Can an AI Algorithm Mitigate Racial Economic Inequality? An Analysis in the Context of Airbnb

Forthcoming at *Marketing Science* Frontiers

Shunyuan Zhang

Harvard Business School, Harvard University

szhang@hbs.edu

Nitin Mehta

Rotman School of Management, University in Toronto

Nitin.mehta@Rotman.Utoronto.Ca

Param Vir Singh

Tepper School University, Carnegie Mellon University

psidhu@cmu.edu

Kannan Srinivasan

Tepper School University, Carnegie Mellon University

kannans@cmu.edu

Can an AI Algorithm Mitigate Racial Economic Inequality? An Analysis in the Context of Airbnb

ABSTRACT

We study the effect of Airbnb's smart-pricing algorithm on the racial disparity in the daily revenue earned by Airbnb hosts. Our empirical strategy exploits Airbnb's introduction of the algorithm and its voluntary adoption by hosts as a quasi-natural experiment. Among those who adopted the algorithm, the average nightly rate decreased by 5.7%, but average daily revenue increased by 8.6%. Before Airbnb introduced the algorithm, white hosts earned $12.16 more in daily revenue than Black hosts, controlling for observed characteristics of the hosts, properties, and locations. Conditional on its adoption, the revenue gap between white and Black hosts decreased by 71.3%. However, Black hosts were significantly less likely than white hosts to adopt the algorithm, so at the population level, the revenue gap increased after the introduction of the algorithm. We show that the algorithm's price recommendations are not affected by the host's race—but we argue that the algorithm's race-blindness may lead to pricing that is sub-optimal, and more so for Black hosts than for white hosts. We also show that the algorithm's effectiveness at mitigating the Airbnb revenue gap is limited by the low rate of algorithm adoption among Black hosts. We offer recommendations with which policy makers and Airbnb may advance smart-pricing algorithms in mitigating racial economic disparities.

Keywords: Pricing Algorithm, Artificial Intelligence, Mitigating Racial Disparity, Economic Gap, Sharing Economy, Airbnb, Algorithm Adoption

1. INTRODUCTION

Machine learning (ML) algorithms can use vast amounts of consumer data and automate key business decisions such as pricing, product offerings, and promotions in real-time and at micro-targeted levels. Large e-tailers such as Amazon and travel websites were early adopters of pricing algorithms, and these tools are now becoming more ubiquitous. For example, Airbnb, a sharing economy platform and the context of this study, created a smart-pricing tool based on an ML algorithm and offered it for free to all Airbnb hosts. When a host "turns on" the algorithm, it automatically adjusts the property's nightly rate to optimize revenue based on an evaluation of a rich set of factors such as the property's characteristics and seasonality that influence the demand of the property.

Pricing algorithms offer many advantages. An algorithm can use massive data on consumers and competitors to predict demand and adjust prices in real time without human intervention. Airbnb has far more data and superior computational resources than any individual host, so the pricing algorithm should be more effective than the average host at setting an optimal price based on demand dynamics (Li et al. 2016). Nevertheless, the algorithm is not guaranteed to benefit hosts, and the opacity of the algorithm makes it difficult to assess. If the incentives of the algorithm designer (Airbnb) and the adopter (host) are not perfectly aligned, then the automated prices may not be optimal for the adopter (Pavlov and Berman 2019). Further, anecdotal evidence suggests that a poorly implemented pricing algorithm can result in worse outcomes (Streitfeld 2018). Hence, our first research question is whether and to what extent Airbnb hosts have benefitted financially from adopting the algorithm.

Airbnb has received attention recently for the disparity in the revenue earned by white and Black hosts (Cox 2017). It is plausible that differences other than race (e.g., education and access to other resources) make it more difficult for Black hosts to determine optimal prices (Bayer and Kofi-Charles 2018, Chetty et al. 2020). If so, then a well-devised pricing algorithm should serve needs of Black hosts and help mitigate racial inequalities. The intuition here is that adoption of the algorithm would make it equally easy for white and Black hosts to optimize their nightly rates.

It is also possible, however, that the cause of the revenue gap cannot be overcome with the pricing algorithm. Perhaps the marketplace holds a racial bias such that guests are willing to pay more for a property owned by a white host than for an equivalent property owned by a Black host. A marketplace bias might manifest as both a lower occupancy rate and a lower nightly rate for Black hosts than for white hosts with similar properties (Edelman and Luca 2014). In other words, Black and white hosts might have different demand curves. In this case, even if the price correction by the algorithm were the same across Black and white hosts, it would have different impact on their revenues. Hence, an algorithm could increase, decrease, or maintain the revenue gap.

Finally, and counterintuitively, the algorithm might not be optimal as far as a racial disparity is concerned, by *not* considering race when setting the nightly rate. US law prohibits the explicit use of protected attributes (such as race) in the construction of algorithmic predictions or thresholds (Barocas and Selbst 2016, Fu et al. 2020). If, however, the demand curve differs between Black and white hosts, then an algorithm that ignores the host's race would effectively average the demand curves and set the same rental price for equivalent properties owned by Black and white hosts. This price would be sub-optimal for both Black and white hosts. Furthermore, as a minority at both the neighborhood and city level (Cox 2017), Black hosts are less represented in the Airbnb data, so the algorithm's prices will skew toward the white demand curve, generating prices that are even more suboptimal for Black hosts than for white hosts. Our second research question is whether Airbnb's pricing algorithm leads to similar changes in revenue among Black and white hosts—and if not, why?

The introduction of the algorithm and its adoption by Airbnb hosts provides us with a quasi-natural experiment to answer these questions. Our main results are as follows.

(1) On average, adoption of the algorithm led to a downward price correction of 5.7% across hosts, which improved their revenues by 8.6%. Before Airbnb introduced the algorithm, Black and white hosts charged similar prices for equivalent properties (in terms of observed host, property, and neighborhood characteristics), but white hosts earned $12.16 more in daily revenue than Black hosts. The revenue gap was caused by a difference in the occupancy rate (rental demand): 20% less for Black hosts' properties than for equivalent white hosts' properties. The algorithm benefited Black adopters more than white adopters, decreasing the revenue gap by 71.3%.

(2) Algorithm adoption led to a similar magnitude of downward price correction across both white and Black hosts, but it led to a greater increase in the occupancy rate for Black hosts. The disparate changes in occupancy explain why Black hosts benefited more from the algorithm

than white hosts. The result also supports our theory that Black and white hosts face different demand curves; the demand for Black hosts' properties is more responsive to price changes than the demand for equivalent properties owned by white hosts. The algorithm reportedly does not use host race to inform the optimal price, and we find that the algorithm sets similar prices for equivalent properties owned by Black and white hosts. We infer that the algorithm effectively averages the distinct demand curves for the properties of Black and white hosts when setting the price.

(3) Even though the algorithm decreased the revenue gap between white and Black adopters, two challenges remain. First, Black hosts were 41% less likely than white hosts to adopt the algorithm. As a result, at the population level (i.e., including adopters and non-adopters), the revenue gap between Black and white hosts increased after the introduction of the algorithm. Second, if Black and white hosts face different demand curves (as our data suggests), then a race-blind algorithm may set prices that are sub-optimal for both Black and white hosts, meaning that the revenue of both groups could be improved by the incorporation of race into the algorithm. Moreover, the prices are likely to be *more* sub-optimal for Black hosts since they are less represented in the data that is used to train the algorithm. We argue that Airbnb can further reduce the revenue gap between Black and white hosts by incorporating race into the algorithm, either directly or indirectly via closely correlated socioeconomic characteristics. While algorithm adoption was less likely among Black hosts than white hosts in all quartiles, it is the lowest for Black hosts in the uppermost quartile of socioeconomic status (SES). On the other hand, it is only the Black hosts in the bottom three quartiles of socioeconomic status who stand to monetarily gain by adopting the algorithm. Thus Airbnb may be able to reduce the revenue gap most efficiently by targeting algorithm promotions to Black hosts in the lower SES quartiles.

Our study has important implications for policy makers and managers. For policy makers, our study shows that when racial biases exist in the marketplace, an algorithm that ignores those biases may have limited effectiveness at reducing racial disparities. Policy makers should consider allowing algorithm designers to incorporate either race or socioeconomic characteristics that correlate with race, provided that the algorithm demonstrates an ability to reduce racial disparities. This recommendation contradicts current US policy but aligns with the emerging literature on the fairness of ML algorithms (e.g., Kleinberg et al. 2018; Williams et al. 2018).

For managers, our results suggest that the revenue gap between Black and white hosts may stem from guests' racial biases. While Airbnb cannot overturn a racial bias that is ingrained in society at large, it could try an intervention that prevents guests from knowing the host's race until they book the property (though masking the host's race could backfire, as we discuss later). Finally, managers should devise strategies to encourage algorithm adoption among Black hosts, especially those in the middle and lower SES quartiles as they reap the largest gain in daily revenue. Otherwise, a racial disparity in algorithm usage may end up increasing the economic disparity rather than alleviating it.

2. DATA

The pricing algorithm was introduced in November 2015. Our data spans July 2015 to August 2017 and pertains to 9396 randomly selected Airbnb properties located in 324 zip codes (405 neighborhoods) concentrated in seven large cities in the US. The data includes key property characteristics and each host's monthly revenue, demographics, and date of algorithm adoption (if any). Of the 9396 properties, 2118 adopted the algorithm at some point during the observation period. Hosts adopted the pricing algorithm at different points in time. Figure 1 shows the distribution of algorithm adoption events over time.

In the next subsections, we describe our dataset and briefly explain how we constructed the key variables; full details appear in Section 2 of the Web Appendix.

Revenue (dependent variable)

We calculated *DailyRevenue* as the product of the average nightly rate and occupancy rate of property i during month t. The average nightly rate of a property in a given month is taken as the average over the number of days in a month, and the occupancy rate is the ratio of the number of days the property was occupied in a month to the number of days the property was available to be booked in the same month. We obtained the average nightly price for each month and monthly occupancy rate of each property from AirDNA, a third party that specializes in collecting Airbnb data. AirDNA uses a proprietary algorithm to calculate the monthly occupancy rate (reported margin of error: 5% or less). We regret that we cannot verify this claim as we do not have access to AirDNA's algorithm. See Section 3 of the Web Appendix for a discussion of the AirDNA data.

Algorithm Adoption (treatment variable)

We define *Smart-Pricing* as a dummy variable that equals 1 if the prices for property i in month t were determined by the algorithm, and 0 otherwise. We determine algorithm usage by obtaining each property's availability calendar of all properties from their property webpages. The calendar indicates the "pricing type" of each property during each month; "*demand_based_pricing*" indicates that the algorithm was used.

Face Data of Airbnb Hosts and Race Classification

We collected the profile photo from each host's page, used a deep learning model to categorize race, and operationalized race with three binary variables: 1) *White,* 2) *Black,* and 3) *Other*; the variable *Race* equals 1 if the host's race was predicted to be the corresponding race, and 0 otherwise.

Section 1 of the Web Appendix contains a detailed explanation of the process. In short, we optimized a deep learning model on over 800,000 face images (from public datasets) with known demographic labels. The trained model categorized race (using the aforementioned three categories) with 92.5% accuracy on the reserved portion (hold-out test set) of the public dataset. Then, we extracted all human faces from the Airbnb profile photos and used the trained model to classify the race of each extracted face.

Dealing with Self-Selection

Hosts decided whether and when to adopt the algorithm, so the treatment and control groups (i.e., adopters and non-adopters) do not satisfy the parallel trends assumption in outcome. We mitigated the self-selection issue with the inverse probability of treatment weighting (IPTW) method, which contains two primary steps: estimate each property's treatment probability as a function of observed covariates and then construct the weights for each property as the inverse of its treatment assignment probability. IPTW constructs a weighted sample of treatment and control groups in which the distribution of measured baseline covariates is independent of treatment assignment (Austin and Stuart 2015; Bitler et al. 2006; Giorcelli 2019). The IPTW-weighted sample is a 'synthetic' sample of treatment and control groups in which the observed covariates are balanced across the two groups, which ensures that the treatment assignment across those two groups is independent of any observed confounder. Using this weighted sample, we then employed Difference-in-Differences (DiD) method to assess the impact of adoption of the algorithm on revenues.

Using IPTW is appealing when we have a rich set of covariates that explain the treatment assignment. To identify these covariates, we browsed the Airbnb host forums for discussions of why hosts were or were not adopting the algorithm. Section 2.4 of the Web Appendix lists the reasons and corresponding variables. In addition to these variables, we included covariates that captured characteristics of the host, property, and neighborhood, to estimate the treatment probabilities of all properties (see Sections 3.1 and 3.2 of the Web Appendix explain the IPTW analysis in detail.

We validated the parallel trends assumption (see Section 6.1 of the Web Appendix); the assumption was not rejected for the IPTW weighted sample but was rejected for the unweighted sample. This suggests that applying IPTW mitigates the self-selection issue and helps in identification to the extent that adequately reduced the systematic differences between the two groups that could influence the outcome variable.

A key limitation of IPTW is that it does not deal with unobserved confounders. We test the robustness of estimates of treatment effects to unobserved confounders by performing the conditional c-dependence sensitivity analysis proposed by Masten and Poirier (2018). Results show that our estimates of treatment effects are fairly robust to unobserved confounders. The details of this analysis are provided in section 3.3 of the Web Appendix.

We report the statistics of the key variables in Table 1. We calculated the revenue gap as the difference in the average daily revenue earned by white and Black hosts. The revenue gap was $19.8 in the pre-launch period (July 2015 to October 2015) and increased to $23.7 in the post-launch period (November 2015 to February 2016). *Prima facie*, the widening of the revenue gap suggests that the introduction of the algorithm worsened the disparity in revenues between white and Black hosts.

3. METHODOLOGY AND RESULTS

The DiD regression in Equation (1) identifies the impact of algorithm adoption on Y_{it}, the average daily revenue from property *i* in month *t*

$$Y_{it} = Property_i + \beta \cdot SmartPricing_{it} + \lambda \cdot Controls_{it} + \gamma . Seasonality_{it} + \varepsilon_{it} \quad (1)$$

The treatment variable, $SmartPricing_{it}$, equals one if the host of property *i* adopted the algorithm in month *t*, and zero otherwise. The parameter β captures the average treatment effect of algorithm adoption on the average daily revenue from property *i* in month *t*, ε_{it} is the random error, $Property_i$ is the property fixed effect, and $Controls_{it}$ represents the time-varying covariates that could impact revenue (see Section 4 of the Web Appendix for the list of covariates).

The DiD regression in Equation (2) identifies the differential effect of algorithm adoption by racial group:

$$Y_{it} = Property_i + \beta \cdot SmartPricing_{it} + \delta \cdot (SmartPricing_{it} \times Race_i) + \lambda \cdot Controls_{it} + \gamma . Seasonality_{it} + \varepsilon_{it} \quad (2)$$

Equation (2) is identical to Equation (1) except that it includes an interaction term, $SmartPricing_{it} \times Race_i$, whose coefficient captures the differential effect of algorithm adoption by race. We set *White* as the reference group for race.

We report the regression results from Equation (1) in column 1 of Table 2. The estimated coefficient of *SmartPricing* indicates that adoption of the algorithm resulted in $6.40 increase in average daily revenues, which amounts to an 8.6% increase. To understand the reason for this result, we estimated Equation (1) twice more with different dependent variables: the average nightly rate (column 2) and monthly occupancy rate (column 3). We find that algorithm adoption decreased the average nightly rate by $9.76 and increased the occupancy rate by 0.0662, suggesting that hosts had been overpricing their properties, and the algorithm adoption led to a downward price correction that attracted more frequent guests and thereby increased revenue.

We report the regression results from Equation (2) in column 1 of Table 3. The estimated coefficients of *SmartPricing* and *SmartPricing*×*Black* indicate that algorithm adoption increased the average daily revenue by \$5.22 among white hosts and \$13.92 (= \$5.22 + \$8.70) among Black hosts. To understand the disproportionate benefit for Black hosts, we estimated Equation (2) twice more with different dependent variables: the average nightly rate (column 2) and monthly occupancy rate (column 3). The main effect of *SmartPricing* on the average nightly rate is -10.15 ($p < 0.05$), but *SmartPricing*×*Black* does not have a significant interaction effect. In other words, algorithm adoption caused the average nightly rate to decrease by \$10.15, and the decrease was similar among Black and white hosts. Meanwhile, algorithm adoption increased the monthly occupancy rate of white hosts by 0.060 and Black hosts increased by 0.131 (= 0.060 + 0.071). Together, the regression results from Equation (2) suggest that algorithm adoption increased the revenue of Black hosts more than white hosts because white and Black hosts face different demand curves for equivalent properties, and the demand for Black hosts' properties is more responsive to price changes than the demand for white hosts' properties.

Next, we examined the impact of algorithm adoption on the racial revenue gap. First, we determined the revenue gap prior to algorithm adoption, which is equivalent to the main effect of *Black* in the regression in Equation (2). However, this main effect cannot be identified from equation (2) because it is absorbed by the property fixed effects, so we ran a second-stage regression. We regressed the estimated property fixed effects on race while controlling for neighborhood fixed effects and other host and property characteristics (see Section 5 of the Web Appendix). We found that *Black* had a main effect of -12.16 ($p < 0.01$), which implies that after controlling for all other observed characteristics, white hosts earned \$12.16 more in average daily revenues as compared to Black hosts, prior to adoption of the algorithm. Second, we determined the effect of algorithm adoption on the revenue gap. We took the difference between the pre-launch revenue gap (\$12.16) and the estimated coefficient of *SmartPricing*×*Black* (8.70; see column (1) of Table 3), yielding a post-launch revenue gap of \$3.46. In other words, algorithm adoption led to a 71.3% decrease in the gap in the average daily revenue of Black and white hosts.

Although the racial revenue gap decreased among hosts who adopted the algorithm, recall from Section 2 that within the full sample, the revenue gap widened from the pre-launch period ($19.8) to the post-launch period ($23.7). The contrasting results underscore that the mitigating effect (i.e., the effect on reducing revenue gap) of the algorithm is contingent on the algorithm's adoption. The algorithm *if adopted* can substantially reduce the revenue gap between Black and white hosts. However, because Black hosts were 41% less likely to adopt the algorithm than white hosts, the introduction of the algorithm widened rather than reduced the racial revenue gap.

We also sought to understand the cause of the racial revenue gap prior to algorithm adoption. We estimated the difference in the average nightly rate offered by Black and white hosts (hereafter, the "price gap") as well as the difference in the monthly occupancy rate ("occupancy gap"), controlling for all other observed characteristics. The price (occupancy) gap was the main effect of *Black* in the regression that is identical to Equation (2) but with the average nightly rate (monthly occupancy rate) as the dependent variable. Since the main effect of *Black* cannot be identified directly from these regressions, we regressed the estimated property fixed effects (calculated above) on race while controlling for neighborhood fixed effects and other host and property characteristics (see Section 5 of the Web Appendix).

In the second-stage regressions, the main effect of *Black* on the average nightly rate was non-significant, while the main effect of *Black* on the monthly occupancy rate was -0.104 ($p < 0.01$). In other words, prior to algorithm adoption, white and Black hosts charged similar prices, but the occupancy for Black hosts was 0.104 less (or 20% lower) than that for white hosts. The difference in the monthly occupancy rate accounts for the racial revenue gap prior to algorithm adoption and suggests that Airbnb guests may be systematically biased against renting from Black hosts.

Additional Analysis:

Dynamic Treatment Effects. Our main model did not allow the treatment effect to vary over time. It is possible that the treatment effect diminishes over time as more competing properties adopt the algorithm or increases over time as Airbnb improves the algorithm. We analyze this issue in Section 6.3 of the Web Appendix, and we do not find any significant dynamic treatment effect over time.

*Seasonal Effect on the Price Correction***.** In Section 7.1 of the Web Appendix, we explore whether the price correction imposed by the algorithm was the same throughout the year, or was it mainly prevalent during certain seasons. We find that algorithm adoption decreased the average nightly rate by $6.23 in all seasons, and it brings prices further down by $11.01 during the off-peak season. The result suggests that hosts tended to overprice their properties throughout the year, but more so during the off-peak season.

Alternative Mechanism for Differential Treatment Effects. We previously proposed that the differential effect of algorithm adoption between Black and white hosts occurs because the demand curve is affected by the host's race. In Section 7.2 of the Web Appendix, we explore an alternative mechanism: that race affects the seasonal variation in the price correction imposed by the algorithm. Perhaps the downward price correction imposed by the algorithm when averaged over the entire year is the same across Black and white hosts, but it is larger among Black (white) hosts during seasons when the demand for Airbnb rentals is more (less) responsive to price changes. If so, algorithm adoption would lead to a greater increase in occupancy and revenue for Black hosts than white hosts when averaged over the entire year. Our analysis, however, finds no significant difference in the downward price correction between Black and white hosts across all seasons of the year.

4. POLICY RECOMMENDATIONS FOR REDUCING THE REVENUE GAP

Although Black and white hosts seem to face different demand curves, we find that the algorithm set similar nightly rates for similar properties owned by Black and white hosts. The algorithm's apparent indifference to the host's race aligns with the methodology published by Airbnb (Ye et al., 2018): the algorithm determines a property's nightly rate by pooling demand information across all properties with similar characteristics in the same neighborhood, without considering host's race. If the demand curve indeed differs by race, such that no property characteristics can explain the disparity, then the algorithm's exclusion of the host's race should result in sub-optimal prices for both Black and white hosts (i.e., revenue would increase more for both Black and white hosts if the algorithm considered race).

Furthermore, because Black hosts are a minority at both the neighborhood and city levels, the algorithm is built on data that represents the demand curve of white hosts more than the demand curve of Black hosts. The race-blind algorithm therefore sets prices that align more closely with the optimal price of the white demand curve. In other words, although the race-blind algorithm sets prices that are suboptimal for both Black and white hosts, the prices are *more* suboptimal for Black hosts.

We suggest that Airbnb could further reduce the revenue gap by incorporating the host's race into the pricing algorithm. The consideration of a protected attribute (such as race) is not currently permitted by US policy. We join other recent studies (e.g., Kleinberg et al. 2018, Williams et al. 2018, Fu et al. 2021) in raising awareness of a counterintuitive reality: a race-blind approach to algorithmic decision-making may worsen racial disparities.

If the revenue gap between Black and white hosts stems from guests' racial biases, then Airbnb could reduce the revenue gap by preventing guests from knowing the host's race, perhaps by masking the host's profile photo until a transaction is made. A race-masking intervention could have unintended consequences, however. Agan and Starr (2018) show that a "Ban the Box" intervention, intended to protect minorities against discrimination by employers, actually *worsened* discrimination. Similarly, if Airbnb guests cannot view the host's photo, then prospective guests who are biased against Black hosts might avoid neighborhoods with a higher concentration of Black residents or make stereotyped inferences about the host's race based on their name, home décor, or other available information.

Lastly, Airbnb can reduce the revenue gap by encouraging Black hosts to adopt the algorithm. To offer a more targeted recommendation, we segmented hosts into four SES quartiles at the neighborhood-level. (Specifically, we proxied SES with educational attainment: the percentage of residents in the host's zip code with a bachelor's degree). We estimated the rate of algorithm adoption and its impact on revenue among Black hosts in each of the four quartiles (see Section 8.1 of the Web Appendix for details). Our results suggest that the rate of adoption was low among Black hosts in all four quartiles. The rate was lowest in the top quartile, but algorithm adoption did not increase the average daily revenue in that quartile, unlike in the lower three quartiles. The data do not offer a clear explanation for the disparity, but we speculate that hosts with more advanced degrees are more proficient at setting optimal prices for their properties without the algorithm's assistance. If Airbnb wishes to address the revenue gap by encouraging algorithm adoption, then it would be most efficient to target Black hosts in the lower SES quartiles.

5. CONCLUSIONS

In this paper, we study the effect of Airbnb's smart-pricing algorithm on the racial disparity in the daily revenue earned by Airbnb hosts. Our empirical strategy exploits the introduction of the algorithm and its voluntary adoption by hosts as a quasi-natural experiment. Our results show that on an average, the algorithm adoption led to 5.7% decrease in prices and 8.6% increase in revenues across hosts. Before Airbnb introduced the algorithm, white hosts earned $12.16 more in daily revenue than Black hosts, after controlling for all other observed characteristics. Conditional on its adoption, the revenue gap between white and Black hosts decreased by 71.3%. However, Black hosts were significantly less likely than white hosts to adopt the algorithm, so at the population level, the revenue gap increased after the introduction of the algorithm. We show that the algorithm is race-blind – its price recommendations are not affected by the host's race and lead to similar price corrections to white and Black hosts. We also show that the algorithm's effectiveness at mitigating the Airbnb revenue gap is limited by the low rate of algorithm adoption among Black hosts. We offer recommendations with which policy makers and Airbnb may advance smart-pricing algorithms as an equalizer of racial economic disparities.

REFERENCES

Agan, A. and S. Star (2018), “Ban the Box, Criminal Records, and Racial Discrimination: A Field Experiment,” *Quarterly Journal of Economics*, 133 (1).

Austin, P.C., E.A. Stuart. 2015. Moving towards best practice when using inverse probability of treatment weighting (IPTW) using the propensity score to estimate causal treatment effects in observational studies. *Statistics in Medicine* 34(28): 3661–3679.

Barocas, S., and A.D. Selbst (2016), “Big Data's Disparate Impact,” *California Law Review*, 104 (3), 671-732.

Bayer, P., and K.K. Charles (2018), “Divergent paths: A new perspective on earnings differences between Black and white men since 1940,” *Quarterly Journal of Economics*, *133*(3), 1459-1501.

Bitler, M.P., Gelbach, J.B., and H.W. Hoynes (2006), "What Mean Impacts Miss: Distributional Effects of Welfare Reform Experiments," *American Economic Review*, 96 (4): 988-1012.

Chetty, R., Hendren, N., Jones, M.R., and Sonya R Porter (2020), “Race and Economic Opportunity in the United States: an Intergenerational Perspective,” *Quarterly Journal of Economic*s, 135(2): 711–783.

Cox, M. (2017), “The Face of Airbnb, New York City: Airbnb as a racial gentrification tool,” http://insideairbnb.com/face-of-airbnb-nyc/

Edelman, B.G. and M. Luca (2014), “Digital Discrimination: The Case of Airbnb.com,” https://ssrn.com/abstract=2377353

Fu, R., Aseri, M., Singh, P.V., and Srinivasan, K. (2020), “’Un’fair Machine Learning Algorithms,” https://ssrn.com/abstract=3408275

Fu, R., Huang, Y., and Singh, P.V. (2021), “Crowd, Lending, Machine and Bias,” *Information Systems Research,* Articles in advnace.

Giorcelli, M. (2019), "The Long-Term Effects of Management and Technology Transfers," *American Economic Review*, 109 (1): 121-52.

Kleinberg, J., Ludwig, J., Mullainathan, S., and Rambachan, A. (2018), “Algorithmic fairness,” AEA Papers and Proceedings, 108: 22-27.

Li, J., Moreno, A., and Zhang, D. (2016), “Pros vs Joes: Agent Pricing Behavior in the Sharing Economy,” https://ssrn.com/abstract=2708279.

Masten, M. and A. Poirier (2018), “Identification of Treatment Effects Under Conditional Partial Independence,” *Econometrica* 86 (1): 317–51.

Pavlov, V., and R. Berman (2019). “Price Manipulation in Peer-to-Peer Markets and the Sharing Economy,” https://ssrn.com/abstract=3468447

Streitfeld, D. (2018), “Amazon’s Curious Case of the $2,630.52 Used Paperback,” *New York Times*.

Williams, B.A., C.F. Brooks, and Y. Shmargad. 2018. How Algorithms Discriminate Based on Data They Lack: Challenges, Solutions, and Policy Implications. *Journal of Information Policy*, 8:78-115.

Ye, P., Qian, J., Chen, J., Wu, C.H., Zhou, Y., Mars, S.D., Yang, F. and L. Zhang (2018), “Customized Regression Model for Airbnb Dynamic Pricing,” *KDD* 18, 19–23.

TABLES AND FIGURES

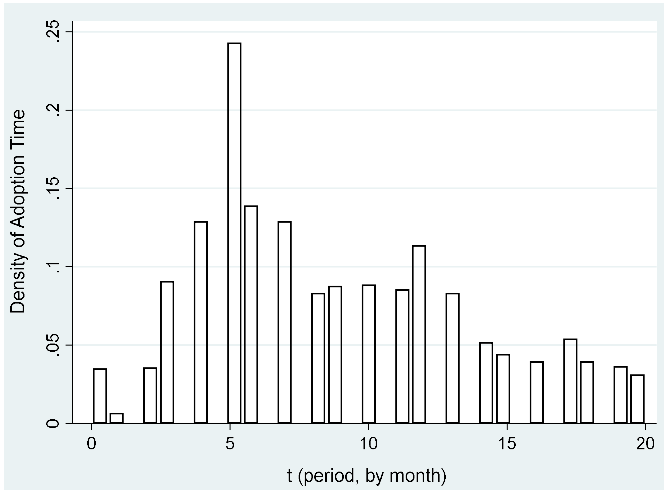

Figure 1 Histogram of the timing of algorithm adoption (t = 0: November 2015). y-axis: density of event; x-axis: month

Table 1. Summary Statistics

	Adopters		Non-Adopters		All Properties	
	Mean	**Std. Dev**	**Mean**	**Std. Dev.**	**Mean**	**Std. Dev.**
# Unique Properties	2118		7278		9396	
White (Race)	0.83	0.37	0.75	0.43	0.77	0.42
Black (Race)	0.07	0.25	0.12	0.32	0.11	0.31
Other (Race)	0.10	0.30	0.13	0.34	0.12	0.32
Age	35.87	10.18	35.36	10.00	35.48	10.04
Revenue	80.31	98.04	67.41	101.70	70.91	100.88
Occupancy Rate	0.50	0.38	0.38	0.39	0.41	0.39
# Reviews	42.88	51.64	28.07	41.78	31.65	44.82
# Photos	18.67	13.23	16.07	11.31	16.70	11.86
Super Host	0.23	0.42	0.16	0.36	0.17	0.38
Instant Book Enabled	0.18	0.38	0.10	0.30	0.12	0.32
Nightly-Rate	179.80	179.30	194.61	190.43	191.03	187.91
Security Deposit	163.63	331.45	140.76	310.24	146.29	315.65
# Minimum Stay	2.97	10.58	3.03	5.54	3.01	7.09
Response Rate (%)	94.06	12.81	92.08	15.40	92.65	14.73
# Host-Owned Listings	2.54	3.92	3.02	8.72	2.90	7.84

Notes: These summary statistics are computed for the adopters (column 1), non-adopters (column 2), and all properties in our sample (column 3). The values for the three race categories (White, Black, and Other) indicate the proportion of properties that belonged to a host of the given race.

Table 2. Impact of Pricing Algorithm Adoption on the Average Daily Revenue, Average Nightly Rate, and Monthly Occupancy Rate

VARIABLE	(1) Avg. Daily Revenue		(2) Avg. Nightly Rate		(3) Occupancy Rate	
	COEFF.	**Std. Err.**	**COEFF.**	**Std. Err.**	**COEFF.**	**Std. Err.**
SmartPricing	6.396***	(1.428)	-9.758*	(4.047)	0.0662***	(0.00526)
Log Number_of_Reviews	11.91***	(1.182)	5.264	(7.596)	0.0458***	(0.00424)
Log Number_of_Photos	6.810*	(2.940)	-49.13	(33.15)	0.0492***	(0.0134)
Log Security Deposit	6.091***	(0.254)	1.554***	(0.446)	0.0257***	(0.000937)
Log # Minimum Stays	-10.43***	(2.292)	-3.588	(3.534)	-0.0520***	(0.00693)
Instant_Book Enabled	11.03***	(1.465)	0.712	(2.060)	0.0636***	(0.00603)
Super Host	9.024***	(1.961)	4.546	(2.987)	0.0316***	(0.00598)
Log # Host-Owned Listings	0.365	(4.267)	4.903	(4.231)	-0.0132	(0.0114)
Host Effort (Response Rate (%))	0.0268	(0.0288)	0.103	(0.0653)	0.000533	(0.00146)
Fixed Effect	Property					
Seasonality	City-Year, City-Month					
Observations	162617		162617		162617	
R-squared	0.51		0.85		0.56	

Notes: The table reports the estimation results from Equation (1). The main DiD model regressed the average daily revenue (column 1), average nightly rate (column 2), and monthly occupancy rate (column 3) on the IPTW-weighted sample of 9,396 properties.
Cluster-robust standard errors are reported at the individual property level.
* $p < 0.05$ ** $p < 0.01$ *** $p < 0.001$

Table 3. Heterogeneous Effects of Algorithm Adoption on Average Daily Revenues, Average Nightly Rates, Occupancy Rates, by the Host's Race

VARIABLE	**(1) Avg. Daily Revenue**		**(2) Avg. Nightly Rate**		**(3) Occupancy Rate**	
	COEFF.	**Std. Err.**	**COEFF.**	**Std. Err.**	**COEFF.**	**Std. Err.**
SmartPricing	5.219***	(1.573)	-10.15*	(4.766)	0.0600***	(0.00569)
SmartPricing×Black	8.700*	(4.383)	5.030	(6.078)	0.0714***	(0.0200)
SmartPricing×Other	6.264	(3.684)	6.350	(5.627)	0.0139	(0.0150)
Log Number_of_Reviews	11.89***	(1.180)	5.245	(7.589)	0.0456***	(0.00424)
Log Number_of_Photos	6.819*	(2.934)	-49.12	(33.13)	0.0492***	(0.0134)
Log Security Deposit	6.088***	(0.254)	1.551***	(0.446)	0.0257***	(0.000938)
Log # Minimum Stays	-10.43***	(2.296)	-3.593	(3.534)	-0.0521***	(0.00694)
Instant_Book Enabled	11.05***	(1.471)	0.740	(2.064)	0.0638***	(0.00605)
Super Host	9.098***	(1.962)	4.633	(3.014)	0.0319***	(0.00599)
Log # Host-Owned Listings	0.390	(4.270)	4.874	(4.220)	-0.0131	(0.0114)
Host Effort (Response Rate (%))	0.0267	(0.0288)	0.103	(0.0653)	0.000533	(0.00146)
Fixed Effect	Property					
Seasonality	City-Year, City-Month					
Observations	162617		162617		162617	
R-squared	0.51		0.85		0.56	

Notes: Table reports the results from estimating Equation (2). This heterogeneous model regresses average daily revenue (column 1), average nightly rate (column 2), and monthly occupancy rate (column 3) on the IPTW-weighted sample of 9,396 properties. Cluster-robust standard errors are reported at the individual property level.
* $p < 0.05$ ** $p < 0.01$ *** $p < 0.001$