

DISCUSSION PAPER SERIES

DP16624
(v. 5)

Visual Representation and Stereotypes in News Media

Elliott Ash, Ruben Durante, Mariia Grebenshchikova
and Carlo Schwarz

POLITICAL ECONOMY

PUBLIC ECONOMICS

CEPR

Visual Representation and Stereotypes in News Media

Elliott Ash, Ruben Durante, Mariia Grebenshchikova and Carlo Schwarz

Discussion Paper DP16624
First Published 08 October 2021
This Revision 01 May 2022

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Political Economy
- Public Economics

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Elliott Ash, Ruben Durante, Mariia Grebenshchikova and Carlo Schwarz

Visual Representation and Stereotypes in News Media

Abstract

We propose and validate a new method to measure gender and ethnic stereotypes in news reports, using computer vision tools to assess the gender, race and ethnicity of individuals depicted in article images. Applying this approach to 700,000 web articles published in the New York Times and Fox News between 2000 and 2020, we find that males and whites are overrepresented relative to their population share, while women and Hispanics are underrepresented. Relating images to text, we find that news content perpetuates common stereotypes such as associating Blacks and Hispanics with low-skill jobs, crime, and poverty, and Asians with high-skill jobs and science. Analyzing news coverage of specific jobs, we show that racial stereotypes hold even after controlling for the actual share of a group in a given occupation. Finally, we document that group representation in the news is influenced by the gender and ethnic identity of authors and editors.

JEL Classification: L82, J15, J16, Z1, C45

Keywords: Stereotypes, gender, race, media, Computer Vision, Text Analysis

Elliott Ash - ashe@ethz.ch
ETH Zurich

Ruben Durante - ruben.durante@upf.edu
ICREA-UPF, IPEG, BSE, CESifo and CEPR

Mariia Grebenshchikova - maria.grebenshchikova@cdtm.de
Center for Digital Technology and Management, and TU Munich

Carlo Schwarz - carlo.schwarz@unibocconi.it
Bocconi University, IGIER, and PERICLES and CEPR

Acknowledgements

We are grateful to Charles Angelucci, Sascha Becker, Levi Boxell, Julia Cage, Karsten Donnay, Thiemo Fetzer, Nicola Gennaioli, Matt Gentzkow, Ro'ee Levy, Maria Petrova, Kurt Schmidheny, Jesse Shapiro, and Tianyi Wang for their helpful comments. We also thank participants at the Online Seminar in Economics + Data Science, the 2nd Monash-Warwick-Zurich Text-as-Data Workshop, Basel Economics Research Seminar, Bocconi Economics Seminar, UPF Applied Lunch Seminar, ETH Z-PASS Seminar, and UCL Political Science Seminar for helpful discussion. D. Ampudia Vicente, D. Cai, S. Dai, R. Jafarian, and D. Vilalta provided excellent research assistance. We thank the staff at Bekodo for assistance with the development and testing of the image classification software. Ruben Durante acknowledges financial support from the European Union's Horizon 2020 research and innovation program [Grant 759885].

Visual Representation and Stereotypes in News Media

Elliott Ash[†]

Ruben Durante[‡]

Maria Grebenschikova[§]

Carlo Schwarz[¶]

May 1, 2022

Abstract

We propose and validate a new method to measure gender and ethnic stereotypes in news reports, using computer vision tools to assess the gender, race and ethnicity of individuals depicted in article images. Applying this approach to 700,000 web articles published in the New York Times and Fox News between 2000 and 2020, we find that males and whites are overrepresented relative to their population share, while women and Hispanics are underrepresented. Relating images to text, we find that news content perpetuates common stereotypes such as associating Blacks and Hispanics with low-skill jobs, crime, and poverty, and Asians with high-skill jobs and science. Analyzing news coverage of specific jobs, we show that racial stereotypes hold even after controlling for the actual share of a group in a given occupation. Finally, we document that group representation in the news is influenced by the gender and ethnic identity of authors and editors.

Keywords: Stereotypes, Gender, Race, Media, Computer Vision, Text Analysis

*We are grateful to Charles Angelucci, Sascha Becker, Levi Boxell, Julia Cage, Karsten Donnay, Thimo Fetzner, Nicola Gennaioli, Matt Gentzkow, Ro'ee Levy, Maria Petrova, Kurt Schmidheny, Jesse Shapiro, and Tianyi Wang for their helpful comments. We also thank participants at the Online Seminar in Economics + Data Science, the 2nd Monash-Warwick-Zurich Text-as-Data Workshop, Basel Economics Research Seminar, Bocconi Economics Seminar, UPF Applied Lunch Seminar, ETH Z-PESS Seminar, and UCL Political Science Seminar for helpful discussion. D. Ampudia Vicente, D. Cai, S. Dai, R. Jafarian, and D. Vilalta provided excellent research assistance. We thank the staff at Bekodo for assistance with the development and testing of the image classification software. Ruben Durante acknowledges financial support from the European Union's Horizon 2020 research and innovation program [Grant 759885].

[†]ETH Zurich, Warwick CAGE, Monash SoDa Labs, and Bocconi PERICLES; ashe@ethz.ch

[‡]ICREA-UPF, BSE, IPEG, CESifo, IZA, and CEPR; ruben.durante@upf.edu

[§]Center for Digital Technology and Management and TU Munich; maria.grebenschikova@cdtm.de

[¶]Università Bocconi, IGIER, and PERICLES; carlo.schwarz@unibocconi.it.

1 Introduction

Mass media play an essential role in popularizing and perpetuating gender and ethnic stereotypes (e.g., Luther et al., 2017). Given the importance of role models, stereotypes can have tangible effects on cultural attitudes, political processes, and real-world outcomes such as human capital investment, fertility choices, and voting (Jensen and Oster, 2009; Chong and Ferrara, 2009; La Ferrara et al., 2012; Kearney and Levine, 2015). Thus, evidence on the representation of identity groups and the presence of stereotypes in the media could have deep policy implications.

Cultural stereotypes are incorporated in media content in multiple, subtle ways. As a consequence, much of the previous work is qualitative (e.g. Sonnett et al., 2015), and the quantitative analysis of media stereotypes presents formidable measurement challenges. Recent work in social science has attempted to fill this gap using tools from machine learning, in particular through automated text analysis of media content (e.g., Garg et al., 2018; Ribeiro et al., 2018; Rozado and Al-Gharbi, 2021). For example, Rao and Taboada (2021) document sizeable differences in the news topics where male and female subjects are discussed.

Yet text is only one dimension of news content, and more can be learned from the analysis of images in addition to text (Jia et al., 2015). Indeed, images play a central role in newspaper narratives as they make up a large fraction of the content, particularly online (see Figure A.1). Images also have a stronger persuasive effect, as attested by extensive research in political psychology (e.g. Huddy and Gunthorsdottir, 2000). This is in part because of the subtle, implicit messaging that can be communicated through images and through text-image associations (e.g. Sonnett et al., 2015). In addition, while text-based methods are effective in studying gender due to the presence of clearly gendered nouns and pronouns, they are less well-suited for the analysis of racial or ethnic stereotypes since these are frequently not explicitly encoded in text.

To overcome these limitations, we propose a novel approach that investigates stereotypes by analyzing images and their relationship to text. Our approach builds on recently developed computer-vision tools, which allow us to measure the perceived identity characteristics (gender/race/ethnicity) of subjects depicted in article images. Note that, the *perceived* identity

characteristics measures the likely judgment of the reader about a person’s identity, rather than the self-identification of individuals in the images. Our deep-learning-based classifier accurately predicts perceived identity characteristics and is easily scalable and applicable to millions of newspaper images (or other image collections).

Using this methodology, we investigate representation of identity groups and the presence of stereotypes in two major U.S. news outlets: the New York Times and Fox News (NYT and Fox henceforth, respectively). We analyze over two million articles published on the web editions of the two outlets between 2000 and 2020, of which 690k are accompanied by an image. Our analysis looks at both gender and race/ethnicity representation and relates it to the text of the articles.

We report three main results. First, we provide some new descriptive evidence on gender and race representation in news media. Male and White individuals are over-represented relative to their population share, while women and Hispanics are under-represented. Representation varies across news sections, with, for example, Blacks being significantly over-represented in the sports section and Asians in the science section.

Second, we study associations between the text of the articles and the identity of individuals in accompanying images. We find that articles about crime and poverty are disproportionately associated with images of Blacks, articles about immigration to images of Hispanics, and articles about science to images of Asians. The association of Blacks with crime and Hispanics with immigration is more pronounced for Fox than for the NYT. Looking to the intersection of race and gender, we find that the association of Blacks with crime is exclusively driven by images of Black men. Turning to occupations, we find that news article images display racial stereotypes in the sense that groups are represented more in articles about stereotypical jobs, even after controlling for a group’s actual share of employment in that occupation.

Third, we examine how the identity of news producers influences group representation in articles’ images. We find that, conditional on news section and time, the identity of the author of an article is significantly related to the gender and ethnic identities featured in the accompanying images. Similarly, when the share of a section’s editors from a given group increases, so does the share of pictures of members of that group featured in that section’s

articles.

Our findings add to the emerging literature on the use of computer vision tools in social science research (e.g., Jia et al., 2015; Peng, 2018; Boxell, 2021; Adukia et al., 2021). The novelty of our study is that it relates image content to text content to examine both representation and stereotypes, looking at both the gender and race/ethnicity dimensions. Our large-scale machine-learning-powered approach expands on the previous small-scale work using hand-curated datasets (e.g. Gilens, 2009). Given the widespread use of images in corporate advertising, political campaign messaging, and social media networks, an understanding of the stereotypes encoded in images opens up many new avenues for research.

2 Data and Methods

This section provides an overview of the data and the computer vision methodology.

Data Sources. Our main analysis uses a set of online articles collected from the web pages of the NYT and Fox. Overall, the sample includes 2,219,267 articles published between 2000 and 2020, 1,284,939 from NYT and 934,328 from Fox. The data contain the headline, the publication date, a range of metadata (e.g. paper section), the full body text of the article, and the accompanying picture.

As we are interested in the relationship between an article’s image and text, we restrict our analysis to the subset of articles that contain at least one image. In total, 463,886 NYT articles and 226,266 Fox articles are accompanied by an image. We also require articles to have some text and therefore exclude video reports. Lastly, we limit our analysis to articles in the news sections (excluding for example opinion columns, editorials, weather reports, and obituaries). Additional details on the data are provided in Appendix A.

Classifying Identity Characteristics in Images. We use computer vision to automatically detect faces in images and assign perceived identity characteristics to those faces. First, the model detects faces and returns the associated subset of image pixels. This part of the classifier is over 90% accurate (Appendix Table B.1) and translates well to the news article images (Appendix Figure B.2). For each detected face, we use the featurizer from (Schroff

et al., 2015) to form an embedded vector representation of the facial features for input to the classifier.

The classifier takes the face vector as input and learns to assign identity characteristics. The training dataset from Karkkainen and Joo (2021) includes 108,000 images collected from the internet and labeled for perceived gender (male and female) and perceived race/ethnicity (White, Black, Hispanic, Asian, Indian, Middle Eastern). When we say *perceived* in this context, we mean to emphasize that the dataset is labeled by annotators viewing the images, rather than the self-identification of individuals in the images. For the study of stereotypes the perceived identity is the more relevant, as it captures the identity a reader is likely to assign to a person in an image.

Importantly, the dataset from Karkkainen and Joo (2021) is designed to be balanced across gender, race/ethnicity, and age categories.¹ We are further bound by the training dataset in assuming binary gender and categorical race/ethnicity, although this is a simplification of the more fluid and complex reality of these identity characteristics.²

The classifier uses a feed-forward neural net architecture (e.g. Goodfellow et al., 2016), with hyperparameters selected by grid search. The model takes as input the face vector, followed by three hidden layers with 300 neurons each, a dropout layer, batch normalization, and finally a softmax output layer. The final classification system takes in images and outputs a list of detected faces with associated predicted probabilities for gender and race/ethnicity. Appendix B.1 provides additional detail on the training data, model structure, and training process.

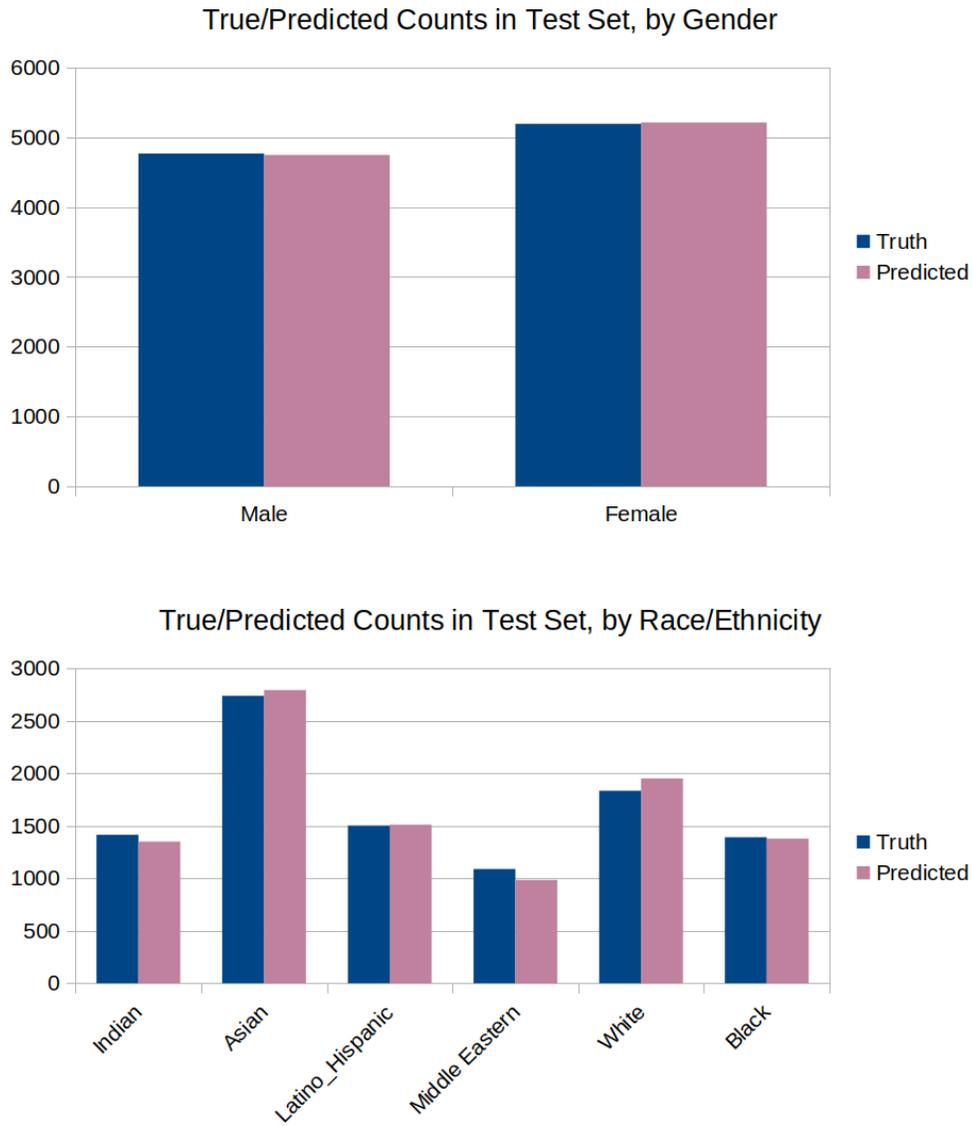
Model performance metrics in the labeled test datasets are reported in Appendix Table B.1. The model is 91% accurate in classifying perceived gender and 73% accurate in classifying perceived race/ethnicity. In a newly annotated in-domain sample of images from our news outlets, the model is 97% accurate for gender and 96% accurate for race/ethnicity.³ Further,

¹Using a large, balanced dataset to train our model is important in light of previous work showing that identity classifiers trained on unbalanced datasets can be systematically biased – for example, gender classifiers may be less accurate for racial-minority faces (Buolamwini and Gebru, 2018).

²Indeed, gender is often understood as a spectrum, and there are many individuals of mixed race or ethnicity. If anything, race/ethnic identity should be thought of as a multidimensional spectrum rather than a set of discrete groups. Besides this limitation of the dataset, we felt that this simplification is acceptable given that identity groups are in many domains still treated as distinct (e.g., in the U.S. census).

³See Appendix Table B.3. This accuracy is conditional on the two annotators agreeing about the gender or race/ethnicity of the face. In our collection, human annotators disagree on gender 6% of the time and

Figure 1: PREDICTED VS TRUE CLASS SHARES IN TEST SET



Notes: These figure show the true image counts (blue) and predicted image counts (pink) by class in the test set, for gender (top panel) and for race/ethnicity (bottom panel).

the model appears to be “fair” in the sense that it is not systematically biased for/against assigning any of the categories. As we show in Figure 1, the classifier faithfully reproduces the distribution of gender and ethnicity in the held-out data. This gives us confidence that the model’s predicted representation shares correctly reflect perceived identity representation in the classified images.

We apply the model to detect and classify faces in all the images in our sample. Appendix Figure A.2 shows the distribution of both the presence and frequency of faces in images. Overall, 66 percent of images contain at least one face. Conditional on the presence of at least one face, an image includes, on average 3.441 faces, with the distribution having a long right tail.

We measure a group’s visual representation in an article as the share of faces of that group detected in the associated image(s). For example, female representation is the share of detected faces classified as female, while Hispanic representation is the share of faces classified as Hispanic.⁴

3 Analysis

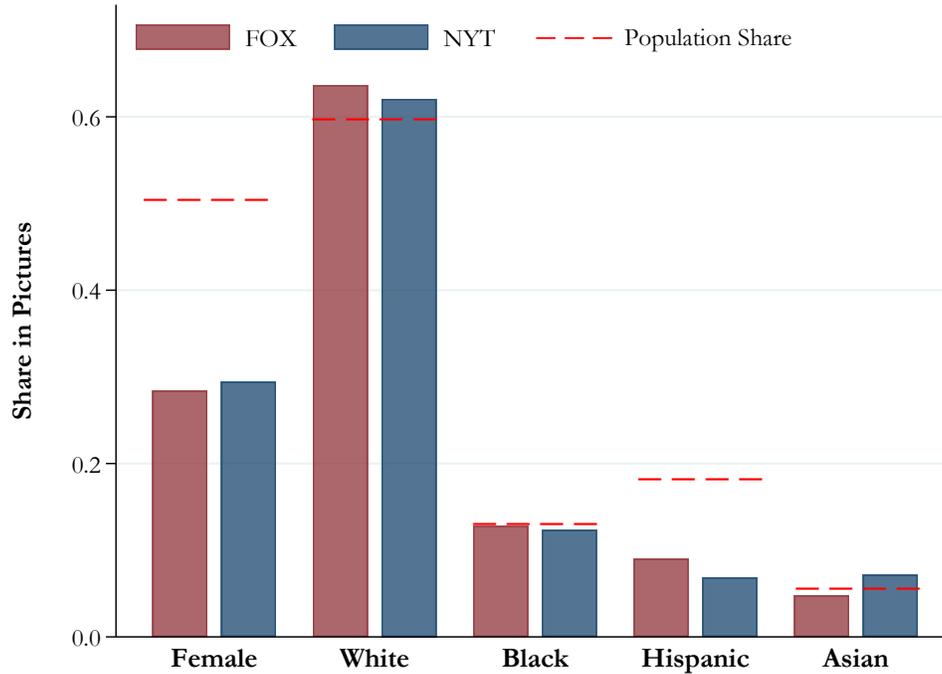
Representation in Newspaper Images. We begin by analysing the representations in newspaper images – i.e., how often people from a given identity group show up in the articles compared to the group’s share in the U.S. population overall.

Figure 2 shows the overall representation by news outlet, with overlaid red lines indicating the group shares in the U.S. population. First, and perhaps most strikingly, women are heavily under-represented in both Fox and NYT articles. While women are slightly more than half of the population, they account for less than a third of the individuals depicted in news articles’ images.

Next, we consider ethnic representation. We find that Whites are over-represented, Hispanics are under-represented, while Blacks and Asians are represented more or less proportionally to their population shares. To the best of our knowledge, this paper is the first to disagree on race/ethnicity 13% of the time.

⁴When calculating the shares we exclude small faces in images (usually in the background), as the classifier performance is worse for these faces. The results remain unchanged if this restriction is not applied.

Figure 2: GENDER AND ETHNICITY SHARES BY NEWS OUTLET



Notes: This figure shows the shares of different gender and ethnic groups in the pictures of Fox and NYT. The red lines indicate the share of each group in the U.S. population.

document the representation of race/ethnic groups in news images. We also observe interesting heterogeneity across the news outlets. Relative to the NYT, Fox includes fewer images of women and more of Whites. Furthermore, Fox shows more images of Hispanics but notably fewer of Asians. In Appendix Figure C.1, we show that these findings are similar if we split the individual ethnic groups by gender. Overall, women are underrepresented for each ethnic group.

These general patterns conceal significant heterogeneity by news section.⁵ Figure 3 shows the heterogeneity in representation across sections relative to the U.S. population shares.⁶ Blue fields indicate those sections in which a group is over-represented, while red fields those

⁵For this analysis, we have harmonized the coding of section names for NYT and Fox. We pool the news outlets to ease exposition. We show the splits by outlet and section in Appendix Figure C.2.

⁶It is of course not clear whether one should compare a group's representation in a news section to its share in the overall population, since, in the real world, a group may be over-represented in some domains and under-represented in others (e.g., Whites in business vs. Whites in sports). Crucially, these real-world differences are themselves also affected by cultural stereotypes, and the goal of our analysis is to examine how image choices may reinforce such stereotypes. In our analysis of occupational stereotypes, presented below, we are able to show that stereotypes affect images choices above-and-beyond job-specific group shares.

in which it is under-represented. We have sorted sections based on the representation of each group.

The first row shows variation in the share of women in images. In line with well-known stereotypes, women have a higher representation in the "soft news" sections - i.e., lifestyle, entertainment, travel - but are under-represented in all but the lifestyle section. The under-representation is particularly stark in the "hard news" sections (politics, world) and male-associated hobbies (sports, cars).

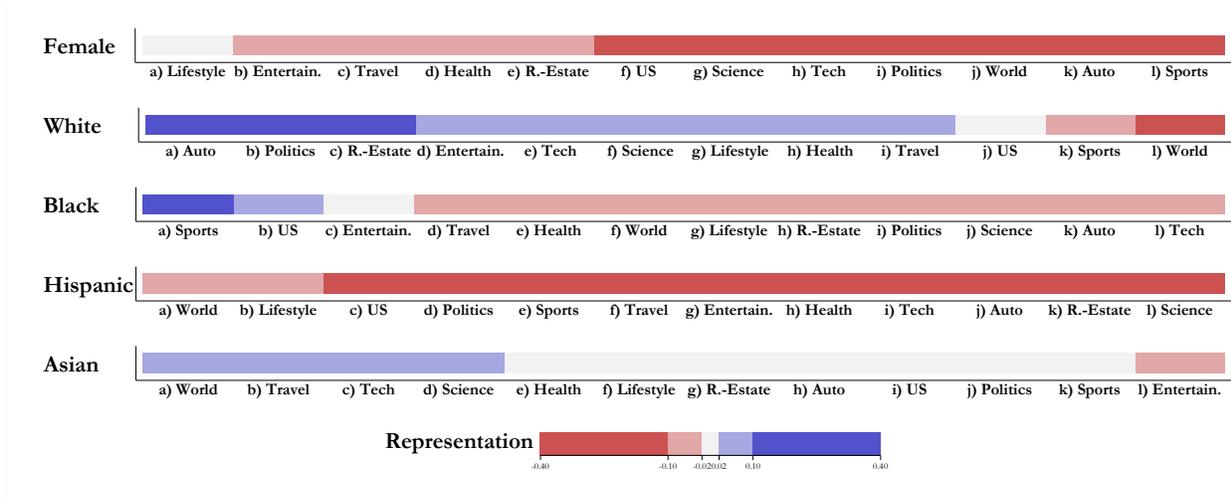
The following rows show the representation for Whites, Blacks, Hispanics, and Asians, respectively. First, we observe that White faces are over-represented in all sections relative to their U.S. population shares. The only exceptions are, unsurprisingly, the sport section and the world section. Whites are particularly over-represented in luxuries (auto, real estate) and politics.

For Blacks, the figure shows a very different ranking of sections. While, as shown in Figure 2 that Blacks are generally proportionally represented, this is particularly driven by their heavy over-representation in the sports section. In contrast, Blacks are under-represented in almost all other sections, especially in luxuries (tech, auto, real estate) and intellectual topics (politics, science, tech).

Hispanics and Asians are both most associated with the world section. This finding, combined with the under-representation of Blacks in this section, supports the view that international news stories tend to concern Latin America and Asia more than Africa. Like Blacks, Hispanics are more heavily under-represented in luxuries (tech, auto, real estate) and science. In contrast, Asians are over-represented in tech and science but under-represented in entertainment.

Association between Text and Images. Next, to explore the relationship between visual and verbal representation, we conduct a combined analysis of images and article text. Specifically, we examine what images tend to accompany news about three divisive policy-relevant issues – crime, poverty, and immigration – which, previous work has shown, tend to be disproportionately associated with disadvantaged minorities (e.g., Gilens, 1996, 2009; Larsen and DeJgaard, 2013; Brown, 2013; Haynes et al., 2016). We also investigate the associ-

Figure 3: GENDER AND ETHNICITY SHARES BY NEWS SECTION



Notes: This figure shows the shares of different gender and ethnic groups in the pictures of Fox and NYT, separately by section. Sections are sorted by frequency. The red lines indicate the share of each group in the U.S. population.

ations between depicted identity groups and science, a high-status activity where stereotypes may be influential.

First, crime is a sensitive topic as Black and Hispanic Americans are highly over-represented in the criminal justice system, partly due to differences in crime rates but also due to prejudicial treatment (Fagan and Ash, 2017; Arnold et al., 2021). Second, poverty and homelessness are the most extreme manifestations of social disadvantage, which affect the Black and Hispanic population disproportionately (Creamer, 2020). Third, immigrants, in particular from Latin America, are often undocumented and targeted by immigration enforcement agencies, and even documented immigrants tend to be economically disadvantaged (Cornelius and Rosenblum, 2005). Fourth, and in contrast, there is a well-known "model minority" stereotype for Asian Americans (Paek and Shah, 2003).

We identify these issues in news articles using a dictionary method (see Appendix C.2 for more detail). Estimating separate linear regressions, we test whether these topics show a stronger association with specific identity groups. The inclusion of fixed effects allows us to compare articles published in the same month, outlet, and, section.

Figure 4 shows some of the most notable associations.⁷ Each bar represents the share

⁷We show the full set of associations in Appendix Table C.2. It is important to note that while the associations between topics and identity characteristics might be driven by differences in real-world group

of faces of a ethnic group in articles about a specific topic, relative to articles in the same section not mentioning that topic. Each row considers a different topic (crime, poverty, and immigration, respectively). Within each row, the left graph includes both genders, the middle graph considers only male faces, while the right graph considers only female faces.

First, take the top row (Panel a), reporting the association between crime-related articles and images of Blacks. In the first image for both genders, we find that both NYT and Fox News are more likely to associate images of Blacks to articles about crime than on other topics. This association is significantly stronger for Fox News (red bar) than for the NYT (blue). This finding is consistent with a stronger racialization of crime in conservative discourse (Gilens, 2009). Analyzing the center and right graphs, we see that there is no association for black females, and that the black-crime relationship is driven exclusively by black male faces.

The association between Black images and articles about poverty is shown in Panel (b) (middle row). Both outlets are more likely to associate images of Blacks to articles about poverty. The relationship is more pronounced for NYT than for Fox, with no difference across genders for NYT. For Fox, instead, male faces are more strongly associated with poverty, with a weaker, barely significant association for female faces. As documented in Appendix Table C.2, further, Fox is more likely than the NYT to display images of Whites when talking about poverty. These two results attest to the different emphasis the two outlets put on “Black” versus “White” poverty, respectively (Gilens, 2000).

Next, Panel (c) (the third row) shows the association of Hispanic with immigration. Both outlets are more likely to use images of Hispanic when talking about immigration, with an association seen for both genders. The immigration-Hispanic association is especially strong for Fox News, consistent with a stronger conservative prioritization of undocumented immigrants from Latin America (Brown et al., 2016). The difference between NYT and Fox is again driven by male faces, consistent with a framing of immigrants as threatening, self-interested economic migrants, rather than harmless, sympathetic asylees (Haynes et al., 2016).

shares, the documented differences between NYT and Fox speak more directly to the different political leanings of the two outlets.

Finally, Panel (d) (bottom row) shows the associations of Asian images with science. Again the associations align with common stereotypes (Paek and Shah, 2003, e.g.). While there do not appear to be any significant differences by outlet, as expected the associations with science is stronger for Asians, especially for men.

Taken together, these findings confirm that news coverage of important policy issues tends to reproduce common racial/ethnic stereotypes, which also vary by gender. Though this generally applies to both outlets, there are significant differences between the two. Those differences are largely consistent with the outlets’ respective ideological stances and political agendas.

Occupational Stereotypes. For the next set of results, we study visual representation for the case of occupations. The advantage of this setting is that we can compare a group’s representation in news images to that group’s true share of workers in each occupation. Through that comparison, we can disentangle to what extent visual representation is explained by occupation-specific demographic differences, as opposed to stereotypes.

For this part of the analysis, we search the text of the articles in our sample and tag all mentions of each of nearly 9000 occupations listed by O*NET (see Appendix C.3 for details). We then estimate the association between the occupations mentioned in an article and the faces in the image associated with it.

Formally, let y_i^k be the share of assigned identity group k (e.g. share of Black faces) in the image from article i , and let the matrix \mathbf{X}_i contain indicator variables for whether each occupation is mentioned in i . The following equation summarizes our strategy:

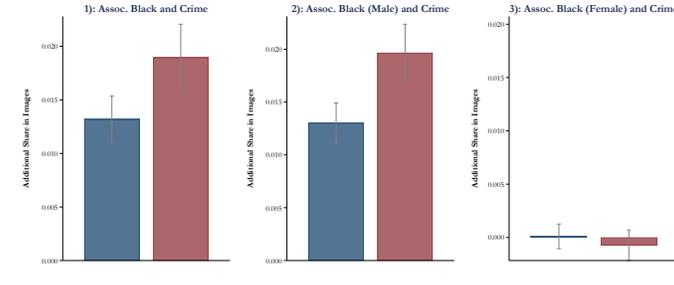
$$y_i^k = \alpha + \mathbf{X}_i' \beta_k + \epsilon_i^k \tag{1}$$

estimated by minimizing mean squared error subject to $\sum_{j=1}^p \beta_{jk}^2 \leq t$, a “ridge penalty” constraint that shrinks large/noisy coefficients and reduces over-fitting.⁸ These regressions allow us to identify which occupations have the highest probability of being associated with a specific group.

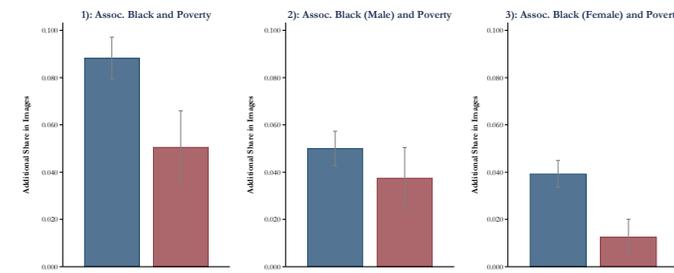
⁸See, e.g., Hastie et al. (2009). The penalty hyperparameter t is selected by ten-fold cross-validation grid search.

Figure 4: ASSOCIATION BETWEEN TOPICS AND VISUAL CHARACTERISTICS

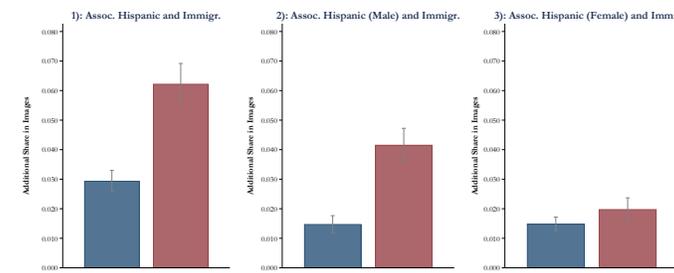
(a) BLACKS AND CRIME



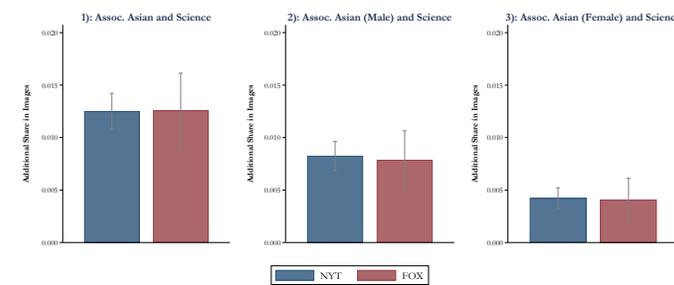
(b) BLACKS AND POVERTY



(c) HISPANICS AND IMMIGRATION



(d) ASIANS AND SCIENCE



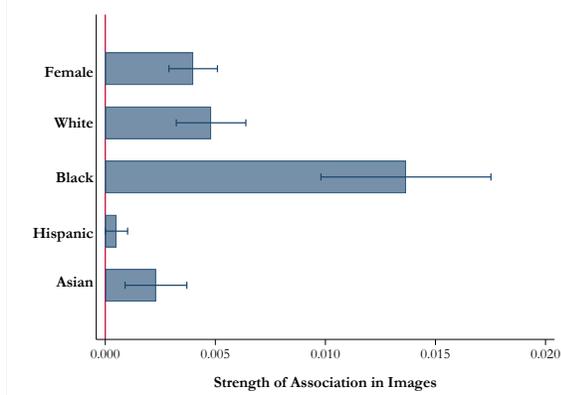
Notes: This figure shows strength of the association of identity groups with articles about specific topics. The bars indicate the additional share of the given identity group that appears in images about the given topic relative to articles in the same outlet, section, and month. The whiskers indicate 95% confidence intervals.

In Appendix Table C.5 Panel A, we validate this method by showing that it derives stereotype-consistent associations between images and nouns, for example by associating “rapper” and “gospel” with images of Blacks, and “swimmer” and “rubles” to Whites. Appendix Table C.5 Panel B reports the occupations most associated with each group. For example, while “hockey player” is associated with images of Whites, “basketball player” is associated with images of Blacks. More generally, we find that images of Whites tend to be associated with high-skill jobs such as intelligence analyst, attorney general, editorial assistant, and safety consultant. In contrast, images of Blacks are associated with lower-skilled service jobs in education and healthcare (i.e., health educator, middle school principal, train operator, health educator, and caster). Images of Hispanics are associated with occupations such as farm worker and garbage collector. Asians, in contrast, are more frequently associated with high-skill occupations including professors, pilot, and engineers.

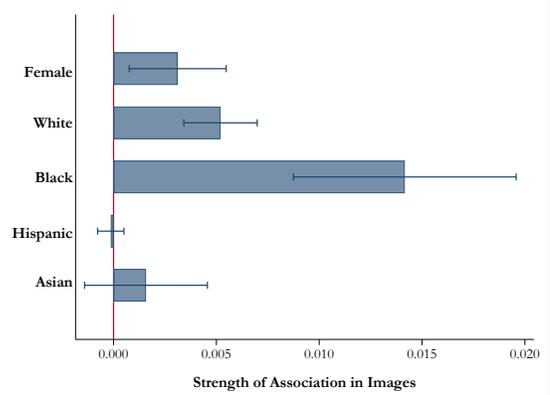
Do these associations reflect real-world occupation shares or visual stereotypes? To more directly answer this question, we match each occupation to its corresponding SOC code, and then link the occupation-identity association coefficients to occupational employment data from the American Community Survey (ACS) (see Appendix C.3 for additional details). From the ACS data, we construct two measures. The first is the group shares $s_{ko} = \Pr(k|o)$, which capture the share of employees in an occupation o that come from identity group k . Second, following Bordalo et al. (2016), we construct a measure of stereotypes based on the “representativeness” of a group for a given occupation. Formally, the representativeness of group k for occupation o is given by $r_{ko} = \Pr(o|k) / \Pr(o| - k)$, where $\Pr(o|k)$ is the probability that a randomly chosen individual of group k works in occupation o , and $\Pr(o| - k)$ is the probability that members of all other groups work in occupation o . This measure, based on relative frequencies, is motivated by the representativeness heuristic (Tversky and Kahneman, 1983) and predicts people’s oversimplified ideas of particular groups (see Bordalo et al., 2016). For example, even though Blacks, due to their size in the overall U.S. population, are a minority in most occupations, they might nevertheless be highly *representative* of some occupations if the relative frequency of Blacks in that occupation is higher than that in the U.S. population overall. Such an occupation, in other words, would be “stereotypically

Figure 5: OCCUPATIONAL STEREOTYPES AND ASSOCIATION IN IMAGES

(a) WITHOUT OCCUPATION SHARES



(b) WITH OCCUPATION SHARES



Notes: This table reports the estimates of Equation (2), in which the dependent variable is the association between an occupation and the share of an identity group in the images as determined by Equation (1). *Represent. Occupation* is the measure of representativeness as suggested by Bordalo et al. (2016). Regressions in Panel (b) additionally include a control variable for the share of an occupation belonging to the indicated identity group (e.g. share of women) based on data from the 2019 American Community Survey. The whiskers indicate 90% confidence intervals.

Black”.⁹ We validate this measure of stereotypes by comparing it to the measures of occupational stereotypes from Garg et al. (2018). As we show in Figure C.3 both measures are strongly positively correlated. The advantage of our measure is that it is available for a far larger set of occupations.

To study the connection between group visual representation and these two measures we estimate the following equation:

$$y_{ko} = \alpha + \beta_1 \cdot r_{ko} + \beta_2 \cdot s_{ko} + \epsilon_{ko} \quad (2)$$

where y_{ko} is the coefficient for group k and occupation o from the Ridge regression (see Equation (1)), which captures the visual representation of the occupation, and r_{ko} and s_{ko} are the representativeness and the occupation shares of group k in occupation o . We estimate the regressions separately by group. The coefficient of interest, β_1 , captures to what extent articles about occupations that are stereotypical of a given group are associated with images of that group, controlling for the group occupation employment share.

⁹Examples of representative occupations by group are shown in Appendix Table C.4 Panel B.

Figure 5 Panel a plots the coefficient β_1 for each of the groups. The estimates are positive for all groups, and statistically significant for all groups except for Hispanics. In Panel b, we additionally control for the occupation shares. The inclusion of this control leaves the estimated coefficients mostly unchanged, just the coefficient for Asians becomes insignificant. These findings suggest that the association between occupations and images is not merely explained by differences in real-world occupation employment shares. On top of true employment variation, visual stereotypes do play a role. Interestingly, the coefficient is the largest for Blacks, suggesting that occupational stereotypes may be strongest for this group.

Effect of Author and Editor Identity on Visual Representation. Finally, we investigate how the identity of newsroom staff affects visual representation in articles. To this end, we coded the gender and ethnicity of authors and editors at the NYT (see Appendix B for details).¹⁰

To analyze the effect of the identity of news producers on visual representation we estimate the following equation:

$$y_{ist}^k = \beta_k \cdot \text{Identity Producer}_{ist}^k + \omega_s + \delta_t + \epsilon_{ist}^k \quad (3)$$

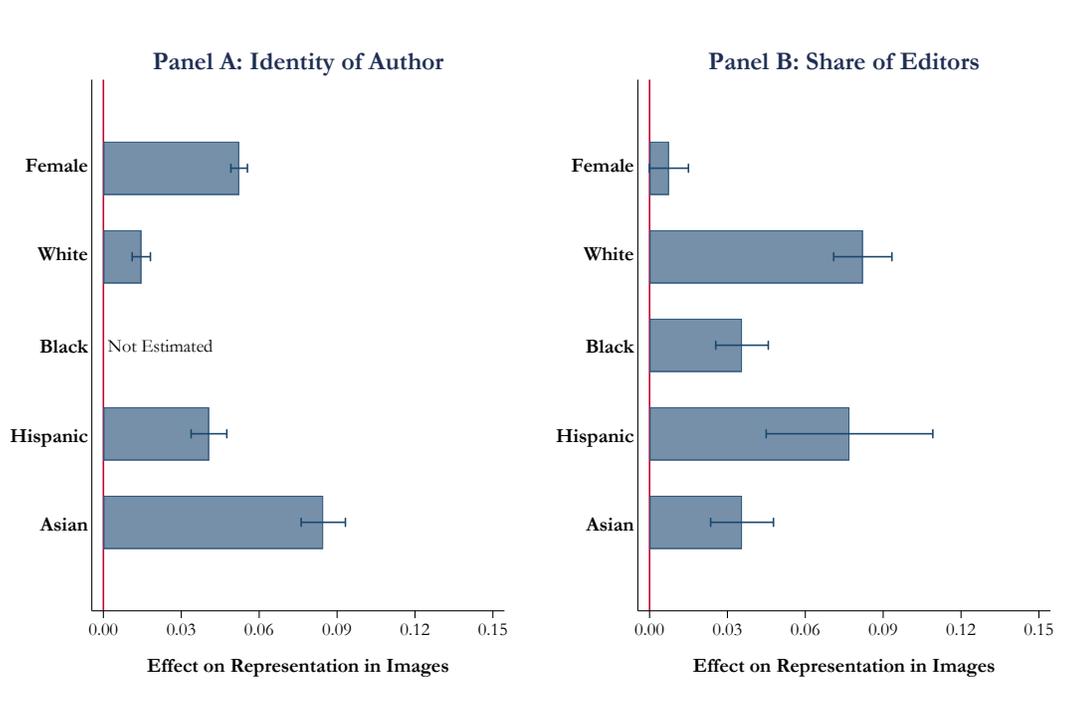
where the dependent variable y_{ist}^k is the image identity share of group k (e.g. females) in article i from section s and month t . $\text{Identity Producer}_{ist}^k$ represents the identity of the relevant news producers, either the author or the editor(s). For the authors, we can relate each article (i) to the identity of its author (k). For editors, instead, we compute the share of editors of each section (s) belonging to a given group (k), e.g., share of female editors of the politics section. All regressions include section and month fixed effects. Standard errors are clustered by section.¹¹

Figure 6 reports the results respectively for authors (panel A) and editors (panel B). In the regressions on author’s identity we always control for the section the author writes in. We find that for all groups, the identity of the author has a positive impact on the visual

¹⁰We are unable to reliably identify Black authors as very few NYT authors have distinctive Black first names. Unfortunately, we could not retrieve information on authors and editors for Fox.

¹¹Using picture captions we rule out the possibility that the effect of the author’s identity is driven by articles that include the picture of the author.

Figure 6: AUTHOR/EDITOR IDENTITY AND REPRESENTATION



Notes: This table reports the estimates of Equation (3) for NYT, in which the dependent variable is the image share of each identity group in an article. The independent variable in Panel A is an indicator variable for the identity group of the article’s author. In Panel B the independent variable is the share of editors in a newspaper section of the respective identity group. All regressions control for section and month fixed effects. Author identity (Panel A) is coded from the name, which does not provide a reliable signal of Black names; hence that estimate is not included. The whiskers indicate 95% confidence intervals.

representation of her/his group.¹² Specifically, articles by female authors are 5.6% more likely to be accompanied by a picture of a woman, and articles by Hispanic and Asian authors are more likely to be accompanied by an image of the respective ethnic group (+4% and +8%, respectively).

We obtain consistent results when estimating the effect of editors' identity. For all groups we find that as the share of a section's editors belonging to a group increases, so does the probability that members of that group appear in pictures associated with that section's articles. The effect is the strongest for White and Hispanic editors for which a 10 percentage point increase in the share of editors is associated with a 0.8 percentage point increase in the share of associated images.

Overall, our findings indicate that the identity of news producers (journalists, editors) has a significant effect on the visual representation of gender and race/ethnic groups. A notable implication is that the choice of images in articles – in general and in their relationship with text – can be attributed, at least in part, to decisions by editors and journalists rather than from automated processes, stock photo databases, or the choices of photographers or image editors. While the latter processes might introduce biases in the image choice, we can show there is also an important element at the stage of news production. Thus, increasing diversity in newsroom staff may help to address issues of under-representation in news media images.

4 Discussion

In this paper, we propose and validate a novel methodology to analyze quantitatively the presence of gender and ethnic stereotypes in news content. Our method – which combines computer vision and natural language processing tools – allows for a comprehensive analysis of the perceived identity characteristics of images, the textual topics of news reports, as well as the interaction between the two. It therefore expands upon previous approaches which considered images or text in isolation.

¹²We note again that our automated method cannot detect Black authors from names. Hence, Black authors are mostly mixed in with White authors, which perhaps explain the smaller coefficient on White author identity.

Using these methods, we have produced a number of substantive findings on representation and stereotypes in two major U.S. news outlets with very different political leanings. Males and Whites are generally over-represented relative to their population share, while females and Hispanics are under-represented. Ethnic minorities are disproportionately associated with crime, poverty, and immigration, a pattern that varies between news outlets consistent with a divergent partisan agenda. News content also reproduces occupational stereotypes. Finally, group representation in media images relates to group representation among newsroom personnel.

Our study illustrates the potential for computer-vision-based methods to tackle questions in social science. Future work could, for example, examine specific events – e.g., the advent of the Black Lives Matter or the #MeToo movements – or the election of charismatic leaders – e.g., Obama or Merkel – to see how those events influence gender and ethnic representation in the media. Moreover, it is important to find out whether and how these media image choices influence attitudes and behavior of news consumers and their peers.

Our approach has broader applicability outside of news media. The image and text analysis methods are flexible and scalable, and they can easily be applied to other types of content, such as social media posts. Further, they could be applied not just to images but to video, for example to analyze visual stereotypes on television.

References

- Adukia, A., A. Eble, E. Harrison, H. B. Runesha, and T. Szasz (2021). What we teach about race and gender: Representation in images and text of children’s books. *Working Paper*.
- Arnold, D., W. Dobbie, and P. Hull (2021). Measuring racial discrimination in algorithms. In *AEA Papers and Proceedings*, Volume 111, pp. 49–54.
- Ash, E. and M. Poyker (2021). Conservative news media and criminal justice: Evidence from exposure to fox news channel.
- Bordalo, P., K. Coffman, N. Gennaioli, and A. Shleifer (2016). Stereotypes. *Quarterly Journal of Economics* 131(4), 1753–1794.
- Boxell, L. (2021). Slanted images: Measuring nonverbal media bias during the 2016 election. *Working Paper*.
- Brown, H. E. (2013). Race, legality, and the social policy consequences of anti-immigration mobilization. *American Sociological Review* 78(2), 290–314.
- Brown, J. A. et al. (2016). The new “southern strategy:” immigration, race, and “welfare dependency” in contemporary us republican political discourse. *Geopolitics, History, and International Relations* 8(2), 22–41.
- Buolamwini, J. and T. Gebru (2018). Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*, pp. 77–91. PMLR.
- Chong, A. and E. L. Ferrara (2009). Television and divorce: Evidence from brazilian novelas. *Journal of the European Economic Association* 7(2-3), 458–468.
- Cornelius, W. A. and M. R. Rosenblum (2005). Immigration and politics. *Annu. Rev. Polit. Sci.* 8, 99–119.
- Creamer, J. (2020). Inequalities persist despite decline in poverty for all major race and hispanic origin groups. *US Census Bureau, September 15*.
- Djourelouva, M. (2020). Media persuasion through slanted language: Evidence from the coverage of immigration.
- Fagan, J. and E. Ash (2017). New policing, new segregation? from ferguson to new york. *Georgetown Law Journal*.
- Galletta, S. and E. Ash (2021). How cable news reshaped local government.
- Garg, N., L. Schiebinger, D. Jurafsky, and J. Zou (2018). Word embeddings quantify 100 years

- of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences* 115(16), E3635–E3644.
- Gilens, M. (1996). Race and poverty in americapublic misperceptions and the american news media. *Public Opinion Quarterly* 60(4), 515–541.
- Gilens, M. (2000). Black poor and the liberal press. *CRJ* 5, 18.
- Gilens, M. (2009). *Why Americans hate welfare: Race, media, and the politics of antipoverty policy*. University of Chicago Press.
- Goodfellow, I., Y. Bengio, and A. Courville (2016). *Deep Learning*. MIT Press. <http://www.deeplearningbook.org>.
- Hastie, T., R. Tibshirani, and J. Friedman (2009). *The Elements of Statistical Learning*. New York, NY: Springer.
- Haynes, C., J. Merolla, and S. K. Ramakrishnan (2016). *Framing immigrants: News coverage, public opinion, and policy*. Russell Sage Foundation.
- Huang, G. B. and E. Learned-Miller (2014). Labeled faces in the wild: Updates and new reporting procedures. *Dept. Comput. Sci., Univ. Massachusetts Amherst, Tech. Rep 14*(003).
- Huddy, L. and A. H. Gunnthorsdottir (2000). The persuasive effects of emotive visual imagery: Superficial manipulation or the product of passionate reason? *Political Psychology* 21(4), 745–778.
- Jensen, R. and E. Oster (2009). The power of tv: Cable television and women’s status in india. *Quarterly Journal of Economics* 124(3), 1057–1094.
- Jia, S., T. Lansdall-Welfare, and N. Cristianini (2015). Measuring gender bias in news images. In *Proceedings of the 24th International Conference on World Wide Web*, pp. 893–898.
- Karkkainen, K. and J. Joo (2021). Fairface: Face attribute dataset for balanced race, gender, and age for bias measurement and mitigation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 1548–1558.
- Kearney, M. S. and P. B. Levine (2015). Media influences on social outcomes: The impact of mtv’s 16 and pregnant on teen childbearing. *American Economic Review* 105(12), 3597–3632.
- La Ferrara, E., A. Chong, and S. Duryea (2012). Soap operas and fertility: Evidence from brazil. *American Economic Journal: Applied Economics* 4(4), 1–31.
- Larsen, C. A. and T. E. Dejgaard (2013). The institutional logic of images of the poor and welfare recipients: A comparative study of british, swedish and danish newspapers. *Journal of European Social Policy* 23(3), 287–299.

- Luther, C. A., C. R. Lepre, and N. Clark (2017). *Diversity in US mass media*. John Wiley & Sons.
- Paek, H. J. and H. Shah (2003). Racial ideology, model minorities, and the” not-so-silent partner:” stereotyping of asian americans in us magazine advertising. *Howard Journal of Communication* 14(4), 225–243.
- Peng, Y. (2018). Same candidates, different faces: Uncovering media bias in visual portrayals of presidential candidates with computer vision. *Journal of Communication* 68(5), 920–941.
- Rao, P. and M. Taboada (2021). Gender bias in the news: A scalable topic modelling and visualization framework. *Frontiers in Artificial Intelligence* 4.
- Ribeiro, F. N., L. Henrique, F. Benevenuto, A. Chakraborty, J. Kulshrestha, M. Babaei, and K. P. Gummadi (2018). Media bias monitor: Quantifying biases of social media news outlets at large-scale. In *Twelfth international AAAI conference on web and social media*.
- Rozado, D. and M. Al-Gharbi (2021). Using word embeddings to probe sentiment associations of politically loaded terms in news and opinion articles from news media outlets. *Journal of Computational Social Science*, 1–22.
- Schroff, F., D. Kalenichenko, and J. Philbin (2015). Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 815–823.
- Sonnett, J., K. A. Johnson, and M. K. Dolan (2015). Priming implicit racism in television news: visual and verbal limitations on diversity. In *Sociological Forum*, Volume 30, pp. 328–347. Wiley Online Library.
- Tversky, A. and D. Kahneman (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological review* 90(4), 293.

Visual Representation and Stereotypes in News Media

Supporting Information

The Appendix presents further details on data, the image classifier, and additional results:

- Appendix A provides details on the data collection.
- Appendix B discusses the computer vision method.
- Appendix C shows additional results.

A Additional Material on Data

New York Times.

We scraped the universe of available articles from the New York Times web site for the years 2000 through 2020 using the NYT API. The API returns the headline, byline, date, and a range of metadata (e.g. paper section). While the API provides the lead paragraph of each article, the full text of the article is not available. We, therefore, visit the links to each article and scrape the full text from the NYT webpage directly.

Importantly for us, the API also provides the links to the images that accompany articles, which we use to obtain the images. The API usually provides links to the image in several sizes and we have always chosen to obtain the largest available image. If an article is accompanied by more than one image, the API will only provide the link to the main image.

To collect the captions of images, we also need to visit the article webpage. Since images on the website get removed from the articles after some time, we are only able to collect the captions for the subset of images that are still available on the website.

We further prepare the NYT data by coding the gender and ethnicity of editors and authors. We collected the history of editorial board membership from the New York Times web site. This history includes the start date, end date, and job title of each editor. A research assistant manually searched for and viewed photos of the editorial board members

to classify gender and ethnicity. We assign editors to sections based on their title, e.g., the sports editor will be assigned to the sports section. For each section and month we then calculate the share of editors of a specific identity group.¹³

Similarly, we code the identity group of authors based on either their first name (female and Black) or their last name (White, Hispanic, Asian).¹⁴ The coding is done based on an extensive list of names that are predominantly used by one specific identity group.¹⁵ We provide some examples of names below. In our coding, we have ensured that no author is assigned to two ethnic groups at the same time (e.g., Black as well as White).

Female first names: *'mary', 'patricia', 'linda', 'barbara', 'elizabeth', 'jennifer', 'maria', 'susan', 'margaret', 'dorothy', 'lisa', 'nancy', 'karen', 'betty', 'helen', 'sandra', 'donna', [...]*

Black first names: *'deion', 'deiondre', 'dele', 'denzel', 'dewayne', 'dikembe', 'duante', 'jamar', 'jevonte', 'kadeem', 'kendis', 'kentay', 'keshawn', 'khalon', 'kofi', 'kwamin', [...]*

White names: *'stoltzfus', 'troyer', 'yoder', 'mast', 'nowak', 'friedman', 'shapiro', 'kowalski', 'stauffer', 'schmitz', 'krueger', 'schwartz', 'kauffman', 'maurer', 'schulte', 'siegel', [...]*

Hispanic names: *'ruvalcaba', 'plascencia', 'bahena', 'chairesz', 'buenrostro', 'valdovinos', 'ceja', 'castrejon', 'mandujano', 'ledezma', 'guadarrama', 'resendiz', 'menjivar', [...]*

Asian names: *'xu', 'zhu', 'zhou', 'xie', 'zhao', 'xiong', 'zhang', 'zheng', 'xiao', 'jiang', 'vue', 'guo', 'luo', 'zeng', 'gao', 'huang', 'ye', 'truong', 'shi', 'phung', 'yang', 'thao', [...]*

Fox News.

Fox News does not provide an API for their articles. We instead scrape the articles from their website by searching for common English words (e.g., “the”, “a”). This allows us to directly obtain the headline, date, section, full text, and associated picture and caption if

¹³We are unable to assign all editors to sections as not all editorial positions indicate a clear section assignment.

¹⁴The coding of authors is based on an automated algorithm, since there are more than 100,000 authors in the NYT data.

¹⁵The lists are obtained from Namecensus.com and Familyeducation.com (both accessed 18.09.2021).

applicable. Similar to the NYT, Fox also archives images after some time. Since we scrape the information directly from the webpage, we can only obtain images that are still available. The Fox News webpage unfortunately does not provide any information on the author.

Harmonizing sections.

For analysis, we additionally harmonize the coding of section between the NYT and Fox News. As Fox News uses the smaller set of sections, which we extract from the article URLs, we recode the sections of the NYT to match those from Fox. This on the one hand involves the simple harmonization of names, e.g., recoding “automobiles” to “auto”. On the other hand we also group sections together, e.g., “arts”, “books”, “movies”, and “theater” are grouped as “entertainment”. We use these recoded sections to exclude opinion columns, weather reports, editorials, and obituaries.

Summary Statistics.

We provide summary statistics for our data by outlet in Table A.1. The first subtable provides information on the number of articles for which we have text, images, images with faces, or a tagged occupation. The second subtable provides summary statistics for the main variables we use in our analysis.

Table A.1: SUMMARY STATISTICS

Nr. of Articles in Sample

	NYT	FOX
Nr. of articles	1,284,939	934,328
Nr. of articles with text	1,281,137	913,837
Nr. of articles with image	463,886	226,266
Nr. of articles with faces	304,573	149,139

Notes: This table presents the number of articles for each newspaper in several categories.

Summary Statistics Main Variables

	NYT					FOX				
	Mean	SD	Min	Max	Obs	Mean	SD	Min	Max	Obs
Image Shares										
Share Female	0.29	0.39	0.00	1.00	281,627	0.28	0.40	0.00	1.00	136,993
Share White	0.62	0.43	0.00	1.00	281,627	0.64	0.44	0.00	1.00	136,993
Share Black	0.12	0.29	0.00	1.00	281,627	0.13	0.31	0.00	1.00	136,993
Share Hispanic	0.07	0.21	0.00	1.00	281,627	0.09	0.25	0.00	1.00	136,993
Share Asian	0.07	0.23	0.00	1.00	281,627	0.05	0.19	0.00	1.00	136,993
Text Shares										
Share Female	0.24	0.30	0.00	1.00	1,132,247	0.23	0.33	0.00	1.00	780,751
Share White	0.72	0.42	0.00	1.00	499,583	0.68	0.44	0.00	1.00	299,727
Share Hispanic	0.12	0.30	0.00	1.00	499,583	0.20	0.38	0.00	1.00	299,727
Share Asian	0.15	0.34	0.00	1.00	499,583	0.11	0.30	0.00	1.00	299,727
Topics										
Crime	0.22	0.41	0.00	1.00	1,284,939	0.20	0.40	0.00	1.00	934,328
Poverty	0.02	0.13	0.00	1.00	1,284,939	0.01	0.09	0.00	1.00	934,328
Immigration	0.04	0.20	0.00	1.00	1,284,939	0.03	0.17	0.00	1.00	934,328
Science	0.27	0.45	0.00	1.00	1,284,939	0.11	0.31	0.00	1.00	934,328
Authors										
Share Female	0.29	0.45	0.00	1.00	1,284,939					
Share White	0.24	0.43	0.00	1.00	1,284,939					
Share Hispanic	0.01	0.12	0.00	1.00	1,284,939					
Share Asian	0.02	0.14	0.00	1.00	1,284,939					
Editors										
Share Female	0.30	0.26	0.00	1.00	1,021,602					
Share White	0.93	0.18	0.00	1.00	1,021,602					
Share Black	0.02	0.10	0.00	1.00	1,021,602					
Share Hispanic	0.00	0.02	0.00	0.25	1,021,602					
Share Asian	0.04	0.14	0.00	1.00	1,021,602					

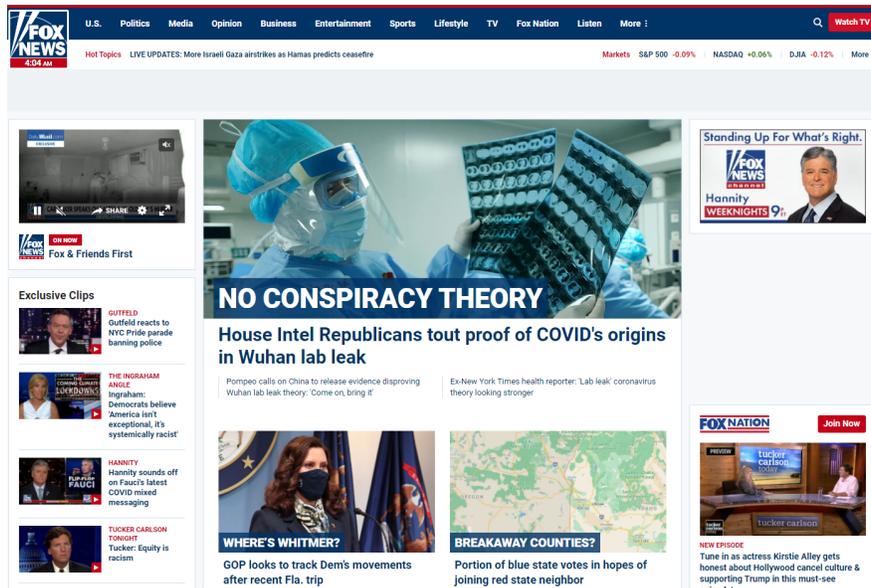
Notes: This table presents descriptive statistics for the NYT and FOX data.

Figure A.1: SCREENSHOTS OF NYT AND FOX NEWS WEBSITES

(a) NYT



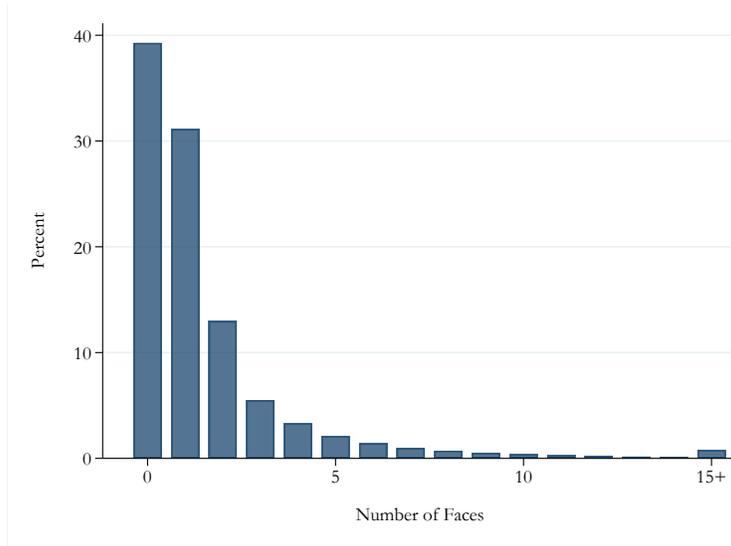
(b) Fox



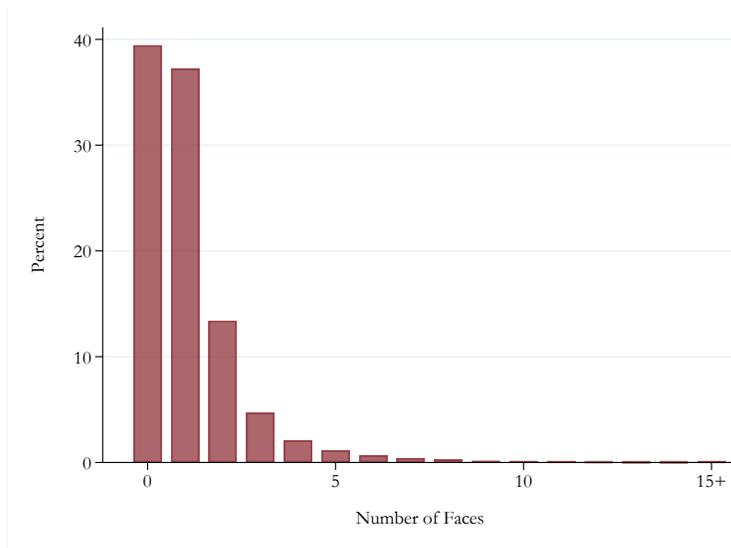
Notes: The figure shows screenshots of the start pages of Fox and NYT taken on the 20th May 2021.

Figure A.2: NUMBER OF FACES IN IMAGES

(a) NYT



(b) FOX



Notes: The figure shows a histogram of the number of faces identified by the neural net in the images of NYT and Fox news. For this figure we winsorized the number of faces at 15 which is above the 99th percentile for both outlets

B Additional Material on Methods

B.1 Image Classifier

Our image classifier has two components: 1) face detection and 2) face group classification.

First, the algorithm detects faces in images, using code available at github.com/ageitgey/face_recognition. The model is trained to detect faces using the Labeled Faces in the Wild (LFW) benchmark dataset (Huang and Learned-Miller, 2014), available at vis-www.cs.umass.edu/lfw. This dataset consists of 13,000 images of faces collected from the internet. The detector gets state-of-the-art performance (99.38% accuracy) on the LFW’s held-out test set.

Each detected face is extracted and represented as a 512-dimensional embedding using FaceNet (Schroff et al., 2015), a pre-trained facial feature extractor. To train the classifier, we start with the FairFace dataset (Karkkainen and Joo, 2021), a collection of 108,501 images with faces designed to be balanced according to gender, ethnicity, and age. These data are available at github.com/joojs/fairface (we used the images with padding=0.25). The FaceNet detector obtains an embedding for each face, and then those embeddings are inputted as predictors in a neural net classifier to produce the gender and ethnicity labels.

We have two classifiers – one to predict gender and one to predict ethnicity. The architecture is a feedforward neural net (or multilayer perceptron). The inputs are the face embeddings, the output is a softmax layer (multinomial logit) across classes. In between there is a sequence of fully connected hidden layers, separated by rectified linear unit activation. We searched over the following architecture choices to discover the best model based on performance in the test set: number of hidden layers (2, 3, or 4), number of neurons in each layer (200, 300, or 400), different levels of dropout at each layer (0.1, 0.3, 0.5, 0.7), and with/without batch normalization at each layer.

The best model on the held-out test set has three hidden layers, 300 neurons in each hidden layer, dropout = 0.5, and with batch normalization. The model’s predicted class probabilities are then calibrated to match the test set.

Table B.1: IMAGE CLASSIFIER PERFORMANCE IN TRAINING DATASET

	Precision	Recall
Face Detection (Aggregate)	0.909	
Gender (Aggregate)	0.912	
Male	0.914	0.918
Female	0.910	0.906
Ethnicity (Aggregate)	0.731	
White	0.701	0.746
Black	0.852	0.844
Hispanic	0.490	0.493
Asian	0.891	0.909
Indian	0.701	0.669
Middle Eastern	0.576	0.521

Notes: This table reports precision and recall for the face detection as well as the gender and ethnicity prediction task for the image classifier.

Table B.2: CONFUSION MATRICES FOR PERFORMANCE IN TRAINING DATA

A. Gender

<i>True Class</i>	<i>Predicted Class</i>		Total True
	Male	Female	
Male	4319	448	4767
Female	428	4765	5193
Total Pred.	4747	5213	
Pred. / True	0.996	1.004	

Notes: This table reports the confusion matrix for the gender prediction task.

B. Race and Ethnicity

<i>True Class</i>	<i>Predicted Class</i>						Total True
	Indian	Asian	Hispanic	Mid East	White	Black	
Indian	946	71	180	102	30	84	1413
Asian	50	2486	113	4	46	37	2736
Hispanic	169	142	740	151	246	53	1501
Middle Eastern	91	6	170	566	242	12	1087
White	22	60	217	149	1367	18	1833
Black	71	26	90	11	19	1173	1390
Total Pred.	1349	2791	1510	983	1950	1377	
Pred. / True	0.95	1.02	1.01	0.90	1.06	0.99	

Notes: This table reports the confusion matrix for the ethnicity prediction task.

Performance metrics for the image classifier in the training domain (the held-out test set from FairFace) are reported in Table B.1. In bold, we report the aggregate accuracy for face detection, gender detection, and ethnicity detection.

We also report the precision and recall by class for the gender and ethnicity categories.

The confusion matrices in Table B.2 show additional detail on the model performance by class.

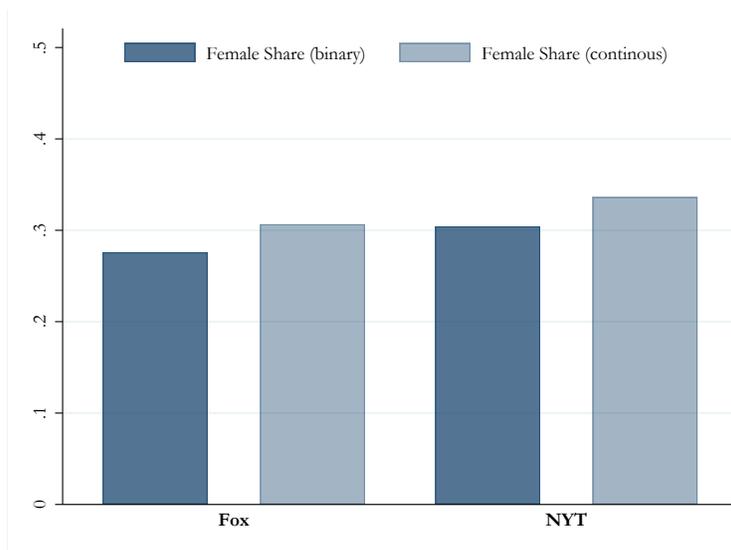
It is reassuring that the model obtains good precision and recall across the classes we are interested in (Appendix Table B.1). The confusion matrices in Table B.2 show in more detail where the model tends to make errors. Model performance is especially good for Black and Asian faces. Meanwhile, Hispanic faces are the most difficult to classify, with the model often classifying them into a different category.

However, the errors for Hispanic ethnicity are "balanced" in the sense that false positives and false negatives happen at about the same rate. Other ethnicities are also wrongfully classified as Hispanic, and the predicted share of Hispanic faces in the test set is very close to (within 1% of) the true share of Hispanic faces. This representative balance can be seen even more clearly, for the Hispanic category as well as the others, in Figure 1. Reassuringly, the classifier faithfully reproduces the distribution of classes in the held-out data. Hence, we have confidence that predicted representation shares in our newspaper image collection fairly reflect identity representation shares and are not systematically biased for or against a particular class.

Notwithstanding some misclassifications, our classifier clearly performs well in measuring identity characteristics in images. In any case, we would not expect a perfect classifier. Neither gender nor ethnicity are perfectly binary categories. Gender is often understood as a spectrum, and there are many individuals of mixed race or ethnicity. It is not unreasonable to think of ethnic identity as a multidimensional spectrum rather than a set of discrete groups. Among these groups, in our data, Hispanic is the least clearly defined category and hence also less visually distinctive.

As a robustness test, we investigate how the female share in images would change if we used a continuous measure based on the probabilities which our classifier estimates. This is akin to treating gender as a spectrum between 0 (100% male) and 1 (100% female). The results in Figure B.1 show that the decisions to treat the categories as binary does not strongly impact our findings, as overall female share is very similar to our baseline metric.

Figure B.1: BINARY VS. CONTINUOUS FEMALE SHARE



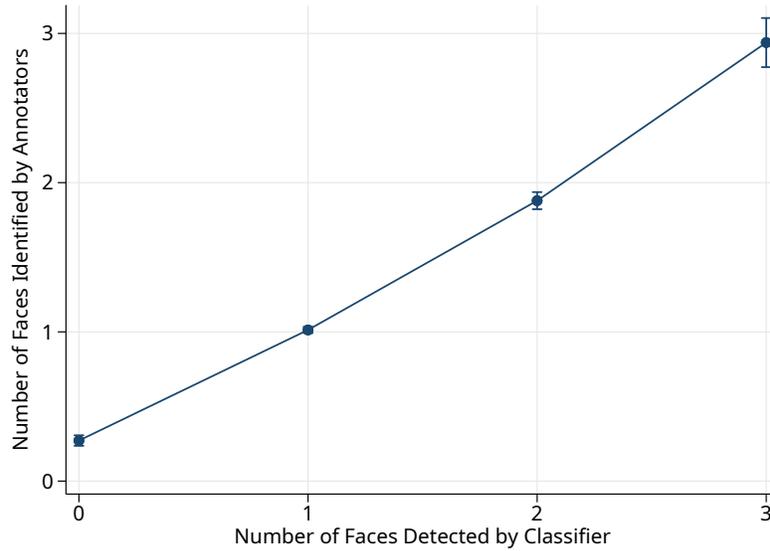
Notes: The figure shows the female share in images for the NYT and Fox if we consider the gender to be a binary or continuous variable.

Because the FairFace dataset includes many images from news websites, we have confidence that the test-set performance metrics are applicable in our empirical domain of the Fox and NYT websites. To double-check this expectation, we performed a small-scale human annotation of images from our collection to check that the classifier is accurate for our purposes.

For this task, we hired four data entry specialists from an online freelancing platform. The annotators had significant data entry experience and worked on an hourly rate. We randomly sampled 500 images from Fox and 500 images from NYT. Each annotator considered 500 images, with 2 annotators assigned to each image. Hence we obtained 2 annotations each for 1000 images – 500 from Fox and 500 from NYT. Each annotator first counted the number of faces. The annotators agreed on the number of faces 83.9% of the time, with a mean absolute deviation of 0.44.

We then assess the classifier’s performance in detecting faces, using the set of images where the annotators agree on the number of faces. Figure B.2 plots the number of faces detected by the classifier, compared to the number of faces identified by annotators. Overall, the classifier reproduces annotator judgment on number of faces 73% of the time, or 80% when there are fewer than three faces. Beyond exact match, the counts are highly correlated, with

Figure B.2: NUMBER OF FACES IDENTIFIED BY ANNOTATORS AND THE CLASSIFIER



Notes: This figure plots the average number of faces identified by human annotators (Y axis) against the number of faces detected by the computer vision algorithm (X axis). Sample limited to images where the annotators agree on the number of faces. Error bars give standard errors of the mean.

an overall Spearman correlation of 0.74 between the classifier face count and the annotator face count.

To get at the accuracy of the classifier, we subset to the set of images with one face – that is, where the classifier and both annotators all agree there is a single face in the image. The annotators were asked to assign a single gender and single race/ethnicity to each face. For race/ethnicity, we exclude the Indian and Middle Eastern categories because they are not part of our empirical analysis and because they are infrequent.

Table B.3 reports the main metrics from annotating gender and race/ethnicity in the newspaper article images. First, we computed agreement as the proportion of faces where the two annotators agreed on the class assignment. As shown in Column 1, the annotators agreed on gender 93.8% and they agreed on race/ethnicity 87.3% of the time. These agreement rates illustrate the significant subjectivity in the assignment of these features.

Next, we measured accuracy after subsetting to the images where the two annotators agreed on the identity assignment. Conditional on the annotators agreeing on a face’s gender, the classifier’s accuracy is 97.9% for gender. Correspondingly, for the images where the

Table B.3: CLASSIFIER ACCURACY IN ANNOTATED NEWSPAPER IMAGES

	(1) Agreement	(2) Accuracy (all)	(3) Accuracy (NYT)	(4) Accuracy (Fox)
Gender (s.e.)	0.938 (0.0105)	0.979 (0.0061)	0.958 (0.0129)	1 (.)
N	518	486	240	246
Race/Ethnicity (s.e.)	0.873 (0.015)	0.96 (0.009)	0.97 (0.012)	0.95 (0.016)
N	463	404	206	198

annotators agreed on one of the four main race/ethnicity classes (white, black, Hispanic, Asian), the classifier accuracy is 96.0%.

Since a key part of our empirical analysis is the comparison between NYT and Fox News, we next check whether the classifier is accurate for the images from both outlets. Encouragingly, Columns 3 and 4 show that they model gets similar performance in both outlets.

B.2 Text Analysis

Parallel with the image analysis, we separately assign identity characteristics to articles using text analysis. For each article, we have the text from the headline and article body. We use a dictionary-based method to detect identity characteristics.

The starting point is the list of distinctive identity words from Garg et al. (2018), listed below. We count distinctive words indicating gender identity and distinctive names indicating ethnic identity.¹⁶ For each article, we then measure the share of gender words that are female – that is, the number of female words, divided by the summed counts of male and female words. We show examples of articles with a high female/male count in Table B.4. Similarly, we measure the share of distinctively ethnic names in each class – for example, the number

¹⁶Ethnic names are used (e.g. “Chang”) rather than explicit ethnic labels (e.g. “Asian”), because the latter are often used to refer to non-ethnic entities, such as geographical regions.

of distinctively Hispanic names, divided by the summed counts of distinctively ethnic names across all ethnic categories.

The following lists of words were detected using regular expression pattern matches. The documents were made all lowercase, and split into words based on white space.

Lists of lexicon words.

Gendered words.

Female words: *'her', 'she', 'mrs', 'ms', 'wife', 'mother', 'women', 'woman', 'sister', 'female', 'herself', 'daughter', 'girl', 'girls', 'daughters', 'queen', 'sisters', 'lady', 'females', 'ladies', 'aunt', 'queens', 'ma', 'wives', 'hers', 'maternity', 'niece', 'mothers', 'princess', 'maid', 'mas', 'maternal', 'chick', 'lesbian', 'maiden', 'dame', 'mom', 'fiancee', 'sis', 'step-mother', 'maids', 'bride', 'dames', 'chicks', 'feminine', 'gal', 'mama', 'aunts', 'lass'*

Male words: *'his', 'he', 'him', 'mr', 'himself', 'mrs', 'man', 'men', 'husband', 'king', 'father', 'son', 'brother', 'brothers', 'pa', 'male', 'sons', 'bros', 'fellow', 'sir', 'boy', 'boys', 'knight', 'guy', 'prince', 'duke', 'mens', 'males', 'uncle', 'grandfather', 'kings', 'nephew', 'boyfriend', 'guys', 'paternity', 'knights', 'dukes', 'fellows', 'stepfather', 'son-in-law', 'fathers', 'grooming', 'manned', 'chap', 'husbands', 'bro', 'grandfathered', 'grooms', 'buddy'*

Garg et al. (2018) Ethnic names.

White/Black names: *'harris', 'nelson', 'robinson', 'thompson', 'moore', 'wright', 'anderson', 'clark', 'jackson', 'taylor', 'scott', 'davis', 'allen', 'adams', 'lewis', 'williams', 'jones', 'wilson', 'martin', 'johnson'*

Hispanic names: *'ruiz', 'alvarez', 'vargas', 'castillo', 'gomez', 'soto', 'gonzalez', 'sanchez', 'rivera', 'mendoza', 'martinez', 'torres', 'rodriguez', 'perez', 'lopez', 'medina', 'diaz', 'garcia', 'castro', 'cruz'*

Asian names: 'cho', 'wong', 'tang', 'huang', 'chu', 'chung', 'ng', 'wu', 'liu', 'chen', 'lin', 'yang', 'kim', 'chang', 'shah', 'wang', 'li', 'khan', 'singh', 'hong'

Table B.4: ARTICLES WITH HIGH FEMALE/MALE COUNT

Newspaper	Article Title	# Female	# Male
Panel A: High Female Count			
NYT	The Real Queen of Wellness	120	0
NYT	She Felt Fine, but Her M.R.I. Showed Several Strokes. What Was Wrong?	111	0
NYT	A Bit of Relief: 'Soup Is Soup'	110	0
NYT	Despite Risks, Freed Reporter Loved Iraq	110	0
NYT	Freed Reporter Returns to U.S. for Joyful Reunion	110	0
NYT	She Had Pain in Her Knee but No Obvious Injury. Why?	109	0
NYT	The Neediest Cases; Retired, but Now Struggling To Support an Older Sister	108	0
NYT	Tavi and the Taviettes	107	0
NYT	A Place Fit for the Boss	107	0
NYT	Living a Life Interrupted by Bipolar Disorder	106	0
Fox	USWNT star Mallory Pugh discusses blazing a trail by going pro [...]	100	0
Fox	Player ratings: How did the USWNT do in the Olympics group stage?	97	0
Fox	Woman who sold handmade bracelets to fund double lung transplant [...]	97	0
Fox	Central American moms, kids stuck in legal purgatory after deportation raids	84	0
Fox	[...] China woman activist more than ever dedicated to feminism [...]	82	0
Fox	Traveling prom dress sisterhood honors friend lost to cancer	78	0
Fox	Notre Dame's All-America point guard, Hall of Fame coach [...]	76	0
Fox	Wozniacki, Jankovic: Outlooks differ, results same	76	0
Fox	Old foes Williams, Henin ready to renew rivalry	76	0
Fox	Veronica Lake, former '40s Hollywood star, appeared 'very damaged' [...]	72	0
Panel A: High Male Count			
NYT	Will Favre Come Back? Should He?	0	303
NYT	Imbroglia Over the Elf-Kohl Payoff: Was It a Hoax to Hide Kickbacks?	0	262
NYT	'We Followed the President's Orders'	0	260
NYT	Day 2 of Senior Bowl, Offensive Notes	0	259
NYT	Life in a Financial Neverland	0	254
NYT	Jim Surdoval: Politician, Dealmaker, Convicted Felon	0	249
NYT	What Happened to the Fortune Michael Jackson Made?	0	247
NYT	The Hardest (and Most Important) Job in Afghanistan	0	240
NYT	Russert to Romney: Are You A Flip-Flopper?	0	238
NYT	The Price of Running Late, Part 1	0	238
Fox	NBA's greatest players of all-time: Who are the top 23?	0	191
Fox	MadFriars' End of the Year Announcer Series - Fort Wayne	0	168
Fox	19 important takeaways from the NFL Combine	0	165
Fox	2012 NFL Combine quarterback rankings	0	159
Fox	The best MLB Draft pick ever for all 30 teams	0	155
Fox	Jensen rates the North American goalies	0	154
Fox	Tiger Tracker: Thursday	0	151
Fox	100 burning questions about 2020 NFL season with kickoff just months away	0	150
Fox	NBA mock draft: It's hard to go wrong in this year's talent grab	0	145
Fox	Mock around the clock: There's no time to draft like this year	0	145

Notes: This table reports example articles with a high/low share of female words for NYT and Fox. See Appendix B.2 for the list of words.

B.3 Validation of Image Representation

Comparing Image and Text Representation. We have constructed measures of identity representation in news articles based on both the images and the text. Thus, we can investigate how the images relate to the text. In particular, we aim to understand whether a text-based approach suffices to identify representation in particular articles.

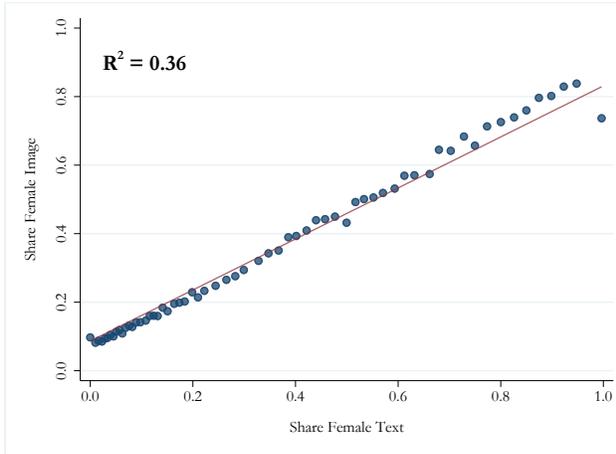
Figure B.3 shows binned scatterplots for the identity shares measured from images (vertical axes) against the identity shares measured from text (horizontal axes). First, as we can see in Panel (a), the image and text measurements for gender are highly correlated ($R^2 = 0.38$). This correlation reflects that detecting gender in English-language text is straightforward, thanks to the presence of gendered pronouns. Thus, text-based measures likely works well in analyzing gender stereotypes in the text of news media. Obviously, the text-based measure does not allow to analyze stereotypes that are embedded in the images.

In contrast, Figure B.3 Panels (b, c, d) show the analogous binned scatterplots for ethnicity – White/Black,¹⁷ Hispanic, and Asian, respectively. Relative to gender, the relationship between the text and image measures of ethnicity is far weaker, with R^2 of 0.007 (White/Black), 0.018 (Hispanic), and 0.01 (Asian), respectively. This weak relationship reflects that name mentions are a relatively poor predictor for race and ethnicity. Our computer vision approach improves on this by measuring ethnicity directly in the images, with the added advantage that we can differentiate between Whites and Blacks.

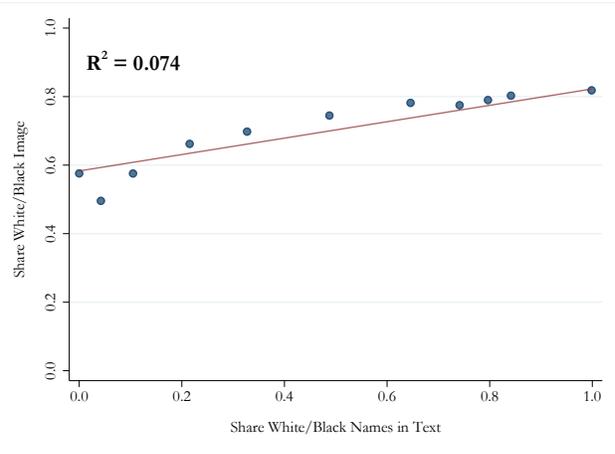
¹⁷We merge White and Black for this validation because the list of distinctive White and Black names from Garg et al. (2018) are identical. Hence the text rankings are identical for those groups. We show the binscatters separately in Appendix Figure B.4.

Figure B.3: IDENTITY CHARACTERISTICS IN IMAGES VERSUS IN TEXT

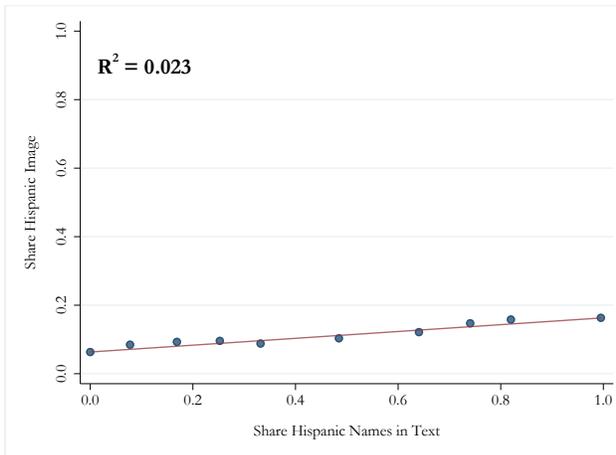
(a) FEMALE



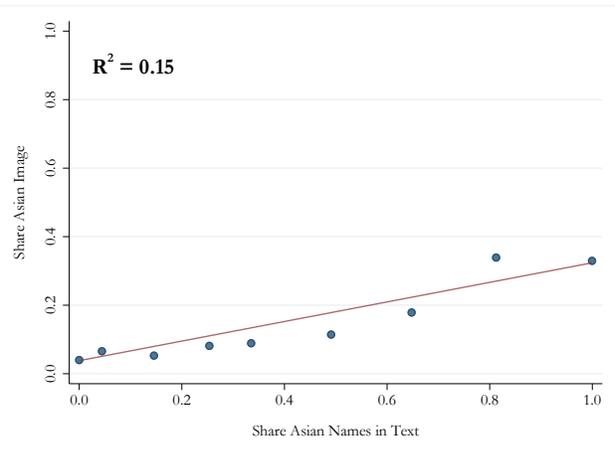
(b) WHITE / BLACK



(c) HISPANIC

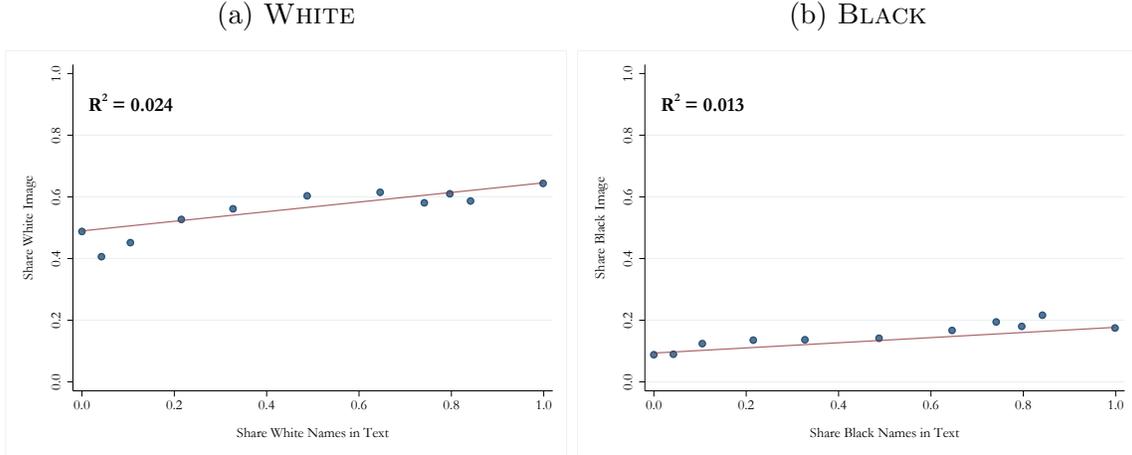


(d) ASIAN



Notes: Binned scatterplots showing the relationship between the image share of the indicated identity groups (vertical axes) and the text-based measures (horizontal axes), based on Garg et al. (2018). The red line indicates the best linear fit. The R^2 of linear regression of the two variables is included in the top left corner of each panel.

Figure B.4: IMAGE VS TEXT SHARES, WHITE AND BLACK SEPARATELY



Notes: These binned scatter plots show the relationship between the share of the indicated identity groups and the text-based measures following Garg et al. (2018). In this figure we consider the share of Whites and Blacks separately. The red line indicates the linear best fit. We further indicate the R^2 of linear regression of the two variables in the top left corner.

Ethnicity Classification and Associated Countries Mentioned in Articles. As an additional validation, we relate the ethnicity from article images with the countries mentioned in the article texts. Formally, let y_i^k be the share of identity group k (e.g. share of faces of Black race) in the image from article i , and let the matrix \mathbf{X}_i contain the number of times a country is mentioned in i .¹⁸ We then estimate linear regressions of the form

$$y_i^k = \alpha + \mathbf{X}_i^t \beta_k + \epsilon_i^k \tag{B.1}$$

subject to $\sum_{j=1}^p \beta_{jk}^2 \leq t$, a “ridge penalty” constraint that shrinks large/noisy coefficients and reduces over-fitting.¹⁹ For each identity group k , we rank each country j by the associated

¹⁸For this analysis we focus on countries with more than 10 million inhabitants. We further remove “Chad” as it is also used as a male firstname. The full list of countries is China, India, United States, Indonesia, Pakistan, Brazil, Nigeria, Bangladesh, Russia, Mexico, Japan, Ethiopia, Philippines, Egypt, Vietnam, DR Congo, Turkey, Iran, Germany, Thailand, United Kingdom, France, Italy, Tanzania, South Africa, Myanmar, Kenya, South Korea, Colombia, Spain, Uganda, Argentina, Algeria, Sudan, Ukraine, Iraq, Afghanistan, Poland, Canada, Morocco, Saudi Arabia, Uzbekistan, Peru, Angola, Malaysia, Mozambique, Ghana, Yemen, Nepal, Venezuela, Madagascar, Cameroon, Côte d’Ivoire, North Korea, Australia, Niger, Taiwan, Sri Lanka, Burkina Faso, Mali, Romania, Malawi, Chile, Kazakhstan, Zambia, Guatemala, Ecuador, Syria, Netherlands, Senegal, Cambodia, Somalia, Zimbabwe, Guinea, Rwanda, Benin, Burundi, Tunisia, Bolivia, Belgium, Haiti, Cuba, South Sudan, Dominican Republic, Czech Republic, Greece, Jordan, Portugal, Azerbaijan, Sweden.

¹⁹See, e.g., Hastie et al. (2009). The penalty hyperparameter t is selected by ten-fold cross-validation grid search.

regression coefficient β_{jk} . We can then investigate which country has the strongest association with each identity group share.

These rankings are reported in Table B.5 Panel A. The association of ethnicities with countries provides further support for the performance of the image classifier, as the top-ranked countries are populated mainly by people of the respective ethnicity. The White ethnicity is associated with European countries, while Black ethnicity has countries from Africa and the Caribbean. The Asian countries are indeed in Asia and Hispanic countries in Latin America. For the Hispanic column, however, there are some deviations, if we move beyond the Top 5 (unreported). While the classifier correctly learns Latin American countries, there are also some Arab countries associated with the Hispanic category. These slight errors reflect what we discussed above in regard to the classifier performance for ethnicity – the classifier has more problems classifying Hispanics in images since Hispanics are a less clearly defined and less visually distinctive ethnic category.

For comparison, we construct similar rankings based on the text measure (see Appendix Table B.5 Panel B). While the text-based rankings do well for Hispanics and Asians, they cannot distinguish that White countries are European and Black countries are African. Further, the text-based measures miss many articles that cover ethnic groups without mentioning them or any of the names from Garg et al. (2018).

Table B.5: COUNTRIES ASSOCIATED WITH ETHNIC GROUPS

Panel A: Image-based Measures

White	Black	Hispanic	Asian
Russia	Zimbabwe	Guatemala	Philippines
Sweden	Nigeria	Venezuela	Thailand
Poland	Benin	Bolivia	Cambodia
Germany	Burundi	Mexico	Japan
Canada	Senegal	Peru	Indonesia

Notes: This table reports the 5 countries with the strongest association with each identity group, indicated in the header of each column. Associations are determined by estimating the ridge regressions from Equation (1).

Panel B: Text-based Measures

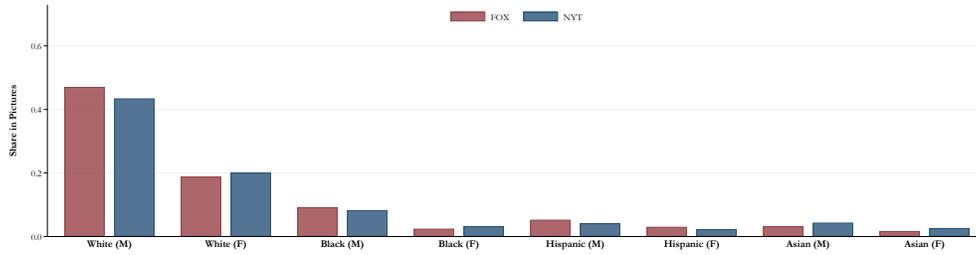
White	Hispanic	Asian
Niger	Cuba	China
Jordan	Guatemala	Taiwan
Benin	Peru	Malaysia
Vietnam	Mexico	Pakistan
Australia	Bolivia	Vietnam

Notes: This table reports the 5 countries which have the strongest association with the identity group indicated at the top as measured based on the names from Garg et al. (2018). The association are determined by estimating a ridge regression based on Equation (1).

C Additional Results

C.1 Additional Representation Results

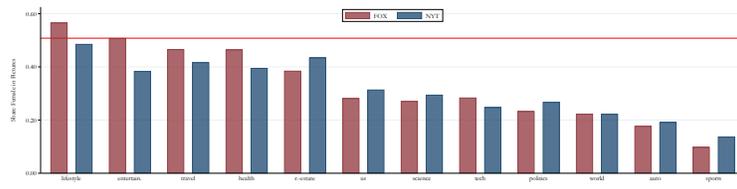
Figure C.1: ETHNICITY SHARES BY GENDER AND NEWS OUTLET



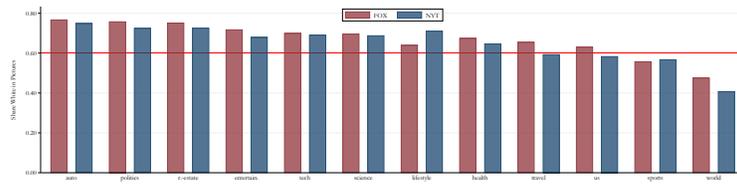
Notes: This figure shows the shares of different ethnic groups split by gender in the pictures of Fox and NYT.

Figure C.2: GENDER AND ETHNICITY SHARES BY NEWS SECTION AND OUTLET

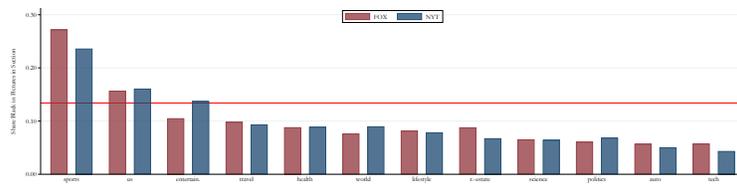
(a) FEMALE



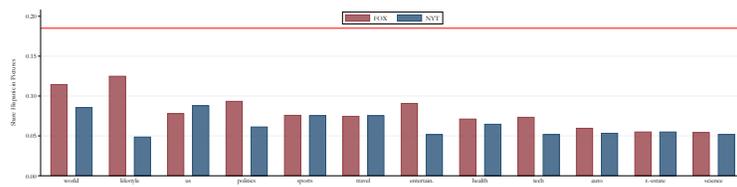
(b) WHITE



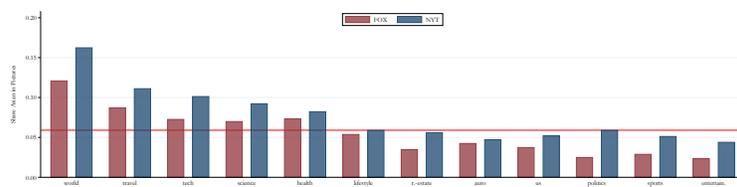
(c) BLACK



(d) HISPANIC



(e) ASIAN



Notes: This figure shows the shares of different gender and ethnic groups in the pictures of Fox and NYT. The shares are shown separately by section and outlet. The sections were harmonized across outlets by the authors. The red lines indicate the share of each group in the U.S. population.

C.2 Associations in Policy-Relevant Topics

In this section, we provide additional evidence for stereotypical narratives with regards to the topics of include crime, immigration, poverty, and science.

As discussed in the main text, we identify articles related to these topics based on the set of the following dictionary words:

Words Crime: *'crime', 'arson', 'assault', 'burglary', 'drugs', 'prison', 'fraud', 'murder', 'homicide', 'theft', 'thief', 'thieves', 'vandal', 'criminal'*

Words Immigration: *'immigra'*

Words Poverty: *'poverty'*

Words Science: *'science', 'scient', 'research', 'professor', 'biologist', 'veterinarian', 'physician', 'physicist', 'engineer', 'chemist', 'mathematician', 'astronomer', 'biochemist'*

We show several example articles with a high topics count in Table C.1. Overall, the dictionary based method appears to perform very well when it comes to identifying relevant articles.

We use the resulting topic words counts in a regression framework. Let W_i^l be an indicator which is 1 if a words related to topic l appears in an article i . We then use the topic measures in a regression model at the article-level of the form:

$$y_{inst}^k = \beta_1^{kl} \cdot W_i^l + \beta_2^{kl} \cdot Fox \times W_i^l + \gamma_{ns} + \delta_t + \epsilon_{inst}^k \quad (C.1)$$

where y_{inst}^k is the image share for identity group k in article i from news outlet n , section s , year t . The specification includes a media outlet times section fixed effect γ_{ns} , and month fixed effect δ_t . We run separate regressions for each topic and each identity group. The main coefficient of interest is β_1^{kl} which measures the strength of the association between topic l and identity group k . Additionally β_2^{kl} measures whether the strength of the association systematically differs for Fox News, relative to NYT. Note that these regression are equivalent to the figures we present in Figure 4.

Table C.1: ARTICLES WITH HIGH TOPICS COUNTS

Newspaper	Article Title	Word Count
Panel A: Crime		
NYT	Could an Ex-Convict Become an Attorney? I Intended to Find Out	227
NYT	What Should Be the Purpose of Prison?	157
NYT	Senate Judiciary Committee Confirmation Hearing	152
NYT	Is Prison Necessary? Ruth Wilson Gilmore Might Change Your Mind	151
NYT	A Surge in Shootings	134
FOX	Patchy reporting undercuts national hate crimes count	57
FOX	Violent crime in US cities surges in summer amid pandemic, protests: [...]	55
FOX	Is the Caribbean more dangerous than you realize?	52
FOX	Prisons of Honduras Ruled by Inmates and Corruption	51
FOX	Death row inmates' last words	50
Panel B: Poverty		
NYT	What Does Obama Really Believe In?	86
NYT	The Poverty Platform	51
NYT	A Gloomy Prediction on How Much Poverty Could Rise	46
NYT	How to Define Poverty? Let Us Count the Ways	39
NYT	U.N. Aims to Cut Poverty in Half as Experts Wonder How to Measure It	36
FOX	Four in five in US face near-poverty, no work	32
FOX	US poverty on track to rise to highest since 1960s	31
FOX	US poverty at new high: 16 percent, or 49.1M	27
FOX	As Economy Rebounds, About 46.5 Million Americans Remain Stuck In Poverty	25
FOX	Latino Poverty Rate Climbs to 28%	25
Panel C: Immigration		
NYT	Why an Heiress Spent Her Fortune Trying to Keep Immigrants Out	140
NYT	The Immigration Equation	123
NYT	The Democrats Have an Immigration Problem	117
NYT	Bankrolling the Anti-Immigration Movement	98
NYT	Legal Help for Immigrants, Even Illegal Ones	77
FOX	Immigration: Key Players On the Road to Reform	67
FOX	Steve King Asks What Part of the 'Rule of Law' Don't We Understand?	52
FOX	Struggling with A Population Decline, Baltimore Pins Its Hopes On Immigrants	48
FOX	Deportation Loophole Lets Thousands Live and Work in US Indefinitely	46
FOX	For Latinos, Immigration More Personal than Political, Poll Says	46
Panel D: Science		
NYT	And Science for All	207
NYT	Why Are There Still So Few Women in Science?	155
NYT	Do Scientists Compete Unethically?	149
NYT	What Shortage of Scientists and Engineers?	146
NYT	Billionaires With Big Ideas Are Privatizing American Science	141
FOX	Declining numbers of blacks seen in math, science	54
FOX	Why (Some) Scientists Avoid the Public	53
FOX	Congressional debate over science funding draws fire from critics	45
FOX	10 Scientists Who Mattered in 2012	44
FOX	China's human-like monkeys spark concerns	40

Notes: This table reports example articles with a high number of topic related words for NYT and Fox. See Appendix B.2 for the list of words.

The regression results on specific topics are reported in Table C.2. We provide evidence of stereotypical representations based on images and policy topics. First, women are under-represented in crime and over-represented in immigration, poverty, and science. Fox has statistically weaker female associations with immigration, poverty, and science. In column (2) we show that the results are similar if we control for the share of female words in the text. This suggests that the images encode associations beyond what would be expected based on the text.

Turning to the ethnicity groups (columns 3-6), we start with the crime topic. Reflecting common stereotypes, Whites and Asians are seen less often in articles about crime, while Blacks and Hispanics are seen more often. The negative relation between White and crime is stronger at Fox, similar to the race-crime messaging differences observed by Ash and Poyker (2021).

For immigration, Whites and Blacks are shown less often in newspaper images, while Hispanics and Asians are depicted more frequently. The association for Hispanics is stronger for Fox, potentially reflecting a policy opposing an increase in immigration from Latin American countries. In contrast, the immigration correlations for Whites and Asians are weaker for Fox.

Next, poverty is less likely to be represented with images of White people. In contrast, we observe a positive association with poverty for all minority groups: Blacks, Hispanics, Asians. Interestingly, Fox is more likely to show images of White poverty and less likely to use images of Black poverty when compared to NYT.

Fourth, when discussing science, we see that NYT shows fewer White, Black, and Hispanic individuals. Instead reflecting a common stereotype, as images of Asians are used. In contrast, Fox associates science articles with images of Whites and is less likely to show images of minority groups.

This more directed analysis of policy-relevant topics has further fleshed out the use of gender and ethnic stereotypes in news media. While the overall variation across topics could be due to differences in real-world based rates, the difference between NYT and Fox is not.²⁰

²⁰While we could in principle control for base rates of the identity groups with regards to crime, immigration, poverty, and science, it is not clear what base rate to use in each case. For example, in the case of immigration, should we use the shares of identity groups with regards to overall, or just illegal immigration?

Fox and NYT have made different editorial choices in how gender and ethnicity are presented along with these socially relevant topics. Exposure to those associations might influence voter attitudes, behavior, and downstream policy decisions (Djourelova, 2020; Galletta and Ash, 2021).

Even within illegal immigration, should we focus on the southern border or people overstaying their visas? This makes the base rate problem somewhat complicated and subjective. Further, the base rates themselves may result from stereotypes and discrimination. Thus, we save this type of comparison for future work.

Table C.2: REPRESENTATION IN ARTICLES ON CRIME, IMMIGRATION, POVERTY, SCIENCE

	Share Female (1)	Share Female (2)	Share White (3)	Share Black (4)	Share Hispanic (5)	Share Asian (6)
Panel A: Crime						
I[# Words>0]	-0.034*** (0.002)	-0.017*** (0.001)	-0.018*** (0.002)	0.020*** (0.001)	0.002** (0.001)	-0.011*** (0.001)
FOX × I[# Words>0]	-0.006* (0.003)	-0.012*** (0.003)	-0.026*** (0.003)	0.011*** (0.002)	0.004** (0.002)	0.000 (0.002)
Female Share Text		0.714*** (0.002)				
Panel B: Immigration						
I[# Words>0]	0.017*** (0.003)	0.012*** (0.003)	-0.046*** (0.003)	-0.005** (0.002)	0.032*** (0.002)	0.005*** (0.002)
FOX × I[# Words>0]	-0.014** (0.005)	-0.008* (0.005)	-0.026*** (0.006)	-0.004 (0.003)	0.039*** (0.004)	-0.004 (0.003)
Female Share Text		0.715*** (0.002)				
Panel C: Poverty						
I[# Words>0]	0.058*** (0.005)	0.020*** (0.004)	-0.125*** (0.006)	0.091*** (0.005)	0.026*** (0.003)	0.008** (0.003)
FOX × I[# Words>0]	-0.026** (0.011)	-0.007 (0.010)	-0.011 (0.013)	-0.040*** (0.009)	0.013 (0.008)	0.015** (0.007)
Female Share Text		0.715*** (0.002)				
Panel D: Science						
I[# Words>0]	0.018*** (0.002)	0.006*** (0.001)	-0.007*** (0.002)	-0.008*** (0.001)	-0.003*** (0.001)	0.019*** (0.001)
FOX × I[# Words>0]	-0.013*** (0.004)	0.001 (0.003)	0.012*** (0.004)	-0.009*** (0.003)	-0.002 (0.002)	-0.004* (0.002)
Female Share Text		0.715*** (0.002)				
Outlet × Section FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Outlet FE	Yes	Yes	Yes	Yes	Yes	Yes
Section FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	406,105	391,119	406,105	406,105	406,105	406,105
Mean of DV	0.29	0.29	0.63	0.12	0.08	0.06

Notes: Regression estimates from Equation (C.1). The outcome variable is the image share of each identity group in an article (indicated by column header). The explanatory variable is an indicator which is 1 if a topic related words appears in an article. Fox is an indicator variable for articles published in the respective outlet. All regressions control for outlet times section, and month fixed effects. Robust standard errors are used in all regressions.

C.3 Words and Text Associations

This sections provides additional details on the analysis of the association between occupations and images. As discussed in the main text, we tag occupations based on a list from O*Net. Table C.3 shows examples of the reported job titles for the occupation of chief executive (SOC code 11-1011). We clean the reported job titles by removing terms in brackets (e.g. “(CEO)”).

We can match occupations to the the ACS data based on the SOC code. For some occupations the ACS does not provide data at the full 6 digit SOC level (e.g., “31-201”). In these cases, we merge at the highest possible level. To avoid many-to-many merges, we exclude occupations that fall into several SOC categories (e.g., “administrative assistant”). The results are similar if we instead merge these occupation to all possible SOC codes.

Table C.3: EXAMPLE OCCUPATION DATA

SOC Code	Title	Reported Job Title
11-1011	Chief Executives	Chief Diversity Officer (CDO)
11-1011	Chief Executives	Chief Executive Officer (CEO)
11-1011	Chief Executives	Chief Financial Officer (CFO)
11-1011	Chief Executives	Chief Nursing Officer
11-1011	Chief Executives	Chief Operating Officer (COO)
11-1011	Chief Executives	Executive Director
11-1011	Chief Executives	Executive Vice President (EVP)
11-1011	Chief Executives	Operations Vice President

Notes: This table reports examples from the occupation data from O*Net.

Based on the ACS data, we calculate the shares of identity groups in the labor force as well as our measure of representativeness. Table C.4 reports examples for occupation with high employment shares (Panel A) or high representativeness (Panel B).

Table C.4: EXAMPLES OF OCCUPATION WITH HIGH SHARES AND REPRESENTATIVENESS

Panel A: High Employment Shares:				
Female	White	Black	Hispanic	Asian
Department secretary	Elevator service technician	Master barber	Grader	County engineer
Staff assistant	Elevator technician	Barber	Artisan plasterer	Avionics engineer
Childcare worker	Elevator mechanic	Stylist	Applicator	Biomedical technician
Child caregiver	Power lineman	Site worker	Plasterer	City engineer
Child care worker	Lineworker	Asbestos worker	Roofer	Research leader
Nanny	Fire management officer	Master cook	Commercial roofer	Environmental engineer
Child care provider	Crew boss	Gas attendant	Housekeeping	Chemical engineer
Barber stylist	Shift commander	Technical inspector	Room attendant	Traffic operations engineer
Cosmetologist	Battalion fire chief	Smog technician	Housekeeping aide	Systems integration engineer
Hairstylist	Fire battalion chief	Parking attendant	Skill labor	Agricultural engineer
Panel B: High Representativeness:				
Female	White	Black	Hispanic	Asian
Department secretary	Racehorse trainer	Mail handler	Grader	Park naturalist
Staff assistant	Dog trainer	Mail processor	Artisan plasterer	Associate professor of genetics
Childcare worker	Horse trainer	Stylist	Applicator	Regional forester
Child caregiver	Grain farmer	Barber	Plasterer	Biologist
Child care worker	Ranch manager	Master barber	Roofer	Food technologist
Nanny	Horticulturist	Certified ophthalmic technician	Commercial roofer	Environmental epidemiologist
Child care provider	Rancher	Veterinary nurse	Housekeeping	Molecular biology professor
Barber stylist	Dairy farmer	Customer service assistant	Room attendant	Epidemiologist
Cosmetologist	Nursery manager	Toll operator	Housekeeping aide	Natural resource specialist
Hairstylist	Greenhouse manager	Director of social work	Skill labor	State epidemiologist

Notes: This table reports examples of occupation with high employment shares for specific identity groups as well as occupations with a high representativeness.

As discussed in the main text, we then relate these measures to the associations between occupations and images which we derive using a ridge regression (see Equation (C.1)). In Table C.5, we list the occupations with the strongest association for each of our 5 identity groups.

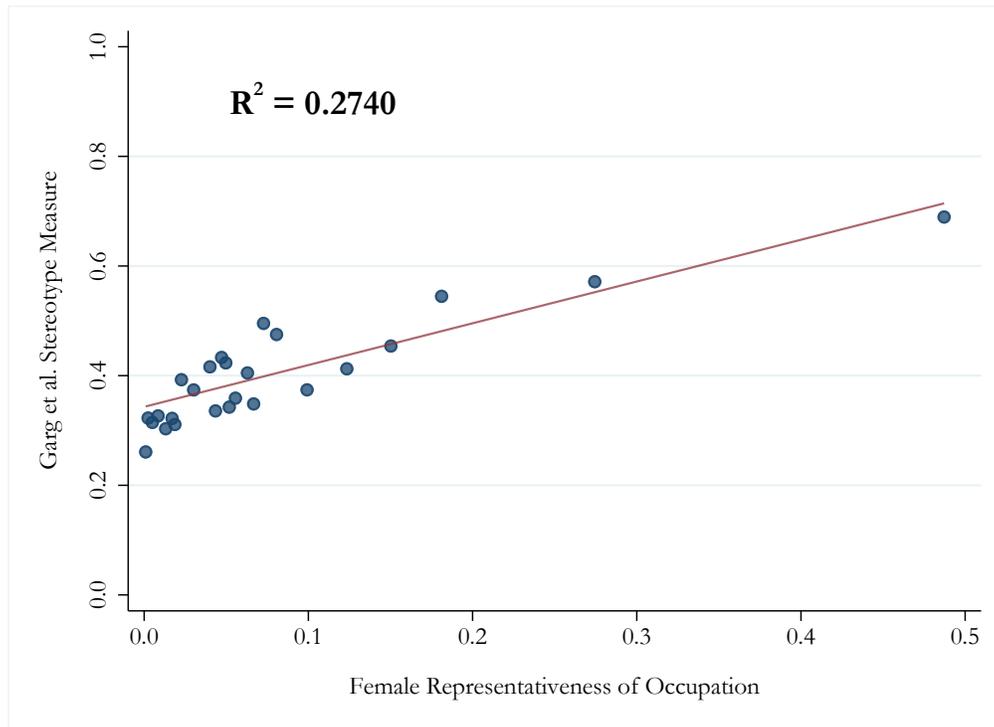
Table C.5: NOUNS AND OCCUPATIONS THAT ARE PREDICTIVE OF IMAGE CHOICE

Panel A: Nouns				
Female	White	Black	Hispanic	Asian
Women	Annexation	Rapper	Morales	Hye
Actress	Hockey	Apartheid	6ix9ine	Mainland
Husband	Degrom	Schoolgirls	Ruiz	Yen
Congresswoman	Chancellor	Blackness	Foxnewslatino	Renminbi
Woman	Nadal	Blacks	Guatemala	Subversion
Chairwoman	Rubles	Gospel	Carey	Microblog
Daughter	Trump	Hop	Detainer	Ocha
Ms	Euro	Continent	Congresswoman	Bak
Boyfriend	Separatists	Basketball	Polizzi	Censors
Heroin	Swimmer	Receiver	Sanchez	Yuan

Panel B: Occupations				
Female	White	Black	Hispanic	Asian
Actress	Hockey player	Middle school principal	Watch leader	Plant operator
Makeup artist	Deputy attorney general	Health educator	Garbage collector	Hydrologist
Facialist	Locker room attendant	Caseworker	Spanish interpreter	Interpreter
Hairstylist	Racehorse trainer	Basketball player	Stock clerk	Office clerk
Case manager	Safety consultant	Train operator	Elevator mechanic	Food science professor
Ballerina	Dairy farmer	Caster	Farm worker	Assembler
Dermatologist	Flight instructor	Case worker	Paper conservator	Materials engineer
Customer service agent	Intelligence analyst	Gospel singer	Service person	International coordinator
Cosmetologist	Editorial assistant	Receiver	Patrol officer	Chief pilot
Midwife	Ride operator	Brick mason	Cardiac sonographer	Tour operator

Notes: This table reports the 10 words which have the strongest association with the identity group indicated in the column header. Associations are determined by estimating a ridge regression based on Equation (1). Panel A restricts the predictive words to the set of nouns. Panel B includes words based on a list of occupations from O*Net.

Figure C.3: VALIDATION OF REPRESENTATIVE MEASURE



Notes: This binned scatter plot show the relationship between our measure of stereotypes (representativeness) vs a measure of stereotypes based on the judgment of people on MTurk. The MTurk measure is based on the data from Garg et al. (2018). We matched the occupations from Garg et al. (2018) to our occupation list based on a string match. We normalize both measures between 0 and 1. The red line indicates the linear best fit. We further indicate the R^2 of linear regression of the two variables in the top left corner.