

NBER WORKING PAPER SERIES

SCAPEGOATING AND DISCRIMINATION IN TIMES OF CRISIS:
EVIDENCE FROM AIRBNB

Michael Luca
Elizaveta Pronkina
Michelangelo Rossi

Working Paper 30344
<http://www.nber.org/papers/w30344>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2022, Revised March 2023

Funding comes exclusively from our universities' research budgets. We would like to thank Andrei Shleifer and Katherine Baldiga Coffman for providing helpful feedback. All errors are our own. E. Pronkina acknowledges that this work was done prior to the author joining Amazon. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2022 by Michael Luca, Elizaveta Pronkina, and Michelangelo Rossi. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Scapegoating and Discrimination in Times of Crisis: Evidence from Airbnb
Michael Luca, Elizaveta Pronkina, and Michelangelo Rossi
NBER Working Paper No. 30344
August 2022, Revised March 2023
JEL No. J15

ABSTRACT

We present evidence that discrimination against Asian-American Airbnb users sharply increased at the start of the COVID-19 pandemic. Using a DiD approach, we find that hosts with distinctively Asian names experienced a 20 percent decline in guests relative to hosts with distinctively White names. In contrast, we do not see spikes in discrimination against Black or Hispanic hosts. Our results suggest that the rise in anti-Asian sentiment in 2020 translated to discrimination in economic activity, highlighting the ways in which scapegoating minority groups can shape markets. Our results also point to the role of platform design choices in enabling discrimination.

Michael Luca
Harvard Business School
Soldiers Field
Boston, MA 02163
and NBER
mluca@hbs.edu

Michelangelo Rossi
Télécom Paris
Département Sciences Économiques et Sociales
19 place Marguerite Perey
Palaiseau 91120
France
michelangelo.rossi@telecom-paris.fr

Elizaveta Pronkina
Universite Paris-Dauphine - PSL
Place du Marechal de Lattre de Tassigny
Paris 75016
France
elizaveta.pronkina@dauphine.psl.eu

A data appendix is available at <http://www.nber.org/data-appendix/w30344>

1. Introduction

The COVID-19 pandemic originated in Wuhan, China, and quickly spread to the rest of the world. Although the COVID-19 related death rate was ultimately much lower in China relative to Europe and North America, the pandemic witnessed an increase in anti-Asian sentiment in the United States. According to a Pew Research Center survey (Ruiz, Edwards and Lopez, 2021), almost half of Asian-American adults experienced at least one racist incident in the first year of the pandemic. Evidence from psychology research shows that anti-Asian sentiment and racial scapegoating grew with COVID-19 (Borja, Jeung, Horse, Gibson, Gowing, Lin, Navins and Power, 2020; Cheah, Wang, Ren, Zong, Cho and Xue, 2020; Tessler, Choi and Kao, 2020; Cheng, Kim, Tsong, Joel Wong et al., 2021).

High profile American conservative politicians and media also frequently associated the virus with China. In notes for a March 2020 speech, for instance, then President Donald Trump, crossing out the word “Corona” in front of virus and replacing it with “Chinese,” went on to refer to COVID-19 as “the Chinese Virus.” Trump was engaging in scapegoating, seeking to place blame for the rapid spread of COVID-19 throughout the United States on China. Google search trends, as seen in Appendix Figure A1, show China Virus and Wuhan Virus to have been particularly common search terms early in the pandemic. This rhetoric mirrored a broader rise in anti-Asian sentiment during the pandemic. Resulting concerns about discrimination against Asian-Americans and the formation of stereotypes associating Asian-Americans with COVID-19 are consistent with a recent body of evidence associating periods of crisis with increased scapegoating and elevated antipathy towards minority groups (Fisman, Hamao and Wang, 2014; Fouka and Voth, 2016; Cantoni, Hagemeister and Westcott, 2019; Bauer, Cahliková, Chytilová, Roland and Zelinsky, 2021; Zussman, 2021; Bursztyrn, Egorov, Haaland, Rao and Roth, 2022).

In this paper, we use the COVID-19 outbreak in 2020 to study how salient events such as a pandemic can at times foster discrimination against a minority group in consumption decisions. To do so, we answer the following question: Did COVID-19 lead to discrimination against Asian-Americans in consumption decisions? We explore this question by looking at short-term housing rentals on Airbnb in New York City. We use a difference-in-differences (DiD) approach to compare the number of guests staying with Asian-American relative to other Airbnb hosts, during the pandemic relative to the previous year.

Should we expect an increase in discrimination? In practice, Asian-Americans have had a lower death rate from COVID-19 relative to White, Black, and Hispanic Americans.¹ Thus, an increase in discrimination would likely not reflect accurate statistical discrimination based on the actual risk of contagion. Yet, a

¹ According to a report by the Kaiser Family Foundation (Latoya and Samantha, 2022), Asian-Americans are the ethnic group with the lowest share of COVID-19 cases and deaths compared to their population share.

growing literature (Jensen, 2010; Bursztyn, González and Yanagizawa-Drott, 2020; Bohren, Haggag, Imas and Pope, 2019) has shown that discrimination can be often rooted in inaccurate beliefs, in contrast with early models of statistical discrimination (Arrow, 1973; Aigner and Cain, 1977). Beliefs about the risks involved in different activities during the pandemic have been noisy (Bundorf, DeMatteis, Miller, Polyakova, Streeter and Wivagg, 2021; Bordalo, Burro, Coffman, Gennaioli and Shleifer, 2022). Stereotypes simplify the representation of groups, which can lead to systematic errors in judgment (Tversky and Kahneman, 1983; Bordalo, Coffman, Gennaioli and Shleifer, 2016). Anti-minority sentiment can be further fueled by political rhetoric and scapegoating (Glaeser, 2005; Bauer et al., 2021). The modern literature on scapegoating dates back at least to the Dreyfus Affair in France in 1899 (Durkheim, 1899), and has been studied within psychology since at least the mid twentieth century (Allport, Clark and Pettigrew, 1954).

Looking at panel data on Airbnb, we find that discrimination against Asian-American hosts sharply increased at the start of the COVID-19 pandemic. After January 2020, hosts with distinctively Asian names, relative to those with White-sounding names, experienced a 20 percent decline in guests as proxied by reviews. An event-study analysis shows Asian and White Airbnb hosts to follow comparable trends during 2019, with a drop in number of guests for Asian hosts first noticed in the spring of 2020 and persisting throughout the year.

Our results are robust to the inclusion of listing and time fixed effects and when we use hosts' pictures to predict their race. To separate anti-Asian bias from changing preferences across neighborhoods (which might be correlated with neighborhood demographics), we use time interacted with neighborhood fixed effects, and coarsened exact matching to ensure comparable properties in the treated (Asian) and control (White) groups. Listings are matched using dwelling characteristics, the pre-pandemic share of Airbnb guests with Asian-sounding names and pre-pandemic official statistics from the American Community Survey about the area in which each listing is located.

As Airbnb allows hosts to reject guests' booking requests, Asian hosts may have reacted to the pandemic by accepting fewer guests, which may account for the reduced number of transactions. To assess this hypothesis, we replicate our analysis using only hosts that have activated the "instant bookable" option. With this option on, hosts automatically accept all guest requests. Neither the sign nor size of the DiD estimates change for this subsample. Moreover, Asian hosts do not charge higher prices than White hosts after January 2020. The drop in the number of guests is thus not driven by changes on the platform supply.

The COVID-19 pandemic has led to a significant reduction in the number of travelers from Asia to the US in 2020. Thus, if Asian travelers are more likely to seek Airbnb hosts, the negative effect that we observed could be the result of a drop in the specific demand for Asian hosts. To test this "host-guest

homophily hypothesis”, we exploit guests’ names to infer the racial composition of each host’s demand. Before COVID-19, guests are more likely to match with hosts belonging to the same race. Yet, the share of Asian guests matched with Asian hosts remains stable during 2020. This suggests that the reduction in the activities suffered by Asian hosts is not primarily due to the absence of Asian guests.

As an additional falsification exercise, we observe numbers of guests during the pandemic for Black and Hispanic hosts. The absence of spikes in discrimination against Black or Hispanic hosts relative to pre-pandemic levels further suggests that our results reflect increased anti-Asian discrimination during the pandemic instead of broad discrimination against hosts from minority groups. We also repeat the analysis by comparing Asian hosts (treated group) with Hispanic and Black hosts (control group), and the statistically significant drop in the number of reviews is confirmed.

Airbnb makes a compelling setting in which to study our research question. Airbnb is of direct interest as the world’s largest short-term accommodation provider, with a market capitalization of 70 billion dollars as of 2022. From a research perspective, online platforms such as Airbnb have important advantages in exploring economic phenomena. Our data track all listings for the year before and after the pandemic’s beginnings in New York City, a major market for Airbnb, enabling us to observe outcomes at the listing level that provide granular insight into booking patterns.

Overall, our results speak to the impact of crises such as the COVID-19 pandemic on discrimination. COVID-19 represented an unprecedented shock in terms of health, job losses, and productivity. Our analysis provides a window into economic activity in one of the world’s largest marketplaces, and demonstrates that the rise in anti-Asian sentiment early in the pandemic translated to significant discrimination in economic activity. Our work contributes to the literature on discrimination and highlights the role of crises in triggering discrimination against scapegoated groups in market settings.

Our results also contribute to the literature on market design. As a platform, Airbnb has made design choices that enable discrimination (Edelman, Luca and Svirsky, 2017; Kakar, Voelz, Wu and Franco, 2018; Ahuja and Lyons, 2019; Ameri, Rogers, Schur and Kruse, 2020; Cui, Li and Zhang, 2020; Laouénan and Rathelot, 2022). For instance, Airbnb makes the name (and the picture) of property owners salient in the process of choosing potential stays. This is in contrast with platforms such as Expedia. By making the ethnicity of users salient before a booking decision is made, Airbnb permits discrimination that would be more difficult had they not showed identifying information about users until after a booking is complete (Fisman and Luca, 2016). Our results provide novel insight into the ways in which increased racial bias can creep into a marketplace with discrimination-permissive design in times of crisis. As a market designer, Airbnb makes important decisions that shape the susceptibility of the platform to shocks that affect discrimination.

2. Empirical Context and Identification Strategy

The data comes from Inside Airbnb, a website tracking Airbnb listings present in many cities over time. We use data about New York City, one of the largest markets for Airbnb in the US, with no relevant policies affecting short-term rentals in its metropolitan area during 2019 and 2020. Our data is composed of monthly snapshots of Airbnb listings present in New York City from January 2019 to November 2020.^{2,3} In total, our data is composed of 23 snapshots. Every month, we observe which listings appear on the Airbnb website at the snapshot date. For each listing observation, we have a set of characteristics such as the listing’s location (longitude and latitude); dwelling’s characteristics; the name and the picture of the host managing the listing; the total number of reviews posted by guests in the past; the names of all guests that have posted a comment; the average star ratings, and the per-night price. We complement this dataset with external official statistics. We use the official neighborhoods from the Department of City Planning of NYC and we assign listings to neighborhoods. Moreover, we add external statistics dated to 2018 (to isolate the endogenous changes to the timing of COVID-19) provided by U.S. Census Bureau in the American Community Survey (ACS) about the census tract in which each listing is located.

We proxy the number of guests staying with each listing in a specific month by the difference in the number of reviews posted between two consecutive snapshots. Thus, we restrict our analysis over listings that are present on the platform for at least two snapshots. Listings on Airbnb may be present on the platform, but not be ready to rent their apartment to guests. We remove this “inactive” part of the market including only listing-snapshot observations if listings have at least one review in the period of analysis (from the beginning of 2019 to the end of 2020) and when they are available for rent in the next 365 days.⁴

In the main analysis, we use information about hosts’ and guests’ names to measure their race.⁵ We apply a machine-learning algorithm (Ye, Han, Hu, Coskun, Liu, Qin and Skiena, 2017; Ye and Skiena, 2019) to identify the probability of a name belonging to different races.⁶ We focus on the four most common races: Asian, Black, Hispanic, and White. We consider a host or a guest to belong to one race if the name-based probability of being of that race is greater than 0.9 - and vary this threshold in our robustness checks.⁷

²A short-term rental regulation enacted by the New York City council, goes into effect in January 2021. We observe a drop in the number of Airbnb listings starting from December 2020 and we restrict our analysis until November 2020.

³The results of our analysis are similar when we study listings located in Los Angeles (see Figure A6). Yet, several short-term rental regulations have been implemented in the county of Los Angeles during the time of the analysis and we prefer to focus on New York City.

⁴The significance and the size of the effects remain the same using alternative sample restrictions, e.g. listings with at least one review in the previous 30, 60 or 90 days. In Appendix, Table A3 and Figure A2 present the results with the alternative restrictions.

⁵We consider only the first listed name when hosts or guests include several names (less than 5 percent of all hosts).

⁶We highly appreciate access to the NamePrism classifier algorithm provided by Prof. Steven Skiena: <http://name-prism.com/about>.

⁷Only 0.3 percent of all Airbnb hosts changed their name in a way altering the name-based prediction. We remove those listings

The aim of this study is to estimate the magnitude of the discrimination against Asian hosts and the associated drop in the number of reviews during the first year of the COVID-19 pandemic. To do so, we compare Asian and White hosts in 2019 and 2020. Before turning on the description of the DiD strategy, we briefly present descriptive statistics about these two groups. In Appendix Table A1, we present summary statistics about all hosts active on the platform between January 2019 and November 2020; and only hosts with distinctive Asian and White names. There are 914 hosts (2.5 percent of the total) with distinct Asian names and 15,730 hosts (41 percent) with distinct White names in New York City during this period. Asian hosts tend to receive more reviews than White hosts. They also tend to select the “instant bookable” option more often. Accordingly, they are more likely to accept guests’ booking requests: this may partially explain the higher number of reviews for this group. Hosts tend to be located in areas with a larger percentage of residents of their race. White hosts are located in areas with fewer minority groups relative to Asian hosts, and in particular, Black and Hispanic hosts, as in Appendix Table A2. In Figure 1, we illustrate the evolution of the number of reviews for Asian and White hosts during 2019 and 2020. On average, Asian hosts tend to have more reviews than White hosts before January 2020. The drop in the demand due to the COVID-19 outbreak affects all Airbnb hosts, but it is much larger for Asian hosts. In particular, after May 2020, White hosts have slightly more reviews than Asian hosts. This preliminary analysis about the number of reviews over time presents the main result to be formalized with our empirical strategy: Asian hosts suffer a strong reduction in their economic activity on the platform during the first months of the COVID-19 pandemic.

To capture discrimination against Asian hosts during the COVID-19 pandemic, we use the following DiD specification:

$$Reviews_{it} = \alpha_1 + \alpha_2 Asian_i + \alpha_3 After_t + \delta Asian_i \times After_t + \beta X_{it} + \varepsilon_{it}, \quad (1)$$

where $Reviews_{it}$ is the number of reviews listing i receives between snapshot t and the previous snapshot in which listing i was present on the platform. In our preferred specification, $Asian_i$ is a dummy variable equal to 1 if the probability to be Asian for the host managing listing i is greater than 0.9; and equal 0 if the probability to be White is greater than 0.9. $After_t$ is equal to 1 for all snapshots after January 2020; and equal to 0 otherwise. The set of controls X_{it} includes listing fixed effects, neighborhood-by-time fixed effects, neighborhood-specific time trends and separate trends for Asian and White hosts.

The coefficient δ captures discrimination against Asian hosts under the assumption that White hosts provide a good counterfactual for the evolution of the number of guests during the COVID-19 pandemic.

from the analysis.

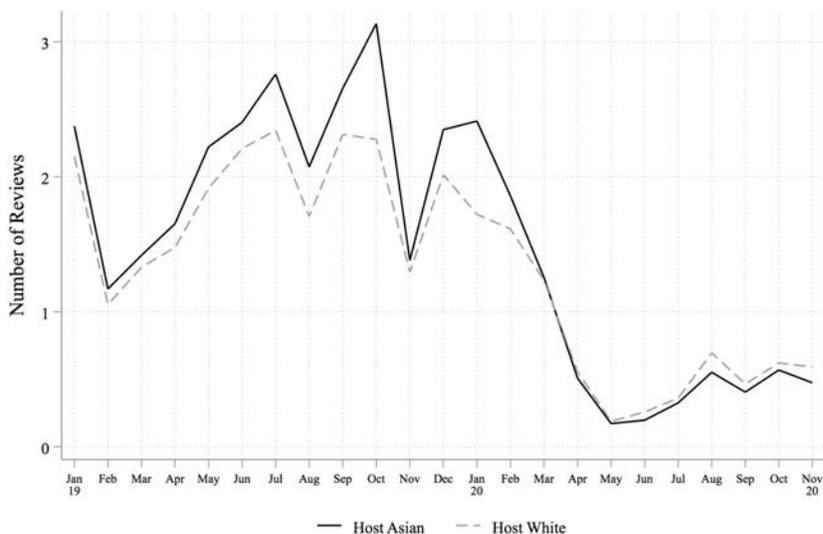


Figure 1. Average Number of Monthly Reviews over Time

Notes: The graph plots the average difference in the number of reviews posted between two snapshots for all Airbnb listings in New York City managed by hosts with the name-based probability to be Asian or White greater than 0.9, active for at least two snapshots between January 2019 and November 2020.

Accordingly, we assume guests’ propensity to write reviews does not vary for Asian hosts relative to White hosts in response to the outbreak of COVID-19. To test the absence of potential pre-trends between treated and control groups, we illustrate the evolution of $Reviews_{it}$ over time with an event-study approach. We consider the following lead-lag model in which $Reviews_{it}$ is regressed over the product between the dummy $Asian_i$ and a full set of dummy variables for each snapshot. The model controls for listing fixed effects; neighborhood-by-time fixed effects and neighborhood-specific time trends:

$$Reviews_{it} = \alpha_i + \rho_t + \sum_{\tau=Jan19}^{Nov20} \delta_{\tau} Asian_i \times 1(t = \tau) + \varepsilon_{it}. \quad (2)$$

We present the results of the estimates of Equation 2 in Figure 2 where we plot the estimated δ_{τ} from January 2019 to November 2020. Before January 2020, the coefficients are close to zero and they do not exhibit a clear trend. Thus, the evolution of the number of reviews for listings managed by Asian and White hosts was similar before the pandemic. This finding supports the parallel trend assumption which is necessary for our analysis. Instead, after January 2020, the coefficients are negative and the number of reviews drops for Asian hosts relative to White hosts. The effect is persistent over 2020. Yet, the reduction is particularly notable in Spring 2020 when the association between COVID-19 and China was the strongest.

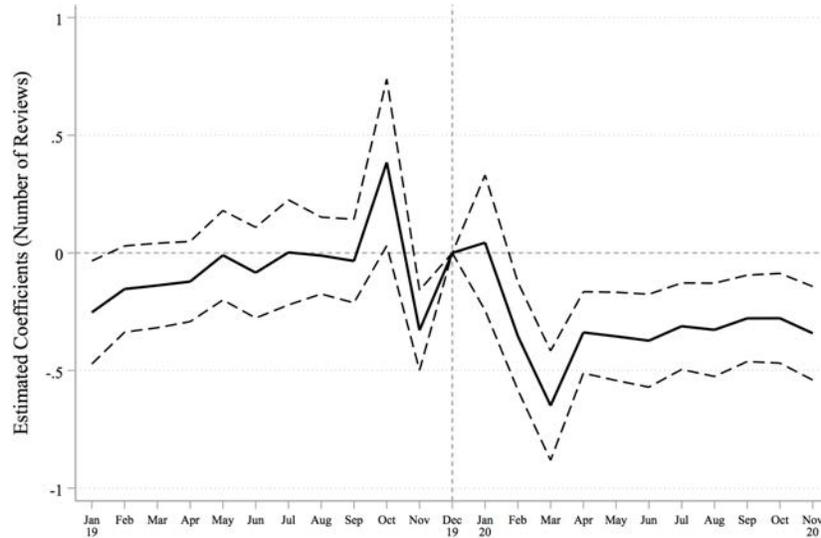


Figure 2. Event study: Number of Reviews - Asian and White Hosts

Notes: In line with Equation 2, $Reviews_{it}$ is regressed on listing fixed effects; neighborhood-by-time fixed effects, neighborhood-specific time trends, and on the products between $Asian_i$ and a full set of dummy variables for each snapshot. The graph plots the estimated coefficients on these products. The value of the coefficient corresponding to December 2019 is normalized to zero. The sample includes snapshots between January 2019 and November 2020 for listings located in New York and managed by hosts with a name-based probability of being Asian or White greater than 0.9. Standard errors (10%) are clustered by listing.

3. Impact on Asian-American Hosts

3.1. Main Results

We now present the main empirical results. Table 1 shows four specifications for the DiD estimates in Equation 1. The outcome variable is the number of reviews received between two snapshots. The sample includes snapshots between January 2019 and November 2020 for listings located in New York City and managed by hosts with a name-based probability of being Asian or White greater than 0.9. We cluster standard errors at listing level.⁸

In Column (1) we use listing fixed effects, time fixed effects, and a linear trend. The listing fixed effects remove all time-invariant elements such as dwellings' and geographical area characteristics; the time fixed effects and the time trend account for time-varying confounders that affect all listings in the same way. For instance, the COVID-19 outbreak and the related lockdowns have determined a sharp drop in demand for all listings in New York. However, this drop might not be the same across neighborhoods in the city. In Column (2) we include neighborhood-by-time fixed effects to allow for differential time variations in the demand for short-term rentals across neighborhoods. On top of mentioned controls, in Column (3) we add separate

⁸The results are robust to clustering the errors at a neighborhood level or at a Census Survey area level.

Table 1. Difference-in-Differences: Number of Reviews - Asian and White Hosts

	(1)	(2)	(3)	(4)
Host Asian \times After Jan 20	-0.291*** (0.077)	-0.264*** (0.076)	-0.327*** (0.107)	-0.255** (0.111)
Listing FEs	✓	✓	✓	✓
Time FEs	✓			
Neighborhood \times Time FEs		✓	✓	✓
Race-Specific Time Trends			✓	✓
CEM				✓
R^2	0.488	0.504	0.504	0.506
N	231,416	231,416	231,416	199,800
Mean Dep. Var.	1.403	1.403	1.403	1.460

Notes: The sample includes snapshots between January 2019 and November 2020 for listings located in New York and managed by hosts with a name-based probability of being Asian or White greater than 0.9. “Neighborhood \times Time FEs” include neighborhood-by-time fixed effects and neighborhood-specific time trends. “Race-Specific Time Trends” include separate time trends for Asian and White hosts. The coarsened exact matching (“CEM”) in Column (4) is based on time-invariant listing characteristics, the share of Asian Airbnb guests in the area before the pandemic, and official statistics from the American Community Survey in 2018 about the area where a listing is located. Standard errors clustered by listing are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

time trends for Asian and White hosts to sort out the demand effect after January 2020 from medium-run tendencies of different ethnic groups. The estimates of the coefficient δ in Equation 1 are negative and significant for all specifications. The drop in the number of reviews experienced by Asian hosts is equal to 0.291 reviews, which accounts for more than a 20 percent decline relative to the mean.

Asian and White hosts may differ, as suggested by descriptive statistics in Section 2. We use coarsened exact matching (Iacus, King and Porro, 2012) to make sure treated and control units are comparable. For the matching, we use time-invariant listings characteristics. To measure the attractiveness of each city area for Asian guests, we use the share of Airbnb guests with Asian-sounding names in all listings located in the same census tract before January 2020. Finally, we use official statistics from the Census American Community Survey in 2018 about the area where a listing is located.⁹ By doing so, almost 14 percent of the sample is removed due to imbalance between Asian and White hosts. In Column (4) we repeat the specification from Column (3) for the balanced subsample. This analysis confirms the evidence of Anti-Asian discrimination with a significant drop in the number of reviews suffered by Asian hosts.

We perform multiple robustness checks to support the estimated decline in the number of transactions

⁹The listing characteristics used for the matching are the following: whether the listing is used for short-term rentals (guests can stay for less than 30 consecutive days); the number of bedrooms; and whether the area in which a listing is located was visited by a percent of Airbnb guests with Asian-sounding names greater than average before the pandemic (the correlation between this measure and the share of Asian in the census tract is 0.53). The area’s official statistics include the median income and the percentage of Asian residents of the census tract where the listing is located.

experienced by Asian hosts. Thus far, we have labeled as Asian or White, hosts whose name-based probability of corresponding to either race group exceeds 0.9. Our results are robust to using different cutoff levels to determine host race. Appendix Table A4 shows DiD estimates for cutoff levels varied from 0.95 to 0.8. Although sample size decreases with narrower cutoffs, names remain a precise signal of host race that enables guests to more easily identify and discriminate against Asian hosts. As a result, the coefficients δ become larger when we use the strictest cutoff of 0.95. Regardless of the chosen level, the DiD estimate remains statistically significant at a 1 percent level.

Additionally, we predict hosts' race based on their pictures using a facial attribute analysis (Serengil and Ozpinar, 2020; and Serengil and Ozpinar, 2021). We consider a host being Asian or White if the host's picture-based probability of being Asian or White is higher than 0.9. When we use this new measure to define treated and control groups, the sample size almost halves. This is in line with the platform's policy that does not require hosts to post a real picture of themselves on the platform. Thus, we may expect to identify fewer hosts with a high picture-based probability of being of a specific race. Yet, the results remain in line with the name-based analysis. Figure A3 in Appendix shows the event study analysis and confirms the absence of pre-trends before 2020 and the drop in the number of reviews for Asian hosts compared to White hosts at the beginning of 2020. To ease the comparison of the magnitude, Appendix Table A5 presents the DiD estimates using only name-based, only picture-based, and both name-based and picture-based predictions. The effect using only picture-based probabilities is comparable in terms of significance and size relative to one using name-based probabilities. However, when we use both names and pictures to predict race, then the size of the effect almost doubles. This suggests that the drop in the number of transactions is the largest for hosts with Asian-sounding names and Asian-looking pictures.

Next, we study if the drop in the number of reviews is driven by the selection of hosts over time. Asian and White hosts enter and exit the platform over time. Thus, variations in the profile of listings that are managed by Asian and White hosts may partially explain our effect. We account for listing selection by restricting our analysis to hosts present on Airbnb before and after January 2020. With this subsample, we match listings using pre-pandemic time-varying listing characteristics like price, rating, and the share of guests with Asian-sounding names. The sample size decreases by almost 50 percent. Yet, the size and significance of the DiD estimates remain similar to the baseline 20 percent drop, as in Appendix Table A6.

Thus far, we have used as the main outcome variable the difference in the number of displayed reviews on a host's webpage over time. Reviews are strongly encouraged but they are not mandatory (roughly 70 percent of stays are reviewed according to Fradkin, Grewal and Holtz, 2021). While review rates may have changed during COVID, our identification still holds unless this change varies across Asian and White hosts.

We use alternative measures of guest demand in Appendix Table A7 that are all robust to the findings in the main analysis. Column (1) repeats for comparison the specification in Table 1; Column (2) uses the number of written comments posted by guests on a listing’s webpage over time; Column (3) adjusts the number of reviews for the number of days between two snapshots. We confirm Asian hosts to experience a negative and significant shock in demand measured in terms of comments or reviews per day. The size of the effects is similar to the baseline estimate.

On a similar point, we study how many nights in the next 30 days are still available to be rented for each listing. A night can become unavailable because it is rented by a guest or because a host decides to deactivate the listing on some specific dates. Accordingly, by subtracting the number of available nights from 30, we define the upper bound of the number of rented nights in the next 30 days for each listing. In Appendix Figure A4 we present an event study analysis as in Equation 2. Also with this proxy, we observe a clear decrease in the number of rented nights for Asian hosts from February to July 2020. Yet, the timing of the decline does not entirely match the decrease in the number of reviews. This difference can stand from the imperfect nature of the proxy used for the number of rented nights, and the mechanical lag between booking and posting a review.

To shed further light on the nature of discrimination in our setting, we analyzed heterogeneous effects in Appendix Figure A5, and found discrimination to persist among hosts with higher and lower ratings, among “superhosts” (those with more experience on the platform), between less and more expensive listings, and among hosts renting out their entire space (as opposed to sharing the listing). This suggests that bias is experienced even by established Airbnb Asian hosts and in settings in which interaction with hosts is minimal. We also check if the results apply to other U.S. cities. Figure A6 in Appendix confirms an evident and statistically significant discrimination against Asian-American hosts in the county of Los Angeles.

3.2. Supply-Side Factors

“Instant Bookable” Hosts. The main results document a significant reduction in the number of reviews suffered by Asian hosts in the first months of 2020. We interpret this result as direct evidence of the surging Anti-Asian discrimination. However, Asian hosts may have reacted to the pandemic, being less willing to accept guests after the outbreak of COVID-19.

To test this hypothesis, Appendix Figure A7 presents the number of Asian and White hosts with the “instant booking” option on and off over time. Airbnb hosts can always activate or deactivate the “instant booking” option. “Instant bookable” listings automatically accept all guests’ requests. With “instant booking” off, guests’ requests to stay are subject to the host’s approval. Although we do not see a massive switch

off of the “instant booking” option right after January 2020, the differential changes in the “instant booking” status between Asian and White hosts over time could partially explain the drop in the number of reviews. “Instant bookable” hosts tend to have more reviews than “non-instant bookable” hosts. In particular, in the year before COVID, Asian “instant bookable” hosts had on average 0.57 more reviews relative to Asian “non-instant bookable” hosts. This is also the case for White hosts, with 0.65 more reviews. The share of “instant bookable” hosts decreased after COVID-19 for both groups, with a 5.3 and 4.0 percentage point drop for Asian and White hosts, respectively. Thus, a rough calibration based on the pre-COVID differences suggests that the reduction in the number of reviews due to the change in “instant bookable” status can only account for 0.004 reviews (equal to less than 2 percent of the detected effect).

An alternative way to check whether this supply-side channel drives our results is to focus our analysis on Airbnb listings that were “instant bookable” during the whole period under analysis. We present the results of the event study for only “instant bookable” hosts in Figure 3 where we plot the estimated δ_τ of Equation 2 from January 2019 to November 2020. The sample size decreases from 231,416 to 66,644 observations. Despite this data loss, the drop in the number of reviews by Asian hosts is still observable, with a notable spike during Spring 2020. Appendix Table A8 reports the DiD results. The reduction in the number of reviews suffered by Asian hosts who cannot reject guests’ requests is statistically significant and of a larger magnitude relative to the baseline specification: a decline of 26 percent in their number of reviews versus a 20 percent reduction for all listings. Accordingly, “non-instant bookable” Asian hosts were less penalized during the pandemic than their “instant bookable” counterparts. A possible interpretation of this result is that “non-instant bookable” Asian hosts became more willing to accept guests than White hosts after the pandemic. Facing a hostile environment with a drop in demand, Asian hosts became less selective to reduce the adverse effects of discrimination.

Prices. Can the drop in the number of reviews for Asian hosts be due to their pricing strategy? If Asian-American hosts increased their prices after COVID-19 more than other groups, they might be less likely to attract guests, resulting in fewer reviews. To explore this hypothesis, we restrict our analysis to hosts present before and after January 2020 and match listings using pre-pandemic average prices and ratings (Appendix Table A6, Column 4). The drop in the number of reviews remains significant. When we control for prices in all specifications in Table 1, the DiD estimates remain similar in size and significance. Moreover, we repeat the DiD analysis for the per-night prices. Appendix Table A9 shows that the DiD estimates are not statistically significant in the first two specifications and are negative and significant when we add race-specific trends and we use coarsened exact matching. Although inconclusive, this result suggests that Asian hosts may have partially reacted to the drop in their operations by reducing prices.

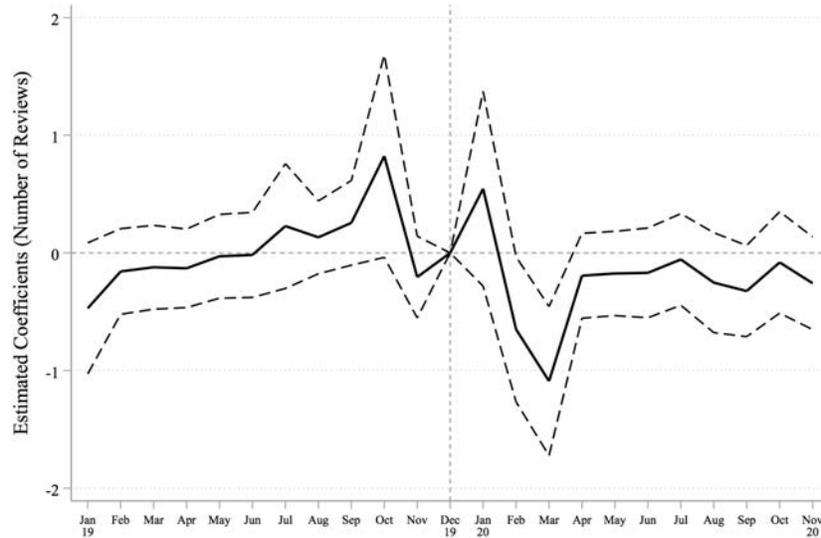


Figure 3. Event study: Number of Reviews - “Instant Bookable” Asian and White Hosts

Notes: In line with Equation 2, $Reviews_{it}$ is regressed on listing fixed effects; neighborhood-by-time fixed effects, neighborhood-specific time trends, and on the products between $Asian_i$ and a full set of dummy variables for each snapshot. The graph plots the estimated coefficients on these products. The value of the coefficient corresponding to December 2019 is normalized to zero. The sample includes snapshots between January 2019 and November 2020 for listings located in New York and managed by hosts with a name-based probability of being Asian or White greater than 0.9. We restrict to hosts who are always “instant bookable” during the period of analysis. Standard errors (10%) are clustered by listing.

3.3. Host-Guest Homophily Hypothesis

If Asian guests are more likely to seek out Asian hosts, our results could reflect homophily - coupled with a disproportionate drop in Asian guests on Airbnb - rather than bias. This homophily would suggest that 1) before COVID, Asian hosts were more likely to match with Asian guests; 2) after COVID, Asian tourists cannot travel to the US, and, as a consequence, Asian hosts suffered a drop in the number of guests.

To address these two points, we use the names of guests to infer their race and to measure the racial composition of each host. In Appendix Table A10 we present the share of guests from different races among hosts before and after COVID-19. *Panel I* of Appendix Table A10 confirms the presence of a modest degree of homophily on Airbnb before COVID-19: guests are more likely to match with hosts belonging to the same race. However, hosts of all races are matched with White guests or guests with names that are not very predictive of their race in more than 90 percent of the cases. Therefore, even if all Asian guests exited the platform in 2020, the impact on Airbnb hosts (Asian or not) would have been limited.

In Appendix Figure A8, we show changes in the share of reviews written by Asian guests on Airbnb over time. The share of Asian guests is lower after COVID-19. This may be due to a relatively stronger impact of the travel restrictions on tourism coming from Asia. Has this change disproportionately affected

Asian hosts? In *Panel II* of Appendix Table A10, we observe that the share of Asian guests matched with Asian hosts remains stable after COVID-19. Conversely, the proportion of White guests and guests from other racial minorities decreased among Asian hosts. To reinforce this last point, in Appendix Figure A9, we compare the share of Asian guests over time for Asian and White hosts with an event study analysis in line with the one in Equation 2 of the main text. Here the main response variable is a dummy variable equal to 1 if a guest is Asian (if the name-based probability to be Asian is larger than 0.9). The results show that Asian hosts do not seem to have suffered from a reduction in the number of Asian guests during COVID-19. The proportion of Asian guests is stable over time (if anything, slightly increasing during the Summer of 2020). These pieces of evidence suggest that the reduction in the activities suffered by Asian hosts during the pandemic is not primarily due to the absence of Asian guests.

3.4. Impact on Black-American and Hispanic-American Hosts

The discrimination narrative during the COVID-19 outbreak was mainly directed against Asians. Other minorities were not the focus of the scapegoating rhetoric meant to associate the spread of the virus with China. We ran a falsification test to check whether the reduction in the number of guests experienced by Asian hosts was the result of rising anti-Asian sentiment in the United States or reflected generalized antipathy towards all minorities. We asked to what extent the decline in demand during 2020 was present only among Asian hosts and whether other minorities experienced a similar drop in number of transactions. To answer these questions, we used the same DiD identification strategy with other minorities as the treated group. In line with the baseline specification, we label hosts as Black or Hispanic if the name-based probability of being of the corresponding race is greater than 0.9. As before, the control group includes White hosts.

Table 2 shows the results. To facilitate comparison of the size of the effects, Column (1) is the baseline estimate for Asian hosts, Columns (2) and (3) the results for Black and Hispanic hosts, respectively, in the treated group. That neither estimate is statistically significant suggests that hosts from other minorities did not experience a reduction in number of guests compared to White hosts. This result is consistent with our conclusion that Asian-American hosts on the platform experienced increased discrimination relative to other groups.

We reinforce this latter point by showing that Asian hosts experienced a reduction in number of reviews relative to other minorities as well during the pandemic. We do this by repeating the baseline DiD design using Black and Hispanic hosts as the control group. Appendix Table A11 shows the result varying the minority group used as a control. Again, to facilitate comparison, Column (1) reports the baseline specification (with White hosts as the control), and Columns (2), (3), and (4), respectively, with Black, Hispanic, and

Table 2. Falsification Test: Number of Reviews - Different Minorities and White Hosts

	(1)	(2)	(3)
Host Asian \times After Jan 20	-0.264*** (0.076)		
Host Black \times After Jan 20		-0.095 (0.122)	
Host Hispanic \times After Jan 20			-0.114 (0.082)
Listing FEs	✓	✓	✓
Neighborhood \times Time FEs	✓	✓	✓
R^2	0.504	0.502	0.498
N	231,416	222,527	228,566
Mean Dep. Var.	1.403	1.394	1.401
Treated Group	Asian	Black	Hispanic

Notes: The sample includes snapshots between January 2019 and November 2020 for listings located in New York and managed by hosts with a name-based probability of being White, or Asian (Column 1), Hispanic (Column 2), and Black (Column 3) greater than 0.9. “Neighborhood \times Time FEs” include neighborhood-by-time fixed effects and neighborhood-specific time trends. Standard errors clustered by listing are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Black and Hispanic hosts combined as the control. There being fewer minority hosts on Airbnb, the sample size decreases significantly for all specifications. When we use only Black hosts as the control, the results are not significant. However, the power of this specification is severely limited by the smaller number of Black hosts (fewer than 200 during the entire period of analysis). Asian hosts experienced a significant drop in the number of reviews when we use Hispanic hosts and Black and Hispanic hosts combined as the control group. For this last specification, Appendix Figure A10 shows the pretend assumption to be satisfied.

4. Conclusion

Crises ranging from wars to pandemics can trigger heightened anti-minority sentiment and scapegoating. Antipathy toward outgroups during times of crisis is not new, and often exacerbated by political actors (Glaser, 2005). For example, the wake of the September 11 attacks saw a rise in anti-Muslim and anti-Middle Eastern sentiment, evidenced, for instance, in a decline in investments in mutual funds managed by Middle Eastern fund managers (Kumar, Niessen-Ruenzi and Spalt, 2015). Our work demonstrates the ways in which Asian-Americans - a scapegoated group in the COVID-19 pandemic - suffered increased discrimination in market transactions in the world’s largest short-term rental marketplace. Previous literature has shown how endemic discrimination is in labor and housing markets (Bertrand and Mullainathan, 2004;

Oreopoulos, 2011; Christensen, Sarmiento-Barbieri and Timmins, 2022), and in various online settings, including Airbnb (Edelman et al., 2017; Kakar et al., 2018). This paper documents the evolving nature of discrimination. In particular, we show that scapegoating alters agents' decisions in a market with an observable negative effect on the scapegoated group.

Our analysis also contributes to a large psychology literature that has explored discrimination, scapegoating, and social comparisons (Lindzey, 1950; Wills, 1981; Rothschild, Landau, Sullivan and Keefer, 2012). Recent theory and lab work have highlighted mechanisms that have the potential to increase discrimination during times of crisis. Crises can “activate” antipathies and prejudice against minorities. They could lead people to seek out “someone to blame” creating favorable circumstances for scapegoating narratives (Bursztyn et al., 2022). During crises, people may form inaccurate beliefs driven by a lack of relevant information. Crises often present new scenarios with considerable uncertainty, which can lead people to rely on associative recall to make decisions (Bordalo et al., 2022). In a model of associative recall, associations between the pandemic and Asians might cue guests to think more about COVID, with the potential to exacerbate discrimination. The fact that the virus was in China before the United States may have contributed to inaccurate beliefs about the virus's ultimate toll. Our analysis points to the potential importance of these forces in managerial and policy relevant settings.

Our results also have important implications for market designers. In contrast with traditional travel websites such as Expedia, Airbnb makes a host's race salient. In response to Edelman and Luca (2014) and Edelman et al. (2017), Airbnb reduced the salience of a host's picture, by removing it off of the initial search results page. However, the name and photograph of a host can still be seen on the main listing page, making it easier for users to discriminate. Our results highlight that design choices that are permissive of discrimination become particularly vulnerable to discrimination when there are spikes in anti-minority sentiment. If Airbnb were to only show a host's name after a booking is made, there would be less scope for discrimination in the booking decision.

References

- Ahuja, Rishi, and Ronan C. Lyons (2019) 'The silent treatment: Discrimination against same-sex relations in the sharing economy.' *Oxford Economic Papers* 71(3), 564–576
- Aigner, Dennis J., and Glen G. Cain (1977) 'Statistical theories of discrimination in labor markets.' *Ilr Review* 30(2), 175–187
- Allport, Gordon Willard, Kenneth Clark, and Thomas Pettigrew (1954) 'The Nature of Prejudice'
- Ameri, Mason, Sean Edmund Rogers, Lisa Schur, and Douglas Kruse (2020) 'No room at the inn? Disability access in the new sharing economy.' *Academy of Management Discoveries* 6(2), 176–205
- Arrow, Kenneth J. (1973) 'The theory of discrimination.' In 'Discrimination in labor markets' (Princeton University Press) pp. 1–33
- Bauer, Michal, Jana Cahlíková, Julie Chytilová, Gérard Roland, and Tomas Zelinsky (2021) 'Shifting punishment on minorities: Experimental evidence of scapegoating.' Technical Report, National Bureau of Economic Research
- Bertrand, Marianne, and Sendhil Mullainathan (2004) 'Are emily and greg more employable than lakisha and jamal? a field experiment on labor market discrimination.' *American Economic Review* 94(4), 991–1013
- Bohren, J. Aislinn, Kareem Haggag, Alex Imas, and Devin G. Pope (2019) 'Inaccurate statistical discrimination: An identification problem.' Technical Report, National Bureau of Economic Research
- Bordalo, Pedro, Giovanni Burro, Katie Coffman, Nicola Gennaioli, and Andrei Shleifer (2022) 'Imagining the Future: Memory, Simulation and Beliefs about COVID.' Technical Report, National Bureau of Economic Research
- Bordalo, Pedro, Katherine Coffman, Nicola Gennaioli, and Andrei Shleifer (2016) 'Stereotypes.' *The Quarterly Journal of Economics* 131(4), 1753–1794
- Borja, Melissa, Russell Jeung, Aggie Yellow Horse, Jacob Gibson, Sarah Gowing, Nelson Lin, Amelia Navins, and Emahlia Power (2020) 'Anti-Chinese rhetoric tied to racism against Asian Americans stop aapi hate report.' *Asian Pacific Policy & Planning Council*.

- Bundorf, M. Kate, Jill DeMatteis, Grant Miller, Maria Polyakova, Jialu L. Streeter, and Jonathan Wivagg (2021) 'Risk perceptions and protective behaviors: Evidence from COVID-19 pandemic.' Technical Report, National Bureau of Economic Research
- Bursztyn, Leonardo, Alessandra L. González, and David Yanagizawa-Drott (2020) 'Misperceived social norms: Women working outside the home in Saudi Arabia.' *American Economic Review* 110(10), 2997–3029
- Bursztyn, Leonardo, Georgy Egorov, Ingar Haaland, Aakaash Rao, and Christopher Roth (2022) 'Scapegoating during crises.' In 'AEA Papers and Proceedings,' vol. 112 pp. 151–55
- Cantoni, Davide, Felix Hagemeister, and Mark Westcott (2019) 'Persistence and activation of right-wing political ideology.' Technical Report, CRC TRR 190 Rationality and Competition Rationality and Competition Discussion Paper Series 143
- Cheah, Charissa SL, Cixin Wang, Huiguang Ren, Xiaoli Zong, Hyun Su Cho, and Xiaofang Xue (2020) 'COVID-19 racism and mental health in Chinese American families.' *Pediatrics*
- Cheng, Hsiu-Lan, Helen Youngju Kim, Yuying Tsong, Y. Joel Wong et al. (2021) 'Covid-19 anti-asian racism: A tripartite model of collective psychosocial resilience.' *American Psychologist* 76(4), 627
- Christensen, Peter, Ignacio Sarmiento-Barbieri, and Christopher Timmins (2022) 'Housing discrimination and the toxics exposure gap in the united states: Evidence from the rental market.' *Review of Economics and Statistics* 104(4), 807–818
- Cui, Ruomeng, Jun Li, and Dennis J. Zhang (2020) 'Reducing discrimination with reviews in the sharing economy: Evidence from field experiments on Airbnb.' *Management Science* 66(3), 1071–1094
- Durkheim, Émile (1899) 'Antisémitisme et crise sociale.' *Textes II: Religion, morale, anomie* pp. 252–254
- Edelman, Benjamin G., and Michael Luca (2014) 'Digital discrimination: The case of airbnb.com.' *Harvard Business School Working Paper*
- Edelman, Benjamin, Michael Luca, and Dan Svirsky (2017) 'Racial discrimination in the sharing economy: Evidence from a field experiment.' *American Economic Journal: Applied Economics* 9(2), 1–22
- Fisman, Raymond, and Michael Luca (2016) 'Fixing discrimination in online marketplaces.' *Harvard Business Review* 94(12), 88–95

- Fisman, Raymond, Yasushi Hamao, and Yongxiang Wang (2014) 'Nationalism and economic exchange: Evidence from shocks to Sino-Japanese relations.' *The Review of Financial Studies* 27(9), 2626–2660
- Fouka, Vasiliki, and Joachim Voth (2016) 'Collective remembrance and private choice: German-Greek conflict and consumer behavior in times of crisis.' *Stanford Center for International Development Working Paper No 587*(6), 26
- Fradkin, Andrey, Elena Grewal, and David Holtz (2021) 'Reciprocity and unveiling in two-sided reputation systems: Evidence from an experiment on airbnb.' *Marketing Science* 40(6), 1013–1029
- Glaeser, Edward L. (2005) 'The political economy of hatred.' *The Quarterly Journal of Economics* 120(1), 45–86
- Iacus, Stefano M., Gary King, and Giuseppe Porro (2012) 'Causal inference without balance checking: Coarsened exact matching.' *Political analysis* 20(1), 1–24
- Jensen, Robert (2010) 'The (perceived) returns to education and the demand for schooling.' *The Quarterly Journal of Economics* 125(2), 515–548
- Kakar, Venoo, Joel Voelz, Julia Wu, and Julisa Franco (2018) 'The visible host: Does race guide airbnb rental rates in san francisco?' *Journal of Housing Economics* 40, 25–40
- Kumar, Alok, Alexandra Niessen-Ruenzi, and Oliver G. Spalt (2015) 'What's in a name? Mutual fund flows when managers have foreign-sounding names.' *The Review of Financial Studies* 28(8), 2281–2321
- Laouénan, Morgane, and Roland Rathelot (2022) 'Can information reduce ethnic discrimination? Evidence from Airbnb.' *American Economic Journal: Applied Economics* 14(1), 107–32
- Latoya, Hill, and Artiga Samantha (2022) 'COVID-19 Cases and Deaths by Race/Ethnicity: Current Data and Changes Over Time.' Available at <https://www.kff.org/coronavirus-covid-19/issue-brief/covid-19-cases-and-deaths-by-race-ethnicity-current-data-and-changes-over-time/>
- Lindzey, Gardner (1950) 'An experimental examination of the scapegoat theory of prejudice.' *The Journal of Abnormal and Social Psychology* 45(2), 296
- Oreopoulos, Philip (2011) 'Why do skilled immigrants struggle in the labor market? a field experiment with thirteen thousand resumes.' *American Economic Journal: Economic Policy* 3(4), 148–171

- Rothschild, Zachary K., Mark J. Landau, Daniel Sullivan, and Lucas A. Keefer (2012) 'A dual-motive model of scapegoating: Displacing blame to reduce guilt or increase control.' *Journal of Personality and Social Psychology* 102(6), 1148
- Ruiz, Neil G., Khadijah Edwards, and Mark Hugo Lopez (2021) 'One-third of Asian Americans fear threats, physical attacks and most say violence against them is rising.' Available at "Pew Research Center" (Washington, DC, USA)
- Serengil, Sefik Ilkin, and Alper Ozpinar (2020) 'Lightface: A hybrid deep face recognition framework.' In '2020 Innovations in Intelligent Systems and Applications Conference (ASYU)' IEEE pp. 23–27
- (2021) 'Hyperextended lightface: A facial attribute analysis framework.' In '2021 International Conference on Engineering and Emerging Technologies (ICEET)' IEEE pp. 1–4
- Tessler, Hannah, Meera Choi, and Grace Kao (2020) 'The anxiety of being Asian American: Hate crimes and negative biases during the COVID-19 pandemic.' *American Journal of Criminal Justice* 45(4), 636–646
- Tversky, Amos, and Daniel Kahneman (1983) 'Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment.' *Psychological Review* 90(4), 293
- Wills, Thomas A (1981) 'Downward comparison principles in social psychology.' *Psychological Bulletin* 90(2), 245
- Ye, Junting, and Steven Skiena (2019) 'The secret lives of names? Name embeddings from social media.' In 'Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining' pp. 3000–3008
- Ye, Junting, Shuchu Han, Yifan Hu, Baris Coskun, Meizhu Liu, Hong Qin, and Steven Skiena (2017) 'Nationality classification using name embeddings.' In 'Proceedings of the 2017 ACM on Conference on Information and Knowledge Management' pp. 1897–1906
- Zussman, Asaf (2021) 'Scapegoating in evaluation decisions.' *Journal of Economic Behavior & Organization* 186, 152–163