

Tjaden, Jasper Dag; Schwemmer, Carsten; Khadjavi, Menusch

Article — Accepted Manuscript (Postprint)

Ride with Me - Ethnic Discrimination, Social Markets and the Sharing Economy

European Sociological Review

Suggested Citation: Tjaden, Jasper Dag; Schwemmer, Carsten; Khadjavi, Menusch (2018) : Ride with Me - Ethnic Discrimination, Social Markets and the Sharing Economy, European Sociological Review, ISSN 1468-2672, Oxford University Press, Oxford, Vol. 34, Iss. 4, pp. 418-432, <https://doi.org/10.1093/esr/jcy024>

This Version is available at:

<https://hdl.handle.net/10419/232061>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Ride with Me - Ethnic Discrimination, Social Markets and the Sharing Economy

Jasper Dag Tjaden¹, Carsten Schwemmer² & Menusch Khadjavi³

Abstract

We study ethnic discrimination in the sharing economy using the example of online carpooling marketplaces. Based on a unique dataset of 16,624 real rides from Germany, we estimate the effects of drivers' perceived name origins on the demand for rides. The results show sizable ethnic discrimination. Further analyses suggest that additional information about actors in this market decreases the magnitude of ethnic discrimination. Our findings broaden the perspective of ethnic discrimination by shedding light on subtle, everyday forms of discrimination in social markets; inform ongoing discussions about ways to address discrimination in an era in which markets gradually move online; and respond to increasingly recognized limitations of experimental approaches to study discrimination.

Keywords

Ethnic discrimination; sharing economy, statistical discrimination; online markets; computational social science

* We are grateful for comments on previous versions from Uri Gneezy, Ruud Koopmans, Jack DeWaard, Jeylan Mortimer, Marc Helbling, Sebastian Wenz and Cornelia Kristen.

¹ Corresponding author: Global Migration Data Analysis Centre, International Organisation for Migration, Taubenstraße 20-22, 10117 Berlin, Germany

² Political Sociology Department, Otto-Friedrich-University Bamberg, Germany

³ Kiel Institute for the World Economy and Department of Economics, Christian Albrechts-University Kiel, Germany

1. Introduction

Decades of social science research provide evidence of ethnic and racial discrimination in various areas of society and in numerous countries (e.g. Pager, 2007; Pager & Shepherd, 2008; Rich, 2014). Despite a long history of policy responses and the introduction of anti-discrimination legislation designed to attenuate ethnic and racial disparities, discrimination appears to persist (Pager et al., 2009). Discrimination studies continue to enjoy attention as discrimination is seen as one of the key mechanisms for explaining enduring economic and social inequality in society. Online markets offer a new perspective on the diverse settings in which ethnic discrimination can occur and provide new channels to test assumptions about why and how members of ethnic or racial groups are being discriminated against (e.g. Zussman, 2013; Edelman et al., 2017).

We join this effort by examining the extent and the drivers of ethnic discrimination in Europe's largest online carpooling market. We compile a new dataset of more than 16,624 carpooling rides offered in Germany by programmatically collecting ride information from an online platform. We estimate the effect of drivers' perceived name origin on the demand for their offered rides (clicks on offer). In order to group names into perceived ethnic backgrounds, we conducted a separate online survey with 1,577 participants who rated a total of 1,381 unique first names to distinguish the associated origin of drivers. Participants distinguished between typically German names and names with an Arab, Turkish or Persian origin. The latter group is the largest and most recognizable immigrant community in Germany.ⁱ Previous studies found that this particular group is disproportionately affected by discrimination (e.g. Blommaert et al., 2014; Diel et al., 2013).

In carpooling markets, private individuals use online platforms to offer seats in their car for a particular ride. Carpooling websites have become serious competitors for conventional bus and train providers across Europe, in particular in low-budget segments of the transportation

market.ⁱⁱ Carpooling offers mid- to long distance rides from city to city rather than short taxi rides within cities (like services such as Uber).

Carpooling platforms are a compelling application for ethnic discrimination studies for several reasons. First, carpooling is not only an economic market where riders select drivers with the best economic value (e.g. price per distance). Carpooling is also a social market. The decision to acquire the service is linked to spending one-off time with a stranger, i.e. the driver. The element of face-to-face personal interaction in a non-professional setting distinguishes carpooling platforms from labor or consumer good markets where ethnic discrimination has been studied (e.g. Doleac & Stein, 2013; Ayres et al., 2015, Ewens et al., 2014). Carpooling may thus help to draw attention to ethnic discrimination in social situations and reveal subtle, everyday forms of discrimination that may otherwise go unnoticed.

Second, online markets such as carpooling are ideally suited to isolate ethnic effects. We are able to observe all characteristics that are visible to the customer including the driver rating, experience, the car et cetera. This setup allows us to ameliorate issues related to both experimental designs – because we do not introduce an artificial treatment that would otherwise not occur in this wayⁱⁱⁱ – and many observational studies which may suffer from omitted variable bias.^{iv} For example, audit studies have often been criticized for introducing additional unobserved factors such as demeanor and socioeconomic background that may ‘pollute’ the treatment. Our analysis controls for every signal available to the consumer.

Third, we exploit variation in the information about drivers to test assumptions about the mechanisms driving ethnic discrimination. Additional information about the driver – such as the rating and experience - could work as a trust signal for consumers (e.g. Abrahao et al., 2017), however, previous evidence is unclear about whether additional information offsets ethnic discrimination (Ahmed et al., 2010; Nunley et al., 2011).

Fourth, we provide tentative evidence that ethnic discrimination effects are not driven by social class bias – a common dilemma given that many ethnic groups in Germany are overrepresented among lower social classes.

Our results indicate large ethnic discrimination effects. Controlling for all observable information, drivers with an Arab/Turkish/Persian sounding name attract significantly less interest in their offers (fewer clicks on the offer) than drivers with typically German names. To achieve the same demand compared to a driver with a typical German name, the average driver with an Arab/Turkish/Persian sounding name would have to offer their ride at 32% less than the price for an average ride. Group differences cannot be explained by any other observable characteristic associated with the driver or the offered ride and are robust against a series of robustness checks.

Consumers appear to use the name as a proxy signal to infer the ‘true’ value of the ride in economic, safety and social terms. When rich information about the driver is available (i.e. high rating, profile picture), ethnic discrimination decreases, as consumers rely less on the name.

Our findings have important implications for policy. First, ethnic discrimination occurs in social online market platforms. This expands the view from traditional discrimination studies in the labor market and housing to more subtle, everyday forms of unequal treatment. The results draw attention to other sectors with stronger social interaction elements, including the service and care sector or group environments such as membership in clubs, associations and interest groups. Second, insights into the mechanisms of discrimination can be the starting point for policy design aimed at reducing disparities (e.g. Guryan & Charles, 2013; Nunley et al., 2011). Our results suggest that providing more relevant context information about market actors may be a powerful strategy to reduce discrimination effects. As such, our results inform the discussion around the need for anti-discrimination efforts in markets that increasingly operate online. While consumers and service providers are often protected against discrimination in

traditional, offline markets (for instance hiring, housing, hospitality and consumer goods), similar provisions do not exist in online markets (see Edelman et al., 2017) and are difficult to prosecute.

2. Evidence and Mechanisms of Ethnic Discrimination

Ethnic and racial discrimination can be defined as differential treatment that leads to unequal outcomes based entirely on ascribed features such as race, ethnic background, name origin, foreign appearance etc. (Blank et al., 2004).

Recent reviews document discrimination effects in employment, housing, credit and commodity markets in many countries (Pager & Shepherd, 2008; Rich, 2014). The strongest evidence for ethnic discrimination is based on studies employing experimental designs (Rich, 2014). Studies show that racial or ethnic groups often are – *ceteris paribus* – disadvantaged in terms of access to labor market (interview invitations, call back rates, wage offers, treatment in interviews) and the housing market (renting, buying or selling apartments and houses).

More recent studies have made advances in two ways: First, they have broadened the application of discrimination studies to other markets (Bryson & Chevalier, 2015; Doleac & Stein, 2013; Edelman et al., 2017; Gneezy et al., 2012; Nunley et al., 2011; Zussman, 2013). Second, researchers have fine-tuned experiments to test hypotheses about why discrimination occurs opposed to whether it occurs (Gneezy et al. 2012; Guryan & Charles, 2013).

In terms of relevant mechanisms, much of the literature across the various domains traditionally attempts to discern whether discrimination stems primarily from taste-based discrimination (racial animus/ prejudice) or from statistical discrimination (asymmetric information).

In the case of taste-based discrimination (Becker, 1971), the driver of discrimination is a negative disposition towards certain groups. In our case, an individual may suffer ‘disutility’

resulting from contact with a specific ethnic group. As such, taste-based discrimination relies on the presence of prejudice. Prejudice in return can loosely be defined as an affective, mostly unfavorable feeling toward a person or group member based solely on their group membership.

In the case of statistical discrimination, differential treatment based on race and ethnic background arises from incomplete or asymmetric information about the productivity of actors (Arrow, 1973; Phelps, 1972).^v When limited information about a product or an individual is available, agents rely on observable group characteristics (such as ethnic group, race) to make inferences about the individual. Another class of statistical discrimination models focuses on the reliability of the information that employers have about individual productivity (Aigner & Cain, 1977; Altonji & Blank, 1999). At the core of both of these strands of statistical discrimination is the notion that a lack of information leads the employer to treat individuals as members of groups (Guryan & Charles, 2013). As assumptions about group characteristics – for example, lower productivity – are applied to the individual, empirical accuracy of the group stereotype is not a necessary requirement for discrimination to affect members of a group (e.g. Kalter, 2006).

Past evidence on the dominant form of ethnic discrimination remained inconclusive as support for either mechanism varies considerably across studies (Ewens et al., 2014). However, recent (experimental) studies point to the importance of statistical discrimination (rather than tastes) for explaining why discrimination persists (Altonji & Pierret, 2001; Bryson & Chevalier, 2015; Ewens et al., 2014; Kaas & Manger 2012; List, 2004; Zussman, 2013). Growing evidence in favor of statistical discrimination may be good news for policy makers given that information asymmetries may more easily be addressed than deep-rooted prejudice. In practice, statistical and taste-based discrimination are difficult to isolate in experimental and non-experimental study designs. Similar to taste-based discrimination, statistical discrimination relies on the concept of stereotypes in the form of certain beliefs associated with a group. It is not clear where

stereotypes originate and whether they are related or congruent with the concept of prejudice. This complicates interpreting evidence for one or the other.

Our study avoids a framing of statistical vs. taste-based discrimination as mutually exclusive mechanisms. We examine the role of information on its own merit while we are aware that previous research has interpreted information effects as indicative of statistical discrimination (e.g. Nunley et al., 2011). Rather than discerning the origins of ethnic discrimination, our primary aim with this study is to examine how additional information affects ethnic discrimination levels. In addition, responding to critiques of experimental studies (e.g. Heckman 1998), we take advantage of rich observational data covering real interactions occurring in a real market.

3. Ethnic Discrimination in Online Carpooling

Our study joins the effort of leveraging online markets – in our case, the largest German online carpooling market – for the study of ethnic discrimination. Carpooling markets match drivers that offer available seats in their car to riders that look for affordable one-off transport between cities. Riders can search rides by departure/ arrival town and date. Besides the place, day and time of departure, the price for a seat and a number of other ride-specific characteristics, carpooling offers show the first names of the drivers. We estimate the effect of an Arab/Persian/Turkish sounding first name on the demand of offered rides as measured by clicks. Arab/Turkish/Persian sounding names are associated with the largest and most recognizable immigrant community in Germany (mostly descendants of low skilled guest workers that arrived since the 1960s). Previous studies have highlighted that members of the Arab/Turkish/Persian community appear disproportionately affected by discrimination (Blommaert et al., 2014; Diehl et al., 2013) in Europe.

We interpret group differences in clicks net off all observable characteristics of the driver and the offered ride as evidence for ethnic discrimination.

In this brief section, we formulate three different possible outcomes to our research question of whether and why ethnic discrimination exists in carpooling markets: (1) no distinct discrimination effects, (2) discrimination effects *unaffected by information* and (3) discrimination effects *sensitive to the information provided*.

Regarding (1) we propose that it is plausible to expect no distinct ethnic discrimination effect owing to the particular context of carpooling and our study design. First, carpooling consumers are on average younger than the general population (Destatis, 2017). Second, rides provide transport between urban centers which suggest that the customers are also more likely to live in urban areas. Third, sharing a ride with a stranger already requires a certain level of trust. Fourth, online market platforms have been shown to reduce information asymmetries associated with productivity and correct biases against certain groups (Agrawal et al., 2013). Accordingly, it would *not* be surprising to find *no* significant discrimination effects given that we control for all observable information about the ride and the driver. In this light any effect that we may find is likely conservatively small compared to ethnic discrimination in other contexts and with other sub-populations in the German society.

Regarding (2), in our setting, traditional taste-based ethnic discrimination approaches suggest that potential consumers discriminate against drivers with a foreign name, because they simply wish to avoid contact with a member of a specific ethnic group. In carpooling, this means that we would expect discrimination effects regardless of variation in other information about the driver. Compared to commodity markets, taste-based ethnic discrimination may be more pronounced in carpooling as the customer is spending several hours with someone from another ethnic group in a narrow space (a car). In this case, simply the fact that the driver is associated with another ethnic group should lead to unequal treatment regardless of other observable

characteristics of the ride or the driver. In other words, we would expect that variation in the information provided about the ‘quality’ of the driver should not affect ethnic discrimination.

Regarding (3), in our setting, statistical discrimination approaches commonly assume that potential riders use the name of a driver as a signal to infer the ‘true value’ of the ride. Following this approach, the value of the ride depends on the degree of provided information about a ride rather than exclusively on the name origin of the driver. One advantage of our large dataset is that we can test for different stereotypes. We hypothesize that a negative ethnic effect on clicks could generally be driven by three different sets of considerations: 1) price and comfort relative to distance, 2) personal safety and 3) the social value.

Based on a narrow economic perspective, consumers simply click on the ride that offers the cheapest price relative to the distance travelled. Other factors may include the car quality as an indication how fast and comfortable the ride will be. In our setting, we control for the distance of the ride, the price and the car comfort.

Other consumers may choose an offer based on how secure they perceive the ride. Security has to be inferred from other available information as there is no objective indicator of security and safety on the platform. We assume that consumers use the name of the driver as two signals for perceived safety of a ride. First, it is a common stereotype that ‘migrants’, especially males, commit more crimes (e.g. Fitzgerald et al., 2011, Trager et al., 2014).^{vi} The other common safety-related stereotype could be that foreigners drive less safely because traffic regulations are less strict or less enforced in their origin countries. To the best of our knowledge there is virtually no reliable comparable data to prove or disprove this stereotype, but surveys suggest that the stereotype exists.^{vii}

Lastly, the value of a particular ride (and as a result, demand for that ride) may be driven by the desire for pleasant social interaction. We know from previous research that certain ethnic

groups are disadvantaged, for example, in flat sharing markets – a market where choices include social interaction (e.g. Przepiorka, 2011). Consumers may click on those offers that suggest the most enjoyable time during the ride. Sharing a ride means sharing private space as car-poolers sit in close proximity. Again, the name of the driver could be a proxy for language. Pleasant conversation is less likely if the driver speaks a different language and possibly listens to ‘foreign’ music. We estimate an interaction effect with music and dialog preferences to test this assumption. Studies on online dating have shown that clear ethnic/racial preferences exist that commonly disadvantage minorities (Jakobsson & Lindholm, 2014, Lin & Lundquist, 2013; Robnett & Feliciano, 2011). Similar to those markets, consumers in carpooling markets may be driven by homophily preferences, i.e. looking to meet drivers who are most like them (see McPherson et al., 2001). Again, the name would signal greater social distance given that the large majority of consumers are Germans.

As there is no direct indicator for safety or sociability, consumers have to rely on other available information, including the name. We argue that the user rating, number of ratings and the driver experience are suitable aggregate proxies for both categories. A bad user rating or low experience suggests that the ride may be less safe and less pleasant. Similar to studies that attempt to test statistical discrimination, we will interact the ethnic indicator with other indicators about productivity signals, in this case, the user rating and experience (Blommaert et al., 2014; Ewens et al., 2014; Nunley et al., 2011). Similar to Nunley et al. (2011), we argue that consumers’ relative weight on beliefs regarding the trustworthiness of drivers with an ethnically distinct name diminishes as other pertinent information about credibility becomes available. The scarcer other information about the ‘true’ safety and ‘fun’ of the ride, the more consumers rely on stereotypes regarding the perceived name origin. In our design, the user rating is based on experience of the driver (how many rides he or she has offered in the past) and the customer satisfaction. The user rating is a strong signal about the trustworthiness of

provided information online and the ‘true’ productivity of the ride. We use additional proxies of safety and sociability to test the effect of additional information including the profile picture (homophily, trust), talking and music preferences (sociability) and gender (safety).^{viii}

4. Data & Methods

In this section we first present details of the data collection process and then elaborate on the empirical methods.

4.1 Data Collection

We compile a new dataset with the aim to achieve a meaningful balance between internal and external validity of discrimination effects. Using one of the largest online carpooling platforms in Germany, we compile a dataset of 17,294 observations (i.e. rides) in Germany that were listed online between 16 July 2015 and 27 July 2015.^{ix} According to the provider, the platform offered 250,000 rides in 2013 and 2014. The platform has 30 million members in 22 countries. According to the company’s website, 10 million users use the website every quarter. Based on access to an Application Programming Interface, in short API, we collected information on all observable information on the offered rides and the drivers. The visual interface (see Figure 1) shows information about age, gender, user picture if available, user rating, car, timing and stops of the ride, price, available seats and some preferences of the driver (smoking, music, talking).

Figure 1

Rides shown in Figure 1 are offered by drivers traveling from Rostock to Berlin. For instance, the first driver offered a seat for his ride at the price of 11 Euros. He has two positive user ratings from earlier interactions on the market. Furthermore, he prefers to talk during the trip

and does not mind riders to travel with their pets. In contrast to the other ride offers in Figure 1, the driver did not upload a picture.

Due to volume and restrictions from the provider, not all rides in the market could be collected. As a result, we selected routes between the largest cities in all 16 German states (Länder). Carpooling is more common between urban centers due to supply and demand for rides. Our strategy allowed us to approximate a balanced geographical representation of observed rides for different regions in Germany. As a second step, we included additional large cities in regions with larger populations, particularly regions with higher shares of ethnic minority residents. Oversampling of urban areas was necessary to ensure a sufficient sample of drivers with foreign-sounding names (see Figure 2).

Figure 2

The data was collected in two steps: First, we pulled data from the API and, second, we used the programming language Python to automatically access the website four times a day at equal intervals between 16 July 2015 and 27 July 2015. Accessing the website in addition to pulling API data was necessary for retrieving information about our main dependent variable, the number of clicks that each ride in our sample received.

4.2 Measurements

We assess group differences regarding demand for offered rides on the carpooling platform. Demand is measured by clicks. We regress the maximum number of clicks that a ride received until departure using a negative binomial regression (see Figure B1 in the Appendix for distribution of clicks).^x In the analysis the number of clicks is adjusted for the number of days that the offer was displayed online until departure.

We restrict our sample in several ways: First, we drop rides that depart after the end of our observation period (right-censoring). This is necessary to obtain an accurate measure of

maximum clicks before departure. Second, we restrict the sample to rides that were uploaded no sooner than seven days before we began the data collection. One week is an appropriate time window given that most rides are uploaded a few days before departure. Third, we drop cross-border rides, as consumers are most likely not German and clicks are inflated as affected rides are also listed in carpooling platforms for neighboring countries, where the provider also operates. Fourth, we limit our sample to routes (e.g. departure city: Munich – arrival city: Berlin) that have more than one offered ride per day and have at least one driver with an Arab/Turkish/Persian name.^{xi} This step is important to ensure that we can observe a counterfactual, i.e. consumers cannot discriminate against drivers from another ethnic group if there are none. It is important to note that all our models additionally control for route and volume (number of offered rides per route and day). The final sample for the analysis of clicks includes 16,624 rides, including 528 rides with an Arab/Turkish/Persian driver.

As our main independent variable regarding discrimination, we use the first name of the driver to infer whether the name is ‘typically Arab/Turkish/Persian’ or ‘typically German’. Names signal membership to particular ethnic group (regardless of whether the signal is true) and ‘ignite’ potential stereotypes (e.g. Bertrand & Mullainathan, 2004, Booth et al., 2012).

Driven by concerns about the objectiveness and reliability of name ratings, we conducted a large online survey in which respondents were asked to rate driver names that we extracted from our carpooling sample. In total, 1,577 student raters participated in the survey. The origin of 1,381 unique first names were on average rated by 20 student raters ($SD=4.6$).^{xii} As carpooling riders are younger than the national average (Destatis, 2017), students represent an appropriate approximation of typical riders. Table 1 shows the most frequent names by perceived name origin.

Table 1

For the analysis, we use an 80% cut-off for determining an Arab/Turkish/Persian name. That is, the driver is considered to have an Arab/Turkish/Persian name if four out of five raters (i.e. 16 out of 20 raters per name on average) considered the name to be typically “Arab, Turkish or Persian”. We also report results for the continuous measure of Arab/Turkish/Persian name origin variable in percentage points (see Figure 4 below). We grouped Arab/Turkish/Persian sounding first names together because they are difficult to distinguish for the average resident in Germany and are commonly associated as being from the same world region.^{xiii} Members of this broad group are associated with the largest and most recognizable immigrant community in Germany (mostly descendants of low skilled guest workers that arrived since the 1960s). Previous studies have highlighted that members of the Arab/Turkish/Persian community appear disproportionately affected by discrimination (Blommaert et al., 2014; Diehl et al., 2013).

For the analysis of mechanisms, we exploit variation in information about each offered ride. We estimate the interaction effect of an Arab/Turkish/Persian first name with the user rating associated with each offered ride. The rating score is an average of past riders’ general evaluation of their ride with the respective driver who offered it. Past riders can rate the driver retrospectively. As a result, the rating is a strong signal of safety, sociability and overall trustworthiness. To test the sociability argument in particular, we estimate an interaction of the name with the profile picture and the music and dialog preference. Drivers that provide a profile picture as personal information likely increase their trustworthiness. The profile picture may also function as a proxy sociability perceptions. The ‘talking preference’ indicates whether or not the driver is interested in talking during the ride which we use as one additional measurement for the sociability argument. Finally, we estimate an interaction with name and gender. We assume that negative stereotypes regarding the safety of rides with Arab/Turkish/Persian drivers largely apply to male drivers.

Controls include all the information that is observable to consumers, including information about the offered ride (route, time, distance) and the driver such as age and gender (see Table A2 and Table A3 in the Appendix for a full description and distribution of all model variables). Regardless of ride and driver information, the clicks on offer may simply be driven by the size of the potential user population which varies considerably across the sampled cities and regions in Germany. For this reason, we control for the demand side using a route identifier for all routes in our sample.

5. Results

The analysis confirms substantial discrimination effects in Germany's online carpooling market. Drivers with Arab/Turkish/Persian names attract less demand (measured in clicks) than drivers with typical German names for the same ride. Controlling for all observable characteristics of the ride and the driver that are visible to consumers, we find that drivers with an Arab/Turkish/Persian sounding name obtain on average 7 clicks less than a driver with a typical German name (significant at $p < 0.01$, see Figure 3 and Table A3 the Appendix). 7 clicks represent approximately 13% of the average number of clicks per offered ride in the sample (51 clicks). In a separate step, we calculate the average discriminatory price premium, i.e. the average willingness to pay to avoid riding with an Arab/Turkish/Persian driver. Dividing the name coefficient by the price coefficient indicates that Arab/Turkish/Persian drivers would have to offer their rides on average 4.20 € cheaper than German drivers to achieve the same number of clicks. This accounts for 32% of the average price of an average ride in our sample. This discriminatory price premium increases to 34% when setting covariates to different values, for example a male, thirty-year-old Arab/Turkish/Persian driver with little experience offering a 300 kilometer ride in a comfortable car on a Sunday afternoon.

Figure 3 uses the 80% cut-off to determine drivers with an Arab/Turkish/Persian name. Figure 4 reports the result for the continuous measure of Arab/Turkish/Persian name origin (the percentage of survey respondents who rated the name to be typically Arab/Turkish/Persian). Disparities between Arab/Turkish/Persian and German drivers increase with the degree of certainty that the name is associated with an Arab/Turkish/Persian background (see Figure 4).

Figure 4

In the second part of the analysis, we turn to the role of information for ethnic discrimination. Our tests suggest that disparities between groups depend on the variation of information about individual drivers.

Figure 5 shows that drivers with an Arab/Turkish/Persian name are disadvantaged against German drivers when they have no or low ratings. Disparities appear to vanish when both drivers have equally high user ratings. The interaction effect is statistically significant at $p < 0.01$ (see Table A3). Similar results for the number of user ratings and the driver experience corroborate these findings (see Figure B2 and Figure B3 in the Appendix).

Figure 5

The findings for the user rating, number of ratings and experience suggest that consumers place less weight on the name of the driver when other relevant information is available. This is consistent with the statistical discrimination hypothesis as the name of the driver may be used as one more source of information about the perceived 'value' of the ride.

We conduct a number of additional indirect tests to approximate different motives (as discussed in section 2.1). There may be at least two possible sources of stereotypes. First, consumers may be concerned with safety since carpooling entails sharing a ride with a stranger. Given stereotypes regarding crime and driving styles among 'foreigners', we hypothesized that statistical discrimination may be driven by safety concerns. Our results provide indirect

evidence that this is the case. First, ethnic discrimination effects are larger for males compared to females (see Figure B4). This is consistent with the assumption that stereotypes regarding crime and driving styles apply particularly to male foreigners (Trager et al., 2014). Consumers may generally feel less safe with a male driver with an Arab/Turkish/Persian sounding name compared to a male German driver. Female drivers with an Arab/Turkish/Persian sounding name may not be disadvantaged compared to female German drivers given that the crime stereotype largely applies to males. However, sample size limitations regarding female drivers with foreign names (N=31) do not allow us to infer that females are not subject to discrimination. The effect is smaller but not zero (see Figure B4 and Table A4 in the Appendix).

Second, Arab/Turkish/Persian drivers without a profile picture are much more disadvantaged than drivers from the same ethnic group with a profile picture (see Figure B5).^{xiv} We interpret the profile picture to be a trust enhancing measure. Ethnic stereotypes regarding safety may simply have more room to engage imagination when users do not know what the driver looks like.

Third, we suspected that consumers may select rides based on sociability considerations. Drivers with an Arab/Turkish/Persian name could be discriminated against when consumers assume that ‘foreign’ drivers may not speak the language or do not share similar music tastes which could make the joint ride less enjoyable. In fact, our findings show that Arab/Turkish/Persian drivers are less disadvantaged when they have indicated a preference for talking during the ride (see Figure B6). We speculate that consumers interpret a talking preference for foreigners as a sign of good German language skills. In contrast, a preference against talking may simply be interpreted as a potential language barrier. This could explain why the positive effect of talking preference on clicks is considerably larger for Arab/Turkish/Persian drivers compared to German drivers.

The results for music preference could be interpreted in a similar vein (see Figure B7). Our findings show that Arab/Turkish/Persian drivers are more disadvantaged when they indicate a preference for music during the ride. Again, consumers may infer that a music preference implies a lack of willingness to talk which, in turn, could be perceived as a language barrier. In addition, consumers may assume that drivers with an Arab/Turkish/Persian name may want to listen to ethnic music which could reduce the enjoyment of the ride for German consumers who may be less likely to share similar tastes.

In sum, our results document substantial ethnic discrimination in Germany's carpooling market. Our findings highlight the power of information about drivers. The more useful information is provided about drivers with foreign-sounding names, the less likely they are to be discriminated. We provided tentative evidence that safety and sociability considerations may drive this information effect. Our results are robust against a series of checks including different samples, variable operationalization, estimators and potential social class bias (see the appendix for details).

One important potential source of bias is perceived social class given that Arab and Turkish migrant communities are overrepresented among lower social classes in Germany. We compare the effects of typically Arab/Turkish/Persian names with the effect of typically Anglo-Saxon names (i.e. Steven, Justin, Kevin). Studies in the German context have shown that Anglo-Saxon names in Germany are associated with low social class (e.g. Kaiser, 2010). The effect of such names, however, is not statistically significant and not negative. As such, we provide tentative findings that the ethnic penalty appears to be robust against social class bias.

6. Summary & Discussion

Recent ethnic discrimination studies increasingly make use of online market data to better understand when and why ethnic discrimination occurs. We aim to contribute to this effort with a novel application of ethnic discrimination in Europe's largest online carpooling market (i.e. Germany). We argue that there are four aspects that make our study a valuable contribution to existing research. *First*, carpooling is a social market that involves one-off, face-to-face interaction in a non-professional setting. This unique setting broadens the scope of ethnic discrimination research to more subtle, diverse and everyday interactions where ethnic minorities may face unequal treatment. *Second*, the social element of this market and the fact that we are able to measure all relevant observable characteristics allows us to test assumptions regarding the mechanisms driving ethnic discrimination. We focus in particular on the effects of information. *Third*, the advantage of our application is that we observe real actors making real decisions in real markets while being able to control all relevant factors that may influence consumer choice. Using observational data while holding all relevant confounders constant is a key advantage of our study and responds to the critique of experimental methods that might create rare and artificial situations (e.g. Heckman 1998). *Fourth*, we can provide tentative analysis to disentangle social from ethnic cues which is a limitation of many ethnic discrimination studies.

We find evidence of substantial ethnic discrimination in Germany's carpooling market. Drivers with Arab/Turkish/Persian sounding names obtain – *ceteris paribus* – less demand (on average 13% less) compared to German drivers. The average Arab/Turkish/Persian driver in our analysis would have to offer his ride 4.20 € cheaper than the average German driver to achieve the same demand, a discriminatory price premium that is equivalent to 32% of the price for an average ride. This finding is robust against a broad range of robustness checks.

Discrimination of drivers with Arab/Turkish/Persian names persists despite the relatively young and urban consumer composition in this particular market. Therefore, the estimated discrimination effect may be conservative compared to other everyday social interactions with ethnic minorities in the German society. Our study is consistent with other recent studies that show ethnic/racial discrimination effects in other online consumer markets (Ayres et al., 2011; Blommaert et al., 2014; Doleac & Stein, 2013; Edelman et al., 2017; Przepiorka, 2011; Robnett & Feliciano, 2011; Zussman, 2013). Discrimination against individuals with Arab, Persian and Turkish sounding names is consistent with findings in other studies across Europe (Blommaert et al., 2014; Gaddis & Ghoshal, 2015; Rich, 2014).

One main result of our study is that ethnic disparities decrease or increase depending on the information that is provided about the service provider (the driver). Higher user ratings, a higher number of ratings and driver experience decrease ethnic discrimination. This shows that stereotypes regarding particular ethnic groups become more salient and active when other information that could signal trust is scarce. In other words, discrimination is more pervasive in information-scarce environments. Consumers appear to use the name origin as a signal for other relevant information that is otherwise not available. Previous studies have argued that such information effects are consistent with statistical discrimination. We are cautious to make strong judgments on the relative importance of statistical and taste-based ethnic discrimination in general, as we do not have the tools to adequately isolate taste-based discrimination. Regardless of this relative importance, the strong effects of information, including user ratings, deserves attention in its own merit and support some previous evidence that suggests that information can ameliorate discrimination effects (Abraham et al., 2017; Kaas & Manger, 2012; Nunley et al., 2011; Ahmed et al., 2010).

Common statistical discrimination assumes that discriminatory behavior is based on stereotypes which are commonly difficult to capture empirically. Our data allowed us to provide a number

of indirect tests of underlying stereotypes that might drive discrimination. Unlike conventional studies in the area of employment, for example, stereotypes in carpooling do not (only) revolve around low productivity, low educational achievement or work ethics. Our analysis suggests that safety and sociability considerations apply. The results show that customers may have – *ceteris paribus* – less trust in a driver with a foreign name and that foreign drivers may signal a lower social value, for example, because they may speak a different language and prefer different music during the ride. Tentative evidence suggests that these effects are not driven by social class bias.

Our findings have implications for policy. Our findings underscore the importance of a general discussion about anti-discrimination legislation in the internet age. It is possible that due to one-to-one communication in online markets, discrimination goes largely undetected and unsanctioned. Moreover, our findings suggest that the type and level of information provided matters for the degree of discrimination, providing a useful leverage point for policy makers. The magnitude of ethnic discrimination *decreases* with an *increase* of available context information about individual actors. Growing evidence in support of statistical discrimination is good news for policy makers as information is often more malleable to policy than deep-rooted prejudice. Another strategy is to remove the information or signal that induces unequal treatment in some users, i.e. the name. Without the name it is harder to assign (ethnic) group membership and thus, harder for stereotypes to be activated. The online apartment sharing platform Airbnb has recently followed this strategy in response to growing evidence of discrimination on their platform (see Edelman et al., 2017)^{xv} The startup adopted new non-discrimination policies and systems to address user complaints. Airbnb is now promising to allow guests to book without prior approval or screening by the host and to reduce the prominence of pictures on guests' profiles in favor of more 'objective', reputation-enhancing information.

Our study also faces certain limitations. Unfortunately, information about the consumers of the rides is not available. As such, we cannot directly correct entirely for ethnic selection on the demand side. Arab/Turkish/Persian riders may be more likely to prefer Arab/Turkish/Persian drivers in which case our estimated discrimination effect is conservative. If there were a noticeable negative in-group bias (unlikely theoretically and given the small relative size of Arab/Turkish/Persian consumers), the effects can still be interpreted as evidence of discrimination in the framework of statistical discrimination). Due to data limitations, we were also not able to match more disaggregated regional population statistics. This information would allow us to disaggregate effects by location and look at how population attitudes may correlate with discriminatory behavior. Further comparative research is also needed to assess discrimination varies across national contexts and different ethnic groups.

In summary, our results have illustrated the power of a name and the information associated with it. We find that the name is used as a proxy for the trustworthiness of actors in a social market environment. Foreign-sounding drivers are trusted less than German drivers when information is scarce. When more information is provided about both drivers, discrimination decreases to levels that are statistically undistinguishable from zero. In other words, when little is known about the quality of a ride, drivers with typical German names enjoy a certain ‘blind trust premium’ that cannot be explained by any relevant quality indicator. Unique to our social market scenario, we were able to provide indication that discrimination in social markets is based on assumptions regarding safety and the social value of spending time with a member of another ethnic group in terms of language barriers and tastes (e.g. music preference). Our findings highlight the role of ethnic discrimination in subtle, everyday social interactions between ethnic groups and the powerful role of information to influence discrimination.

References

































- Abrahao, B., Parigi, P., Gupta, A., & Cook, K. S. (2017). Reputation offsets trust judgments based on social biases among Airbnb users. *Proceedings of the National Academy of Sciences*, **114** (37), 9848–9853
- Ahmed, A. M., Andersson, L., & Hammarstedt, M. (2010). Can discrimination in the housing market be reduced by increasing the information about the applicants?. *Land Economics*, **86**(1), 79-90.
- Al Ramiah, Ananthi, and Miles Hewstone (2013). Intergroup Contact as a Tool for Reducing, Resolving, and Preventing Intergroup Conflict: Evidence, Limitations, and Potential, *American Psychologist*, **68**(7), 527.
- Agrawal, Ajay, Lacetera, Nicola, and Elizabeth Lyons. (2013). Does Information Help or Hinder Job Applicants from Less Developed Countries in Online Markets? *National Bureau of Economic Research* No. w18720.
- Aigner, Dennis, and Glen Cain. (1977). Statistical Theories of Discrimination in Labor Markets. *Industrial and Labor Relations Review*, **30**(2), 175-187.
- Altonji, Joseph, and Charles Pierret. (2001). Employer Commitment and Statistical Discrimination. *The Quarterly Journal of Economics* **116**, 313-350
- Altonji, Joseph, and Rebecca Blank. (1999). Race and Gender in the Labor Market. *Handbook of Labor Economics* **3**, 3143-3259.
- Anderson, Lisa, Fryer, Roland and Charles Holt. (2006). Discrimination: Experimental Evidence from Psychology and Economics. Pp. 97-118 in *Handbook on the Economics of Discrimination* edited by William Rodgers. Edward Elgar Publishing.
- Arrow, Kenneth. (1973). The Theory of Discrimination. Pp. 3–33 in *Discrimination in Labor Markets*, edited by Orley Ashenfelter and Albert Rees. Princeton University Press.
- Ayres, Ian, Banaji, Mahzarin and Christine Jolls. (2015). Race Effects on eBay. *The RAND Journal of Economics* **46** (4): 891-917.
- Becker, Gary. (1971). *The Economics of Discrimination*. Chicago: University of Chicago Press.
- Bertrand, Marianne, and 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.
- Blank, Rebecca, Dabady, Marilyn, and Constance Citro. (2004). Measuring Racial Discrimination: National Research Council Panel on Methods for Assessing Discrimination. National Academy of Sciences.
- Blommaert, Lieselotte, Coenders, Marcel, and Frank van Tubergen. (2014). Discrimination of Arabic-named Applicants in the Netherlands: An Internet-based Field Experiment Examining Different Phases in Online Recruitment Procedures, *Social forces* **92** (3), 957-982.
- Booth, Allison, Leigh, Andrew, and Elena Varganova. (2012). Does Ethnic Discrimination Vary Across Minority Groups? Evidence from a Field Experiment. *Oxford Bulletin of Economics and Statistics* **74**(4), 547-573.
- Bryson, Alex, and Arnaud Chevalier. (2015). Is There a Taste for Racial Discrimination Amongst Employers? *Labour Economics*, **34**, 51-63.

- Diehl, Claudia, Andorfer, Veronika, Khoudja, Yassine, and Karolin Krause. (2013). Not in My Kitchen? Ethnic Discrimination and Discrimination Intentions in Shared Housing among University Students in Germany, *Journal of Ethnic and Migration Studies*, **39**(10), 1679-1697.
- Doleac, Jennifer, and Luke Stein. (2013). The Visible Hand: Race and Online Market Outcomes, *The Economic Journal*, **123**(572), 469-492.
- Edelman, Benjamin, Luca, Michael, and Dan Svirsky. (2017). Racial Discrimination in the Sharing Economy: Evidence from a Field Experiment, *American Economic Journal: Applied Economics*, **9**(2), 1-22.
- Ewens, Michael, Tomlin, Bryan and Liang Choon Wang. (2014). Statistical Discrimination or Prejudice? A Large Sample Field Experiment, *Review of Economics and Statistics*, **96**(1), 119-134.
- Fitzgerald, Jennifer, Curtis, Amber, and Catherine L. Corliss. (2011). Anxious Publics - Worries About Crime and Immigration, *Comparative Political Studies*, **45** (4), 477 – 506.
- Gaddis, Michael, and Raj Ghoshal. (2015). Arab American Housing Discrimination, Ethnic Competition, and the Contact Hypothesis. *The ANNALS of the American Academy of Political and Social Science*, **660**(1), 282-299.
- Gneezy, Uri, List, John, and Michael Price. (2012). Toward an Understanding of Why People Discriminate: Evidence From a Series of Natural Field Experiments, *National Bureau of Economic Research* No. w17855
- Guryan, Jonathan, and Kerwin Charles. (2013). Taste-based or Statistical Discrimination: The Economics of Discrimination Returns to its Roots, *The Economic Journal*, **123**(572), 417-432.
- Heckman, James. (1998). Detecting Discrimination, *The Journal of Economic Perspectives*, **12**(2): 101-116.
- Lin, Ken-Hou, and Jennifer Lundquist. (2013). Mate Selection in Cyberspace: The Intersection of Race, Gender, and Education, *American Journal of Sociology*, **119**(1), 183-215.
- Jakobsson, Niklas, and Henrik Lindholm. (2014). Ethnic Preferences in Internet Dating: A Field Experiment. *Marriage & Family Review* **50**(4), 307-317.
- Kaas, Leo, & Manger, Christian. (2012). Ethnic discrimination in Germany's labour market: a field experiment. *German Economic Review*, **13**(1), 1-20.
- Kaiser, Astrid. (2010). Vornamen: Nomen est omen? Vorerwartungen und Vorurteile in der Grundschule. *Schulverwaltung. Zeitschrift für Schulleitung und Schulaufsicht*, **21**(2), 58-59.
- Kalter, Frank. (2006). In Search of an Explanation for the Specific Labor Market Disadvantages of Second Generation Turkish Migrant Children. *Zeitschrift für Soziologie*, **35**(2), 144-160.
- List, John. 2004. The Nature and Extent of Discrimination in the Marketplace: Evidence from the Field, *The Quarterly Journal of Economics*, **119**(1), 49-89.
- McPherson, Miller, Smith-Lovin, Lynn, and James M. Cook. (2001). Birds of a Feather: Homophily in Social Networks, *Annual Review of Sociology*, **27**(1), 415-444.
- Neumark, David. (2012). Detecting Discrimination in Audit and Correspondence Studies, *Journal of Human Resources* **47**(4), 1128-1157.

- Nunley, John, Owens, Mark, and Stephen Howard. (2011). The Effects of Information and Competition on Racial Discrimination: Evidence from a Field Experiment, *Journal of Economic Behavior & Organization*, **80**(3), 670-679.
- Pager, Devah, and Hana Shepherd. (2008), The *Sociology of Discrimination: Racial Discrimination in Employment, Housing, Credit, and Consumer Markets*, *Annual Review of Sociology*, **34**, 181-209.
- Pager, Devah., Bonikowski, Bart, and Bruce Western (2009). Discrimination in a Low-wage Labor Market: A Field Experiment, *American Sociological Review*, **74**(5), 777-799.
- Phelps, Edmund (1972). The Statistical Theory of Racism and Sexism, *The American Economic Review*, **62**(4), 659-661.
- Przepiorka, Wojtek. (2011). Ethnic Discrimination and Signals of Trustworthiness in an Online Market: Evidence from two Field Experiments, *Zeitschrift für Soziologie*, **40**(2), 132-141.
- Rich, Judith. (2014). What Do Field Experiments of Discrimination in Markets Tell Us? A Meta Analysis of Studies Conducted Since 2000, *IZA Discussion Paper*, **8584**(1), 1-68.
- Robnett, Belinda, and Cynthia Feliciano. (2011). Patterns of Racial-Ethnic Exclusion by Internet Daters, *Social Forces* **89**(3), 807-828.
- Trager, Glenn and Charis E. Kubrin. (2014). Complicating the Immigration-Crime Nexus: Theorizing the Role of Gender in the Relationship between Immigration and Crime. Pp. 527-548 in Rosemary Gartner and Bill McCarthy (Eds.), *The Oxford Handbook of Gender, Sex, and Crime*. New York: Oxford University Press.
- Zussman, Asaf. (2013). Ethnic Discrimination: Lessons from the Israeli Online Market for Used Cars, *The Economic Journal*, **123**(572), 433-468.

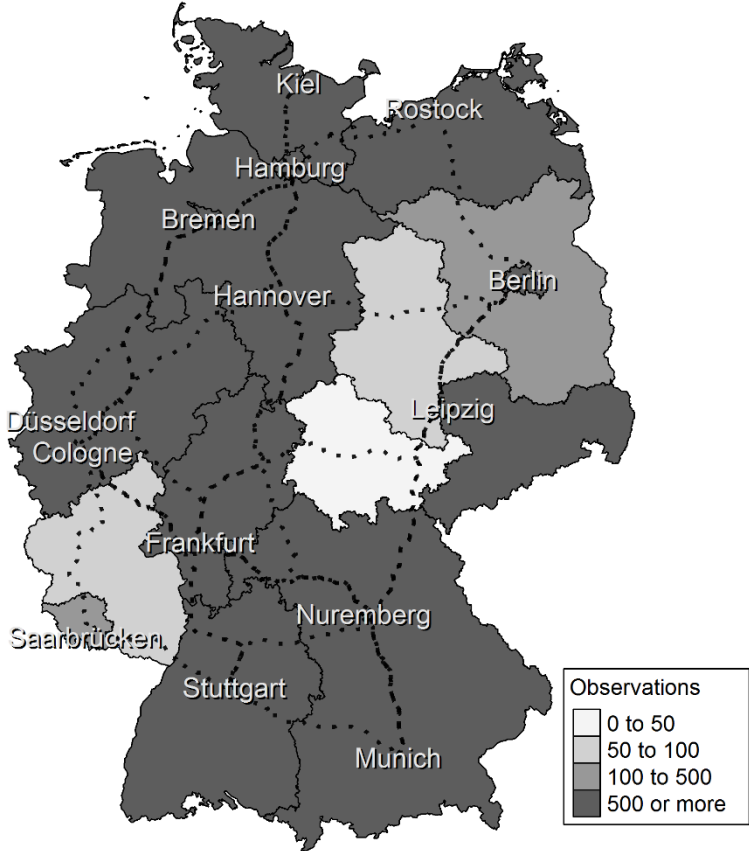
TABLES & FIGURES (in order of appearance)

Figure 1: User Interface of Online Carpooling Platform

Entfernung: 233 km - Dauer: 2Std.18Min.		Ordnen nach	🕒	☰	€
 <p>Female M 20 Jahre Aufsteiger/in</p> <p>★ 4.0 - 2 Bewertungen</p> <p> </p>	<p>Morgen - 21:00 Uhr Rostock → Berlin</p> <p> Rostock  Berlin</p> <p>Fahrzeug: OPEL CORSA ★★</p>	<p>11 € pro Mitfahrer/in</p> <p>3 Plätze frei</p>			
 <p>Female M 20 Jahre</p> <p>   </p>	<p>Mittwoch 11. November - 13:30 Uhr Rostock → Berlin</p> <p> Rostock, Deutschland  Mollstraße 19, 10249 Berlin, Deutschland</p> <p>Fahrzeug: SKODA OCTAVIA Combi ★★</p>	<p>11 € pro Mitfahrer/in</p> <p>2 Plätze frei</p>			
 <p>Female M 20 Jahre Fortgeschrittene/r</p> <p>★ 4.7 - 3 Bewertungen</p> <p></p>	<p>Mittwoch 11. November - 14:00 Uhr Rostock → Berlin</p> <p> Rostock Hbf, Rostock  S+U Hermannstr. (Berlin), Neukölln</p> <p>Fahrzeug: VOLKSWAGEN PASSAT ★★★★★</p>	<p>12 € pro Mitfahrer/in</p> <p>4 Plätze frei</p>			
 <p>Male M 20 Jahre Aufsteiger/in</p> <p>★ 5.0 - 1 Bewertung</p> <p> 161 Freunde/innen</p> <p>   </p>	<p>Mittwoch 11. November - 15:10 Uhr Rostock → Berlin → Ansbach</p> <p> Rostock Hbf, Rostock  Ankunft: Berlin (Bitte sprechen Sie die Details mit dem Fahrer/der Fahrerin ab.)</p> <p>Fahrzeug: SKODA OCTAVIA Combi ★★</p>	<p>11 € pro Mitfahrer/in</p> <p>3 Plätze frei</p>			
 <p>Female M 20 Jahre Aufsteiger/in</p> <p>★ 5.0 - 2 Bewertungen</p> <p> 237 Freunde/innen</p> <p>   </p>	<p>Mittwoch 11. November - 17:00 Uhr Rostock → Berlin → Bernau Bei Berlin</p> <p> Rostock, Deutschland  Ankunft: Berlin, Deutschland (Bitte sprechen Sie die Details mit dem Fahrer/der Fahrerin ab.)</p> <p>Fahrzeug: VOLKSWAGEN GOLF V ★★</p>	<p>11 € pro Mitfahrer/in</p> <p>3 Plätze frei</p>			

Source: Carpooling Data Germany 2015 (compiled by authors). Images, names and age of drivers pixelated. See main text for description in English.

Figure 2: Sampled rides in online carpooling market



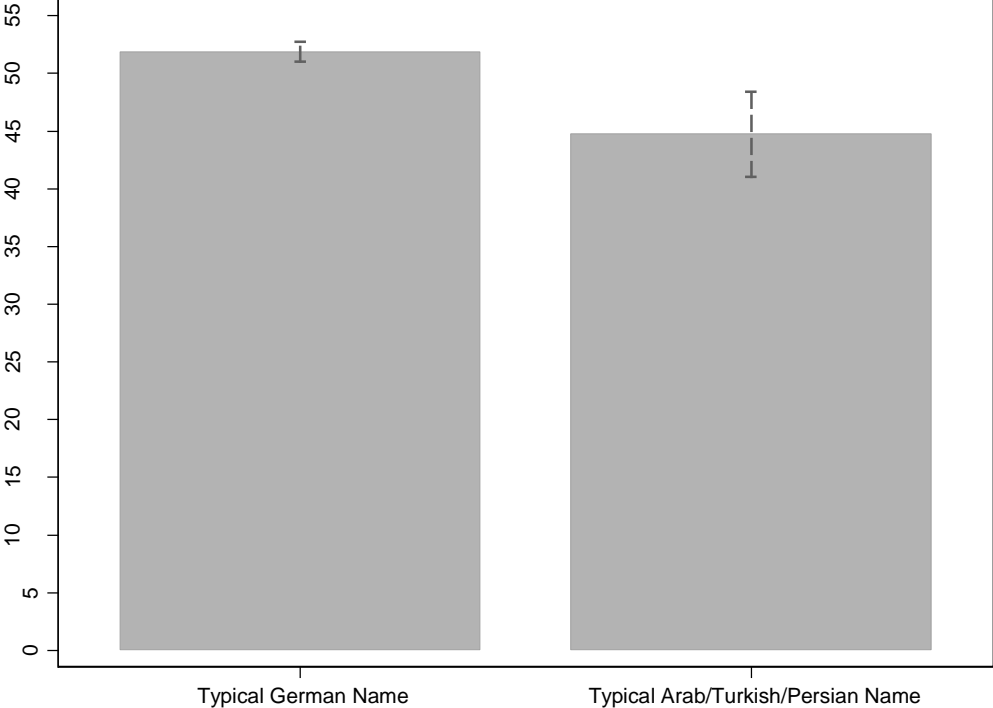
Source: Carpooling Data Germany 2015 (compiled by authors)

Table 1: Most frequent name origins with high origin certainty

#	Male names		Female names	
	Arab/ Persian/ Turkish	German	Arab/ Persian/ Turkish	German
1	Ali	Thomas	Sanam	Julia
2	Mohammed	Christian	Halime	Sarah
3	Süleyman	Daniel	Sahar	Johanna
4	Seref	Martin	Hülya	Lisa
5	Mohamed	Michael	Taman	Anna
6	Kadir	Alexander	Büsra	Katharina
7	Serdar	Andreas	Dersimgül	Grit
8	Ismail	Sebastian	Gülli	Maria
9	Mustafa	Markus	Gülten	Laura
10	Cem	Jens	Husna	Anne
11	Osman	Peter	Nasrin	Franziska
12	Salman	Tobias	Nesrin	Lena
13	Yusuf	Christoph	Senem	Stefanie
14	Amir	Matthias	Sinem	Alexandra
15	Ercan	Stefan	Özlem	Anja
16	Mehdi	Chris	Abdul	Annika
17	Oguz	Robert	Cigdem	Nadine
18	Rami	Jan	Elif	Sandra
19	Ahmad	Volker	Fatemeh	Miriam
20	Ersin	Friedrich	Gülcan	Carolin

Note: Most frequent driver names by name origin based on online survey ratings (N=1,577 survey participants; 20 ratings per name on average).


Figure 3: Predicted Number of Clicks on Offer Ride by Name Origin of the Driver



Note: Carpooling Data Germany 2015 (compiled by authors). Predicted clicks are adjusted (see full model in Table A3 in the Appendix). N= 16,624. Group differences are statistically different ($p < 0.01$).


Figure 4: Predicted Number of Clicks on Offer by Certainty of Arab/Turkish/Persian Name

Origin



Note: Carpooling Data Germany 2015 (compiled by authors). Predicted clicks are adjusted (see Table A3 in the Appendix). N= 16,624. Effect of continuous name measure is statistically significant ($p < 0.01$). 95% confidence interval.

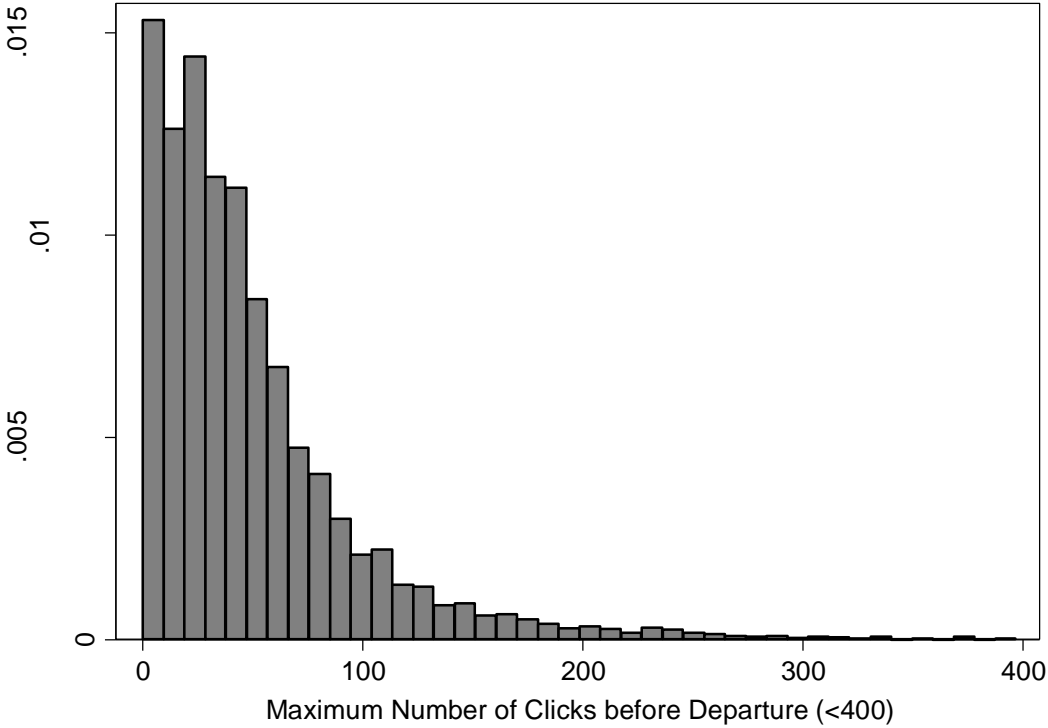
Figure 5: Predicted Number of Clicks on Offered Ride by Name Origin of the Driver and the User Rating



Note: Carpooling Data Germany 2015 (compiled by authors). Predicted clicks are adjusted for all observable characteristics (see Table A3 in the Appendix). N= 16,624. 95% confidence interval.

APPENDIX (in the order of appearance)

Figure B1: Distribution of Clicks



Source: Carpooling Data Germany 2015 (compiled by authors) (N=16,624, clicks above 400 clicks excluded for visualization purposes)

Table A1: Operationalization of model variables

Variable	Description	Operationalization (see also Table A3)
Dependent Variables		
Clicks	The maximum number of clicks an offered ride received before departure	Continuous
Independent Variable		
Name origin	Rated origin of the driver's name	Categorical: Arab/Turkish/ Persian vs. German (Continuous scale in Figure 4)
Controls		
Time of day	The time during a day when the ride departs	Categorical: night, morning, midday, afternoon, evening
Day of week	The day of the week when the ride departs	Categorical: Monday - Sunday
Distance in km	The distance between departure city and arrival city	Continuous
Price in euro	Price to be paid for one seat on the ride	Continuous
Gender	Gender of the driver	Categorical: Female, male
Age	Age of the driver	Continuous (in years)
Smoking preference	Smoking preference of the driver	Categorical: yes, no
Music preference	Music preference of the driver	Categorical: yes, no
Dialog preference	Dialog preference of the driver	Categorical: yes, no, maybe
Rating	Rating of the driver by previous customers	Categorical: 0-5 stars
Experience	Experience of the driver based on the number of offered rides in the past	Categorical: 0 – no experience to 4 – high experience
Picture	Availability of a profile picture for the driver	Categorical: yes, no
Comfort	Comfort of the ride conditional of the type of car	Categorical: simple/normal, comfortable, luxury, score not available
Auxiliary Variables		
Route ID	Control for routes between cities	Dummy variable
Ride Volume	Control for number of rides offered per route	Continuous variable
Log days until departure	Log number of days before departure (i.e. the time the offer was online before departure)	Continuous variable

Table A2: Summary Statistics of Clicks Model (see Figure 3 and Table A3)

	German				Arab/Turkish/Persian				T-test		
	<i>mean</i>	<i>sd</i>	<i>min</i>	<i>max</i>	<i>mean</i>	<i>sd</i>	<i>min</i>	<i>max</i>	<i>b</i>	<i>T</i>	
Maximum Number of Clicks on Offer	51.4	53.7	0	629	45.8	49.4	0	331	5.53*	(2.50)	
Traffic volume	16.1	11.9	2	58	12.4	9.8	2	58	3.67***	(8.29)	
Days online until departure	3.6	2.7	1	11	4.3	3.2	1	11	-0.78***	(-5.51)	
Female	0.3	0.4	0	1	0.1	0.2	0	1	0.21***	(18.76)	
Age*	31.3	9.7	18	101	31.3	7.0	18	54	-0.07	(-0.24)	
Number of Ratings	5.3	11.8	0	154	4.3	10.1	0	79	1.03*	(2.28)	
User Rating	2.9	2.3	0	5	2.6	2.3	0	5	0.36***	(3.58)	
Experience	1.1	1.2	0	4	0.9	1.1	0	4	0.20***	(3.89)	
Profile picture	0.4	0.5	0	1	0.4	0.5	0	1	0.01	(0.27)	
Smoking preference	0.0	0.2	0	1	0.1	0.3	0	1	0.06***	(4.68)	
Music preference	0.4	0.5	0	1	0.4	0.5	0	1	-0.06***	(-4.68)	
Talking preference									-0.01	(-0.56)	
	maybe	0.9	0.3	0	1	0.8	0.4	0	1	0.01	(0.56)
	yes	0.1	0.3	0	1	0.2	0.4	0	1	0.10***	(5.56)
	no	0.0	0.2	0	1	0.1	0.2	0	1	-0.08***	(-4.95)
Car comfort										-0.02*	(-2.02)
	simple/normal	0.5	0.5	0	1	0.3	0.5	0	1	0.18***	(8.58)
	comfortable	0.2	0.4	0	1	0.2	0.4	0	1	0.01	(0.42)
	luxury	0.0	0.2	0	1	0.1	0.3	0	1	-0.04***	(-3.54)
	n/a	0.2	0.4	0	1	0.4	0.5	0	1	-0.14***	(-6.59)
Nighttime	0.0	0.2	0	1	0.1	0.3	0	1	-0.07***	(-4.99)	
Morning	0.2	0.4	0	1	0.1	0.3	0	1	0.06***	(3.83)	
Midday	0.2	0.4	0	1	0.2	0.4	0	1	0.04*	(2.37)	
Afternoon	0.4	0.5	0	1	0.4	0.5	0	1	0.00	(0.08)	
Evening	0.1	0.4	0	1	0.2	0.4	0	1	-0.03	(-1.88)	
Sunday	0.3	0.4	0	1	0.2	0.4	0	1	0.03	(1.35)	
Monday	0.1	0.3	0	1	0.1	0.3	0	1	-0.01	(-0.66)	
Tuesday	0.1	0.2	0	1	0.1	0.2	0	1	-0.01	(-1.19)	
Wednesday	0.1	0.2	0	1	0.0	0.2	0	1	0.01	(1.22)	
Thursday	0.1	0.3	0	1	0.2	0.4	0	1	-0.04*	(-2.36)	
Friday	0.3	0.4	0	1	0.2	0.4	0	1	0.03	(1.73)	
Saturday	0.1	0.3	0	1	0.1	0.3	0	1	-0.01	(-0.60)	
Distance in km	260.9	154.2	36	1178	290.0	168.8	39	653	-29.11***	(-3.87)	
Price in Euro	13.3	7.9	1	58	14.9	8.8	1	40	-1.61***	(-4.13)	
N	16107				517						

* We decided to keep age outliers in the estimation model. Drivers that reported an age over 80 years old represent 0.004% of the sample. Excluding outliers from the estimation does not change the results

Table A3: The effect of name origin on clicks – all model coefficients (Average Marginal Effects)

	Empty Model	Full Model	Full model including interaction between rating and name origin (see Fig 5)
	AME	AME	Coef.
Arab/Turkish/Persian name	-7.725*** (1.87)	-7.178*** (1.92)	-0.334*** (0.06)
Number of rides per route	0.0613 (0.04)	-0.0973 (0.05)	-0.00194* (0.00)
Log days until departure	-3.838*** (0.53)	-3.668*** (0.54)	-0.0714*** (0.01)
Time of day (ref. night)			
morning		-3.044 (2.13)	-0.0494 (0.04)
mid-day		-2.408 (2.10)	-0.0384 (0.04)
afternoon		-2.889 (2.17)	-0.0463 (0.04)
evening		-9.319*** (2.10)	-0.178*** (0.04)
Day of week (ref. Sunday)			
Monday		-2.752 (1.62)	-0.0533 (0.03)
Tuesday		-11.02*** (1.78)	-0.219*** (0.04)
Wednesday		-7.369*** (1.85)	-0.143*** (0.04)
Thursday		-10.17*** (1.33)	-0.204*** (0.03)
Friday		-2.425* (1.10)	-0.0452* (0.02)
Saturday		-6.907*** (1.45)	-0.132*** (0.03)
Distance in km		0.277*** (0.02)	0.00537*** (0.00)
Price in euro		-1.695*** (0.19)	-0.0330*** (0.00)
Female		4.365*** (0.93)	0.0825*** (0.02)
Age		-0.00684 (0.04)	-0.000107 (0.00)
Number of user ratings		0.118** (0.04)	0.00227** (0.00)
Rating		-0.0301 (0.31)	-0.00280 (0.01)
Experience (ref. Newcomer)			
Intermediate		-0.891 (1.53)	-0.0169 (0.03)
Experienced		1.026 (1.80)	0.0200 (0.03)
Expert		0.970 (1.88)	0.0213 (0.04)
Ambassador		6.204* (2.85)	0.119* (0.05)
Profile picture available		1.113	0.0195

Smoking preference (ref. No/ Maybe)		(0.85)	(0.02)
Yes		2.830 (1.96)	0.0488 (0.04)
Music (Maybe/No)			
Yes		1.025 (0.86)	0.0206 (0.02)
Dialog (ref. Maybe)			
Yes		5.451*** (1.48)	0.0984*** (0.03)
No		1.032 (1.96)	0.0205 (0.04)
Car comfort (ref. simple/normal)			
Comfortable		-1.374 (0.97)	-0.0269 (0.02)
Luxury		1.999 (2.08)	0.0366 (0.04)
n/a		-2.463* (1.00)	-0.0475* (0.02)
Arab/Turkish/Persian name x user rating			0.0678*** (0.02)
Observations	16624	16624	16624
<i>AIC</i>	163749.4	163297.5	163286.0
<i>BIC</i>	164521.2	164285.5	164281.7

Note: AMEs based on Negative Binomial Regression model, standard errors in parentheses. Coefficients for auxiliary variables (i.e. route id) not reported. *** p<0.01; ** p<0.05; * p<0.1. Source: Carpooling Data Germany 2015 (compiled by authors)

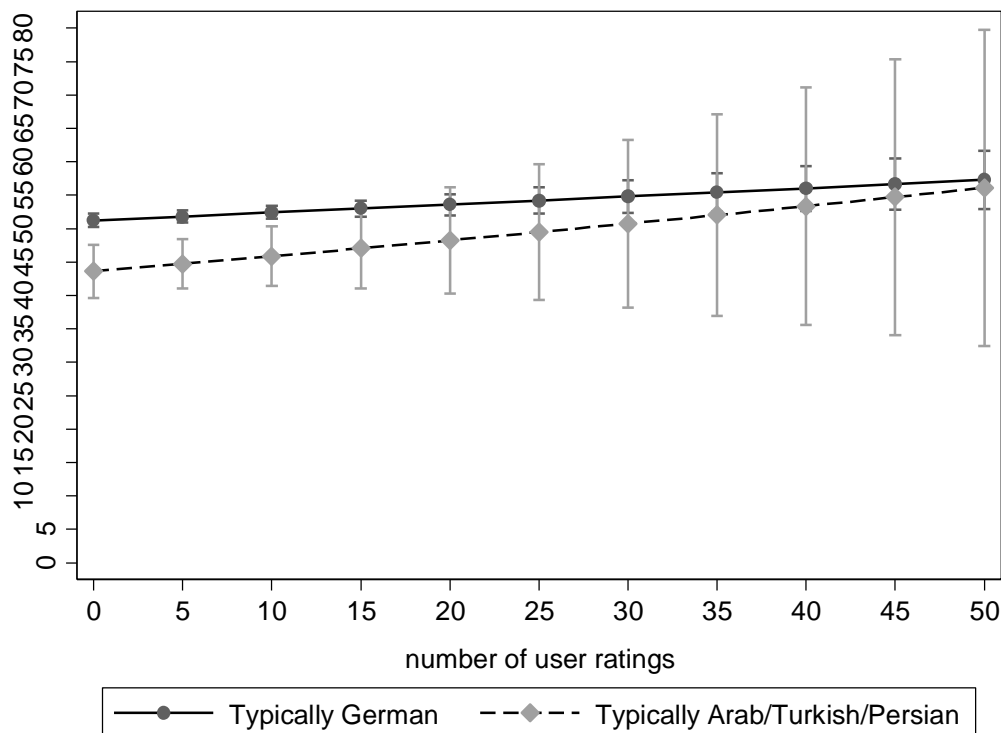
Table A4: Sub-group effects by sex, region and profile picture.

	Men only	Women only	East Germany only	West Germany only	With Profile Picture	Without Profile Picture
Arab/Turkish/Persian (ref. German names)	-7.113*** (2.02)	-4.287 (8.12)	-4.088 (4.97)	-7.820*** (2.15)	-3.826 (3.21)	-7.625** (2.51)
Observations	12299	4325	3872	12752	7396	9228
Arab/Turk./Pers. (N)	486	31	52	465	227	290
<i>AIC</i>	120583.3	42771.91	36561.04	126713.7	73114.31	90226.97
<i>BIC</i>	121517.9	43542.95	36855.33	127555.9	73970.99	91103.96

Standard errors in parentheses


* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Figure B2: Predicted Number of Clicks on Offered Ride by User Name Origin and Number of User Ratings




Note: Carpooling Data Germany 2015 (compiled by authors). Predicted clicks are adjusted (see Table A3 in the Appendix). N= 16,624. 95% confidence interval.

Figure B3: Predicted Number of Clicks on Offered Ride by User Name Origin and Driver Experience




Note: Carpooling Data Germany 2015 (compiled by authors). Predicted clicks are adjusted (see Table A3 in the Appendix). N= 16,624. 95% confidence interval.

Figure B4: Predicted Number of Clicks on Offered Ride by Name Origin and Sex



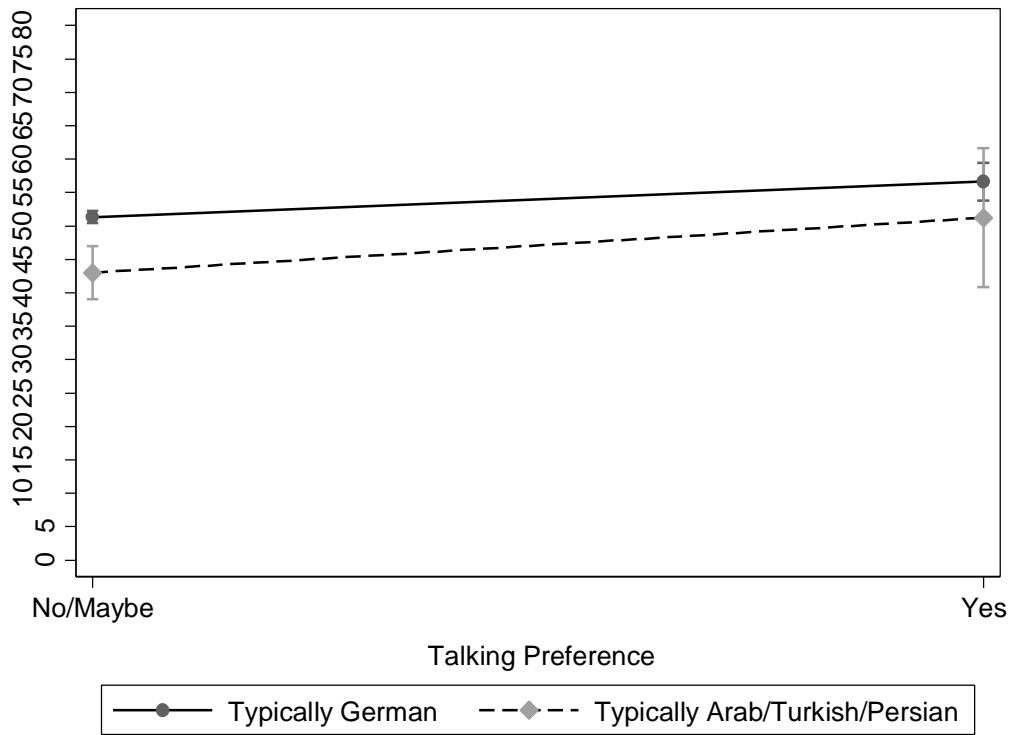
Note: Carpooling Data Germany 2015 (compiled by authors). Predicted clicks are adjusted (see Table A3 in the Annex). N= 16,624. 95% confidence interval.

Figure B5: Predicted Number of Clicks on Offered Ride by Name Origin and Profile Picture




Note: Carpooling Data Germany 2015 (compiled by authors). Predicted clicks are adjusted (see Table A3 in the Annex). N= 16,624. 95% confidence interval.

Figure B6: Predicted Number of Clicks by Name Origin and Talking Preference



Note: Carpooling Data Germany 2015 (compiled by authors). Predicted clicks are adjusted (see Table A3 in the Annex). N= 16,624. 95% confidence interval.

Figure B7: Predicted Number of Clicks by Name Origin and Music Preference



Note:

Carpooling Data Germany 2015 (compiled by authors). Predicted clicks are adjusted (see Table A3 in the Annex). N= 16,624. 95% confidence interval.

Robustness checks

As references throughout the paper, we conducted a series of robustness checks to assess the sensitivity of our results:

First, to assess the robustness of our price and distance information, we varied how price and distance enter the model. Including both separately or including a ‘price per km’ measure does not change the results.

Second, we applied propensity score matching to assess the robustness of the key ethnic name effects. Nearest neighbor matching yields slightly larger ethnic penalties and confirms our regression results.

Table A5: Propensity Score Matching

Nearest neighbor matching	Coeff.	AI Robust SE	z	p-value	95% Confidence Interval	
Average Treatment Effect	-10.15991	2.594111	-3.92	0.000	-15.2442	-5.07554
Average Treatment Effect on the Treated	-8.577434	2.407805	-3.56	0.000	-13.2966	-3.85822

Third, we investigate the effect of name origin for different sub-groups including by gender, region and availability of profile picture (see Table A4 in the Appendix).

Fourth, we attempt to disentangle ethnic from social discrimination effects by comparing the effect of drivers with an Arab/Turkish/Persian name and drivers with an ‘Anglo-Saxon name’. Several studies have shown that Anglo-Saxon names (i.e. Steven, Justin, Kevin) signal low social class in Germany (Kaiser, 2010). Table A6 presents the results for ethnic and social discrimination. We do not find any significant effects for the first names indicated by Kaiser (2010).

Table A6: Ethnic vs. Social Cues

Name origin	AME	SE	z	p-value	95% conf. Interval	
Arab/Turkish/Persian name	-7.178	1.915	-3.75	0.000	-10.932	-3.423
Anglo-Saxon name	2.753	2.968	0.93	0.354	-3.063	8.569

ENDNOTES

ⁱ The exact origin of Arab, Persian and Turkish names is difficult to distinguish for a lay person. However, members of all three groups in Germany are commonly associated to be of the same migrant group with assumed cultural similarities.

ⁱⁱ See article in the German newspaper *die Welt* entitled „BlaBlaCar und Co. vor diesen hippen Mitfahrdiensten zittert die Bahn“ (i.e. these are the carpooling services that the train companies are afraid of). Accessible at <https://www.welt.de/wirtschaft/article129721188/Vor-diesen-hippen-Mitfahrdiensten-zittert-die-Bahn.html>.

ⁱⁱⁱ As a pre-test, we uploaded a limited number of artificial rides on the route Munich to Cologne, varying profiles by the name origin only (using ‘Mehmet’ and ‘Serkan’ as typical Turkish first names and ‘Johannes’ and ‘Tobias’ as typical German first names). The pre-test indicated large discrimination effects which further strengthened our objective to collect real market data.

^{iv} See Heckman (1998) and Neumark (2012) for discussions on the limitations of audit and correspondence studies.

^v Other forms of discrimination include implicit, unintentional biases (e.g. Anderson, Fryer, & Holt, 2006). We will focus our discussion on statistical and taste-based discrimination as the decision to car-pool (our application) involves conscious weighing of numerous alternatives (other competing rides) and evaluation of several characteristics (location, price, timing, age, experience, rating et cetera).

^{vi} See news articles referring to the issue of crime rates and foreigners in Germany: <http://www.strafrecht-wi.de/auslaenderkriminalitaet/>; <http://www.bpb.de/politik/innenpolitik/innere-sicherheit/76639/auslaenderkriminalitaet?p=all>;

<http://www.spiegel.de/lebenundlernen/schule/kriminalitaet-von-migranten-laut-gutachten-nicht-hoher-a-983536.html>

^{vii} See news report covering a survey on the reputation of car drivers in different European countries: <https://www.welt.de/motor/news/article108612704/Europaweite-Umfrage.html>

^{viii} Stereotypes regarding high crime rates for ethnic minorities largely affect males (Trager et al. 2014).

^{ix} We would like to thank the provider for supporting academic research by allowing access to this data.

^x Goodness of Fit tests revealed that negative binomial regression is superior to other count models. Robustness checks reveal similar results for the OLS estimator.

^{xi} There is generally less demand for carpooling in rural areas. Ethnic minority drivers are also less often offer rides on rural routes. Including rural routes would hence bias the average number of clicks for German drivers downward.

^{xii} Raters were paid for their efforts and consisted mostly of students from laboratory pools at the University of (...) and the University of (...) (all in Germany). We thereby guaranteed that no rater participated more than once. The raters did not receive information on the aim of the study to avoid demand effects. More technical information available upon request.

^{xiii} To support this fact, we collected additional data from 38 respondents. Each respondent was asked to allocate a particular origin to the most frequent names in the Arab/Turkish/Persian group. On average, neither origin was chosen with more than 50% certainty. In comparison, the German origin was allocated to typical German names with a certainty of over 90%.

^{xiv} We also estimated the baseline model based on a subsample of rides with and without a profile picture (total N = 7,664; Arab/Turkish/Persian = 233). The effect of an Arab/Turkish/Persian name is smaller and not significant.

^{xv} See <http://blog.airbnb.com/an-update-on-the-airbnb-anti-discrimination-review>; <https://www.theguardian.com/technology/2016/jul/20/airbnb-hires-eric-holder-racial-discrimination-bias>