, due to IRB restrictions, we could not capture driver characteristics beyond retaining a hash that allowed us to check whether we called the same driver multiple times. This limits our analysis in that while we have no indication that drivers during peak times come from a different population, we cannot completely rule out this possibility. It would also be interesting to determine which drivers may be more likely to exhibit biased behavior and explore ways to reduce this bias. It is possible that making the bias salient (rather than implicit) for the driver could remove the

bias. Training, app notifications, and/or providing information about the drivers' cancellation rates should be explored as options for targeted bias reduction. Another potential mechanism to reduce this bias could be to determine how much income drivers are foregoing as a result of the bias. Making this cost explicit may also reduce biased behavior but would require knowledge of what happens after a driver cancels on a rider.

Third, the measure of price we observe is actually a range of prices, e.g., \$16 to \$20. We do not observe the actual price charged to the rider. This range is fairly constant over time within Peak and Non-Peak periods. The near collinearity means we cannot cleanly identify both of these variables in the same regression. While we do not believe it is likely as our survey of ridesharing drivers does not find meaningful differences across Peak and Non-Peak driver populations, it is possible that other differences about peak times lead to differences in the cancellation rates.

Fourth, we limit our analysis to one pick-up point and one destination in one U.S. city. Our intent was to minimize variation for the purpose of exploring our primary research questions. If the experiment was run in different cities, we cannot rule out that results may differ. Future research should try to understand where and why heterogeneity in bias exists and how this can be remedied.

Fifth, we document bias against someone who signals support for a social cause. In our case, this is LGBT community supporters, but it is possible that drivers do not want to give a ride to anyone who signals support for *any* social cause. Future research should determine whether/what other biases exist? Do signals of religious, political, or cultural affiliations result in drivers' biased behavior? We expect that it will, but it is unclear which social causes will result in (relatively) large or small biases.

Sixth, we do not have a large enough sample to explore passive cancellation by drivers, i.e., do drivers exhibit biased behavior through choosing either slow routes to or actively driving away from riders with certain characteristics in an attempt to get the rider to cancel themselves? If this behavior exists, it is possible that some riders are subjected to waiting longer for a ride to be canceled in addition to experiencing higher cancellation rates. Future research should explore whether this behavior exists.

Finally, our research brings up an important policy/legal question. If a ridesharing company has drivers who consistently favor a certain demographic, to what extent is the platform complicit and potentially legally responsible for the biased behavior? 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 behavior. Regulators and legal scholars should address this issue.

References

- Aker, Jenny C, Christopher Ksoll, Travis J Lybbert. 2012. Can mobile phones improve learning? Evidence from a field experiment in Niger. *American Economic Journal: Applied Economics* 4(4) 94–120. URL <http://pubs.aeaweb.org/doi/abs/10.1257/app.4.4.94>.
- Axt, Jordan R. 2018. The best way to measure explicit racial attitudes is to ask about them. *Social Psychological and Personality Science* 9(8) 896–906.
- Badgett, MV, Holming Lau, Brad Sears, Deborah Ho. 2007. Bias in the workplace: Consistent evidence of sexual orientation and gender identity discrimination .
- Bandi, Chaithanya, Antonio Moreno, Zhiji Xu. 2017. The hidden costs of dynamic pricing: Empirical evidence from online retailing in emerging markets. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2970346.
- Bendoly, Elliot, Karen Donohue, Kenneth L Schultz. 2006. Behavior in operations management: Assessing recent findings and revisiting old assumptions. *Journal of operations management* 24(6) 737–752.
- Benson, Peter L, Stuart A Karabenick, Richard M Lerner. 1976. Pretty pleases: The effects of physical attractiveness, race, and sex on receiving help. *Journal of Experimental Social Psychology* 12(5) 409–415.
- Berg, Nathan, Donald Lien. 2002. Measuring the effect of sexual orientation on income: Evidence of discrimination? *Contemporary economic policy* 20(4) 394–414.
- Bertrand, Marianne, Esther Duflo. 2017. Field experiments on discrimination. *Handbook of Economic Field Experiments* 1 309–393.
- Bertrand, Marianne, Sendhil Mullainathan. 2004. Are Emily and Greg more employable than Lakisha and Jamal? a field experiment on labor market discrimination. *The American Economic Review* 94(4) 991–1013.
- Blendon, Robert J., Logan S. Casey, John M. Benson, Justin M. Sayde, Tiffany Chan, Carolyn Miller, Jordan Reese, Dwayne Proctor, Anne Gudenkauf, Joe Neel, Keith Woods, Sara Goo, Vickie Walton-James, Luis Clemens, Alison Macadam, Alison Kodjak, Rae Ellen Bichell. 2017. Discrimination in america: Experiences and views of african americans URL <https://www.npr.org/assets/img/2017/10/23/discriminationpoll-african-americans.pdf>.
- Brown, James R, Gustav Martinsson. 2018. Does transparency stifle or facilitate innovation? *Management Science* 65(4) 1600–1623.
- Buell, Ryan. 2019. Last place aversion in queues. URL <https://ssrn.com/abstract=3090591>.
- Buell, Ryan, Ethan Porter, Michael Norton. 2019. Surfacing the submerged state: Operational transparency increases trust in and engagement with government. URL <https://ssrn.com/abstract=2349801>.
- Buell, Ryan W, Tami Kim, Chia-Jung Tsay. 2016. Creating reciprocal value through operational transparency. *Management Science* 63(6) 1673–1695.

- Buell, Ryan W, Michael I Norton. 2011. The labor illusion: How operational transparency increases perceived value. *Management Science* **57**(9) 1564–1579.
- Burtch, Gordon, Seth Carnahan, Brad N Greenwood. 2017. Can you gig it? An empirical examination of the gig-economy and entrepreneurial activity. URL <https://dx.doi.org/10.2139/ssrn.2744352>.
- Chen, Christopher, Nicos Savva. 2018. Unintended consequences of hospital regulation: The case of the hospital readmissions reduction program. URL <https://ssrn.com/abstract=3236983>.
- Chinander, Karen R, Maurice E Schweitzer. 2003. The input bias: The misuse of input information in judgments of outcomes. *Organizational Behavior and Human Decision Processes* **91**(2) 243–253.
- Choudhary, Vivek, Masha Shunko, Serguei Netessine. 2018. Does real-time feedback make you try less hard?: A study of automotive telematics. URL <http://dx.doi.org/10.2139/ssrn.3260891>.
- Cohen, Peter, Robert Hahn, Jonathan Hall, Steven Levitt, Robert Metcalfe. 2016. Using big data to estimate consumer surplus: The case of Uber. URL <http://www.nber.org/papers/w22627>.
- Cryder, Cynthia, Simona Botti, Yvetta Simonyan. 2017. The charity beauty premium: Satisfying donors versus should desires. *Journal of Marketing Research* **54**(4) 605–618.
- Cui, Ruomeng, Jun Li, Dennis J Zhang. 2016. Discrimination with incomplete information in the sharing economy: Field evidence from Airbnb URL <https://ssrn.com/abstract=2882982>.
- Cunningham, Scott, Gregory DeAngelo, John Tripp. 2017. Craigslist’s effect on violence against women. URL <http://scunning.com/craigslist50.pdf>.
- Davis, James A, Tom W Smith. 1991. *The NORC general social survey: A user’s guide*, vol. 1. SAGE publications.
- Dewey, Caitlin. 2015. More than 26 million people have changed their facebook picture to a rainbow flag. heres why that matters. *The Washington Post* URL <https://www.washingtonpost.com/news/the-intersect/wp/2015/06/29/more-than-26-million-people-have-changed-their-facebook-picture-to-a-rainbow-flag-heres-why-that-matters/>. Last accessed May 13, 2019.
- Di Caro, Martin. 2017. How many people in washington region drive for uber or lyft? it is still hard to say. *WAMU* URL <https://wamu.org/story/17/08/24/many-people-washington-region-drive-uber-lyft-still-hard-say/>. Last accessed May 13, 2019.
- Edelman, Benjamin, Michael Luca, Dan Svirsky. 2017. Racial discrimination in the sharing economy: Evidence from a field experiment. *American Economic Journal: Applied Economics* **9**(2) 1–22.
- Edelman, Benjamin G, Michael Luca. 2014. Digital discrimination: The case of airbnb.com .
- Freeman, Michael, Nicos Savva, Stefan Scholtes. 2017. Gatekeepers at work: An empirical analysis of a maternity unit. *Management Science* **63**(10) 3147–3167.
- Friedrich, James, Paul Barnes, Kathryn Chapin, Ian Dawson, Valerie Garst, David Kerr. 1999. Psychophysical numbing: When lives are valued less as the lives at risk increase. *Journal of Consumer Psychology* **8**(3) 277–299.

- Ge, Yanbo, Christopher R Knittel, Don MacKenzie, Stephen Zoepf. 2016. Racial and gender discrimination in transportation network companies.
- Gino, Francesca, Gary Pisano. 2008. Toward a theory of behavioral operations. *Manufacturing & Service Operations Management* **10**(4) 676–691. URL <https://pubsonline.informs.org/doi/abs/10.1287/msom.1070.0205>.
- Gino, Francesca, Lisa L Shu, Max H Bazerman. 2010. Nameless + harmless = blameless: When seemingly irrelevant factors influence judgment of (un) ethical behavior. *Organizational Behavior and Human Decision Processes* **111**(2) 93–101.
- Goldin, Claudia, Cecilia Rouse. 2000. Orchestrating impartiality: The impact of “blind” auditions on female musicians. *American Economic Review* **90**(4) 715–741.
- Greenberg, Daniel, Emma Beyer, Maxine Najle, Oyindamola Bola, Robert P. Jones. 2019. Americans show broad support for lgbt nondiscrimination protections. Tech. rep. URL <https://www.prrri.org/research/americans-support-protections-lgbt-people/>.
- Greenwood, Brad N., Idris Adjerid, Corey M. Angst. 2017. Race and gender bias in online ratings: An origins story. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2949511.
- Greenwood, Brad N, Sunil Wattal. 2017. Show me the way to go home: An empirical investigation of ride-sharing and alcohol related motor vehicle fatalitie. *MIS Quarterly* **41**(1).
- Hansen, Hans Krause, Mikkel Flyverbom. 2015. The politics of transparency and the calibration of knowledge in the digital age. *Organization* **22**(6) 872–889.
- Herek, Gregory M., Kevin A. McLemore. 1998. Attitudes toward lesbians and gay men scale. *Handbook of sexuality-related measures* 392–394.
- Herek, Gregory M, Kevin A McLemore. 2011. The attitudes toward lesbians and gay men (ATLG) scale.
- Joseph, W Benoy. 1982. The credibility of physically attractive communicators: A review. *Journal of advertising* **11**(3) 15–24.
- Kono, Daniel Y. 2006. Optimal obfuscation: Democracy and trade policy transparency. *American Political Science Review* **100**(3) 369–384.
- Kruger, Justin, Derrick Wirtz, Leaf Van Boven, T William Altermatt. 2004. The effort heuristic. *Journal of Experimental Social Psychology* **40**(1) 91–98.
- Libich, Jan, et al. 2006. Should monetary policy be transparent? *Policy: A Journal of Public Policy and Ideas* **22**(1) 28.
- Lurie, Nicholas H., Jayashankar M. Swaminathan. 2009. Is timely information always better? the effect of feedback frequency on decision making. *Organizational Behavior and Human Decision Processes* **108**(2) 315 – 329. doi:<https://doi.org/10.1016/j.obhdp.2008.05.005>. URL <http://www.sciencedirect.com/science/article/pii/S0749597808000745>.

- Marlowe, Cynthia M, Sandra L Schneider, Carnot E Nelson. 1996. Gender and attractiveness biases in hiring decisions: Are more experienced managers less biased? *Journal of applied psychology* **81**(1) 11.
- Mejia, Jorge, Shawn Mankad, Anand Gopal. 2015. More than just words: Latent semantic analysis, online reviews and restaurant closure. *Academy of Management Proceedings*, vol. 2015. Academy of Management, 13912.
- Milkman, Katherine L., Modupe Akinola, Dolly Chugh. 2012. Temporal distance and discrimination: An audit study in academia. *Psychological Science* **23**(7) 710–717.
- Minear, Meredith, Denise C Park. 2004. A lifespan database of adult facial stimuli. *Behavior Research Methods, Instruments, & Computers* **36**(4) 630–633.
- Mohan, Bhavya, Ryan Buell, Leslie John. 2018. Lifting the veil: The benefits of cost transparency. URL <http://dx.doi.org/10.2139/ssrn.2498174>.
- Mohr, Lois A, Mary Jo Bitner. 1995. The role of employee effort in satisfaction with service transactions. *Journal of Business Research* **32**(3) 239–252.
- Morales, Andrea C. 2005. Giving firms an “e” for effort: Consumer responses to high-effort firms. *Journal of Consumer Research* **31**(4) 806–812.
- Moritz, Brent B., Arunchalam Narayanan, Chris Parker. 2019. Unraveling behavioral ordering: Relative costs and the bullwhip effect: An experimental study. URL <https://ssrn.com/abstract=3195282>.
- Movement Advancement Project, Family Equality Council, and Center for American Progress. 2011. All children matter: How legal and social inequalities hurt LGBT families. Tech. rep. URL <https://www.lgbtmap.org/file/all-children-matter-full-report.pdf>.
- Narayanan, Arunachalam, Brent B. Moritz. 2015. Decision making and cognition in multi-echelon supply chains: An experimental study. *Production and Operations Management* **24**(8) 1216–1234. doi:10.1111/poms.12343. URL <http://dx.doi.org/10.1111/poms.12343>.
- NYC Administration for Children’s Services. 2016. Respectfully asking sexual orientation and gender identity (sogi) questions. Tech. rep. URL https://www1.nyc.gov/assets/acs/pdf/lgbtq/Respectfully_Asking_SOGI_Questions.pdf.
- Oldham, Greg R, J Richard Hackman. 1978. Work design in the organizational context. Tech. rep., YALE UNIV NEW HAVEN CONN SCHOOL OF ORGANIZATION AND MANAGEMENT.
- Ozeren, Emir. 2014. Sexual orientation discrimination in the workplace: A systematic review of literature. *Procedia - Social and Behavioral Sciences* **109**(Supplement C) 1203 – 1215. URL <http://www.sciencedirect.com/science/article/pii/S187704281305252X>.
- Parker, Chris, Kamalini Ramdas, Nicos Savva. 2016. Is IT enough? Evidence from a natural experiment in India’s agriculture markets. *Management Science* **62**(9) 2481–2503. URL <https://doi.org/10.1287/mnsc.2015.2270>.

- Pope, Devin G, Justin R Sydnor. 2011a. Implementing anti-discrimination policies in statistical profiling models. *American Economic Journal: Economic Policy* **3**(3) 206–31.
- Pope, Devin G, Justin R Sydnor. 2011b. Whats in a picture? evidence of discrimination from prosper. com. *Journal of Human resources* **46**(1) 53–92.
- Rich, Steven, Ted Mellnik, Kimbriell Kelly, Wesley Lowery. 2018. Murder with impunity: Where murders go unsolved. *The Washington Post* URL <https://www.washingtonpost.com/graphics/2018/investigations/unsolved-homicide-database/>. Last accessed May 13, 2019.
- Roberts, John. 2009. No one is perfect: The limits of transparency and an ethic for intelligent accountability. *Accounting, Organizations and Society* **34**(8) 957–970.
- Small, Deborah A, George Loewenstein. 2005. The devil you know: The effects of identifiability on punishment. *Journal of Behavioral Decision Making* **18**(5) 311–318.
- Stiglitz, Joseph E. 2002. Information and the change in the paradigm in economics. *American Economic Review* **92**(3) 460–501.
- Stiglitz, Joseph E. 2003. Information and the change in the paradigm in economics, part 1. *The American Economist* **47**(2) 6–26.
- Strathern, Marilyn. 2000. The tyranny of transparency. *British educational research journal* **26**(3) 309–321.
- Tilcsik, Andrs. 2011. Pride and prejudice: Employment discrimination against openly gay men in the united states. *American Journal of Sociology* **117**(2) 586–626. URL <http://www.jstor.org/stable/10.1086/661653>.
- US Department of Housing and Urban Development. 2013. Housing discrimination against racial and ethnic minorities 2012. Tech. rep. URL https://www.huduser.gov/portal/Publications/pdf/HUD-514_HDS2012.pdf.
- Weichselbaumer, Doris. 2003. Sexual orientation discrimination in hiring. *Labour Economics* **10**(6) 629–642.
- Younkin, Peter, Venkat Kuppuswamy. 2017. The colorblind crowd? founder race and performance in crowd-funding. *Management Science* .
- Yu, Jessica. 2018. Search, selectivity, and market thickness in two-sided markets.
- Zhang, Dennis J., Hengchen Dai, Lingxiu Dong, Fangfang Qi, Nannan Zhang, Xiaofei Liu, Zhongyi Liu. 2017. How does dynamic pricing affect customer behavior on retailing platforms? Evidence from a large randomized experiment on Alibaba. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3029707.

Appendix

A. AMT Questionnaire Details

Table 3 Questions asked in the AMT driver questionnaire

Panel A: Attitudes toward Gay and Lesbian people			
Item	Question (1 = strongly disagree to 7 = strongly agree)	Scale	Adapted from
1	Sex between two men is just plain wrong.	ATG-S5	Herek and McLemore 2011
2	I think male homosexuals are disgusting.	ATG-S5	Herek and McLemore 2011
3	Male homosexuality is a natural expression of sexuality in men. (Reverse scored)	ATG-S5	Herek and McLemore 2011
4	Male homosexuality is a perversion.	ATG-S5	Herek and McLemore 2011
5	Male homosexuality is merely a different kind of lifestyle that should not be condemned. (Reverse scored)	ATG-S5	Herek and McLemore 2011
6	Sex between two women is just plain wrong.	ATL-S5	Herek and McLemore 2011
7	I think lesbians are disgusting.	ATL-S5	Herek and McLemore 2011
8	Female homosexuality is a natural expression of sexuality in women. (Reverse scored)	ATL-S5	Herek and McLemore 2011
9	Female homosexuality is a perversion.	ATL-S5	Herek and McLemore 2011
10	Female homosexuality is merely a different kind of lifestyle that should not be condemned. (Reverse scored)	ATL-S5	Herek and McLemore 2011
Panel B: Attitudes toward racial minorities			
Item	Question	Scale	Adapted from
1	If you could find the housing that you would want and like, would you rather live in a neighborhood that is (1 = all Black, 7 = all White)	GSS-Race	Davis and Smith (1991)
2	Which statement best describes you?(1 = "I strongly prefer African Americans to European Americans", 7 = "I strongly prefer African Americans to European Americans").	ExpPref	Axt (2018)

B. Driver Attitudes Towards URM and the LGBT Community

Figure 10 Peak and Non-Peak drivers are similar in terms of attitude towards URM.

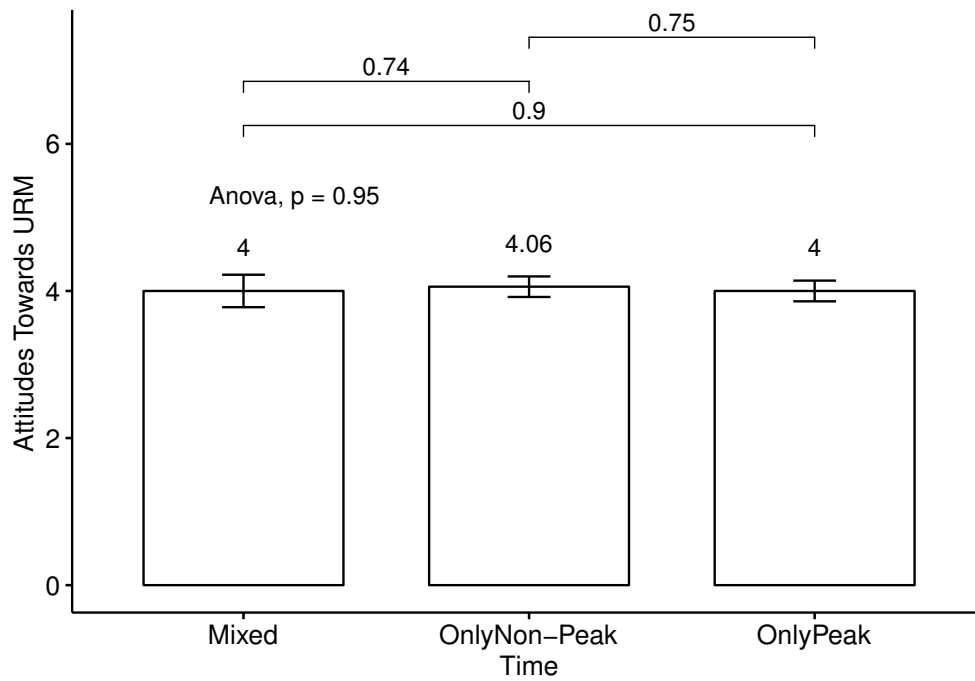


Figure 11 Peak and Non-Peak drivers are similar in terms of attitude towards the LGBT community.

