

Demographic Correlates of Humanizing Media Coverage of Homicide: Evidence from the *Boston Globe*, 1976-84

Abstract

Media coverage of violent crime both reflects and shapes how society perceives members of different demographic groups and subgroups. Measuring how coverage varies across these groups through an intersectional lens can thus provide nuanced insight into public perceptions of different groups' humanity. We merge FBI data on Massachusetts homicides from 1976-1984 with *Boston Globe* articles to study how race, age, and sex of homicide victims and offenders—as well as the interactions of these characteristics—correlate with humanizing media coverage, conditional on detailed circumstantial factors. We uncover patterns of differential coverage across groups: (i) Among male victims, whites are significantly more likely to receive humanizing coverage than blacks; (ii) there is no significant difference in coverage across races for female victims, except those aged 18-29, among whom blacks receive less humanizing coverage; (iii) female and juvenile victims receive more humanizing coverage; and (iv) race/sex of the offender do not correlate with coverage.

Keywords— Media distortion analysis | Natural language processing | Intersectional analysis | Ideal victim theory

17 Crime stories make up 10-50% of all newspaper articles (Chermak 1995), with dispropor-
18 tionate attention devoted to stories about homicide in particular despite its relatively low prevalence
19 (Paulsen 2003). Among homicides, the quality and quantity of coverage varies substantially, and
20 this variation in coverage may plausibly relate systematically to the characteristics of the individuals
21 involved in nuanced ways that have important social implications (Chermak 1995; Meyers 1996;
22 Clayman and Reisner 1998). For example, depictions of crime in newspapers can provoke public
23 anxieties and direct those anxieties towards certain groups (Lowry, Nio, and Leitner 2003; Boulahanis
24 and Heltsley 2004), change jurors' decisions and trial outcomes (Tans and Chaffee 1966), and affect
25 attitudes of law enforcement officials towards both victims and perpetrators of crime (Graham and
26 Lowery 2004; Dukes and Gaither 2017). Furthermore, in addition to influencing public perceptions
27 and treatment of different groups, discrepancies in coverage of homicide also *reflect* the complex
28 social forces from which they emanate.

29 Societal attitudes towards victims are affected by whether a victim fits the "ideal victim"
30 prototype (weak, vulnerable, etc.) (Christie 1986). Empirical evidence has shown that, when
31 describing ideal victims, more sympathetic language is used (Lewis, Hamilton, and Elmore 2021;
32 Wilson and O'Brien 2016) and, in the case of newspaper reports of homicide, there is an increased
33 level of coverage and sympathy (Bouchard, Wong, and Gushue 2020). Principal among the
34 characteristics influencing whether a victim is perceived as an ideal victim are demographic factors
35 such as the victim's age, race, and sex (Long 2021) as well as the *interactions* between such factors.
36 It is therefore important to analyze ideal victim typifications through an intersectional lens (Wilson
37 and O'Brien 2016; Steffensmeier, Ulmer, and Kramer 1998). Failing to do so risks flattening both
38 the lived experience of certain groups and our understanding of the complexity of the operative
39 social forces. Accordingly, this paper aims to paint a nuanced picture of how demographic factors
40 interact to shape the qualitative and quantitative coverage afforded to victims of violent crime.

41 We are not the first to study differential media coverage across demographic or socioeconomic
42 groups. An important existing literature has sought to analyze media distortions by measuring, e.g.,
43 differences in news coverage of crime based on characteristics of victims and offenders (Johnstone,
44 Hawkins, and Michener 1994; Weiss and Chermak 1998; Sorenson, Manz, and Berk 1998; Lundman
45 2003; Gruenewald, Pizarro, and Chermak 2009; Schildkraut and Donley 2012). The studies that
46 comprise this literature can be understood as varying principally along three dimensions: (i) How
47 media coverage is measured, (ii) which demographic groups are considered in the analysis, and (iii)
48 data availability, including which other homicide-specific factors the analysis is able to control for

49 (and thus the plausibility of a causal interpretation), which homicides are chosen for analysis (and
50 how they are matched with media coverage), and the geographic scope of the analysis.

51 Regarding the first dimension, measurement of coverage, coverage is traditionally measured
52 through readily available quantitative metrics, such as whether or not any articles are written, the
53 number of articles written, or an article’s “prominence,” (defined variously as the total number of
54 words, column length, or where the article appears within the paper). However, several important
55 recent papers have instead applied content analysis to understand the qualitative nature of homicide
56 coverage (Humphries 1981), focusing on whether offenders and victims are described in positive or
57 negative terms (Barlow, Barlow, and Chiricos 1995; Pritchard 1985; Entman 1992; White, Stuart,
58 and Morrissey 2021) or coding articles according to a strict standardized criterion, such as the
59 number of close relations of the victim (e.g., family, friends, etc.) mentioned in the article (White,
60 Stuart, and Morrissey 2021). Regarding the second dimension, demographic groups considered, the
61 most common finding in this literature has been reduced coverage of black victims, usually without
62 a distinction between sexes or age groups. Regarding the third dimension, data availability, existing
63 studies usually attempt to match homicide case reports to specific news articles by searching for
64 victim or offender names based on a single large metropolitan area and its corresponding major
65 newspaper; what geographic, temporal, and homicide-specific controls are able to be included in the
66 analysis varies significantly across studies. As we describe below, we build on this literature along
67 each of these three dimensions.

68 A central contribution of our analysis concerns our approach to measuring coverage, and
69 specifically how we determine whether a news article humanizes a victim. Our approach takes a
70 comprehensive view of the textual content of the articles in our sample rather than focusing on a
71 single dimension of humanizing language. We do this by applying natural language processing
72 (NLP) techniques to automate the classification of each article as humanizing or impersonal with
73 respect to a particular victim. Specifically, our analysis uses GPT-3, a large language model (LLM)
74 that has been pre-trained with 300 billion tokens on 45 terabytes of data obtained from a variety of
75 sources including books and public web pages (Brown, Mann, Ryder, Subbiah, Kaplan, Dhariwal,
76 Neelakantan, Shyam, Sastry, Askell, et al. 2020).¹ The value of language models in text classification
77 lies in their ability to derive meaning out of precise textual context (Jensen, Karell, Tanigawa-Lau,
78 Habash, Oudah, and Fani 2022), thus affording a great deal of flexibility in determining whether or

¹NLP has shown utility in automating and standardizing content analysis (Wang, Tian, Ones, Landers, et al. 2022; Brown 2018; Rezapour 2021).

79 not language is humanizing.² It provides a standardized and repeatable approach to classification
80 that replaces the more time-consuming method of using human coders. To limit the effect of
81 individual subjectivity, such coding approaches usually apply very strict criteria (e.g., counting the
82 number of family member references). The use of automated NLP classification allows for a more
83 holistic analysis of the text while allowing for repeatable and testable outcomes. In our use case,
84 the automation of news article textual content analysis incurs additional complexity. We wish to
85 analyze the specific language used to characterize homicide victims, and as such, the automation
86 technique must recognize and extract the precise portions of the text that refer to the victim (and in
87 the case of articles which include more than one homicide or more than one victim, the language
88 referring to one particular victim at a time). We selected GPT-3, a newer generation *generative*
89 LLM, meaning that it is specifically designed and pre-trained to generate natural language as well
90 as encode and analyze text input. The model can use its inherent knowledge of natural language
91 to analyze the entire article and generate an extract of only the portions referring to a particular
92 victim, then from the extract analyze the specific language used to describe the victim and determine
93 whether it humanizes them.

94 A second central contribution of our analysis concerns how we understand the role of
95 demographics in governing media coverage: we focus on the importance of understanding differential
96 media coverage of crime through an intersectional lens. To do this, our regression analysis explicitly
97 accounts for how race, age, and sex all *interact* in relation to media coverage. Put differently, our
98 analysis allows us to observe the conditional relationship between race on media coverage within
99 age-sex cells, and thus paints a nuanced picture of how media coverage differs across groups. This
100 is important because it avoids flattening the experience of broadly defined groups as discussed
101 above, but also because it provides much richer empirical discipline needed to discriminate between
102 various sociological theories of media coverage and crime.

103 Finally, our analysis is built around a new dataset that we construct from the FBI's Supple-
104 mentary Homicide Reports (SHR) database and the universe of *Boston Globe* articles between
105 1976 and 1984. This allows us to include nearly all homicides in an entire state, and to match
106 those homicides to nearly all corresponding articles in a major regional newspaper. Furthermore, in
107 addition to the demographic information described above (both for victim and offender), the FBI's
108 SHR data includes detailed information on the circumstances of each homicide, as well as when

²A language model predicts the most likely next words to follow given a previous text, based on examples given during the model's training phase.

109 and where it occurred. This allows us to control for a wide range of factors that might otherwise
110 confound our analysis of how demographics relate to media coverage of crime.

111 **Study design**

112 **Overview**

113 We examine demographic correlates of newspaper coverage of homicides occurring in Massachusetts
114 between 1976 and 1984. Homicide data, including date as well as demographics of both victims
115 and suspected offenders, when known, is obtained from the publicly available FBI crime data. We
116 match individual homicide reports to articles published in the *Boston Globe* morning editions about
117 each homicide and determine whether each resulted in one or more published articles. Using NLP,
118 we determine whether the newspaper articles include humanizing details about the victims or simply
119 provide basic information in an impersonal way. We apply a suite of general linear models to identify
120 the relationship between demographic factors and the amount and quality of coverage.

121 **Homicide data and news articles**

122 The universe of homicides available for our study includes all homicides reported anywhere in
123 Massachusetts in the years 1976-1984. Rather than relying on police records from individual
124 localities, we utilize data from the SHR (Regoeczi and Jarvis 2021), which captures homicide data
125 from all local and state jurisdictions. The location and particular years were selected because they
126 represent the confluence of dates between when case-level information is published in the SHR
127 (starting in 1976) and when the digitized textual transcripts of the article content of a major regional
128 newspaper (*Boston Globe*) were available in bulk format for data mining (through 1984).

129 Our analysis includes each of the 1,800 Massachusetts homicides reported in the SHR as
130 well as every newspaper article published in the *Boston Globe* for that time frame (1,215,056
131 articles, including non-news items such as advertisements and stock quotes). We develop and
132 test a filtering algorithm to automate the identification of Massachusetts homicide (MH) articles,
133 based on keyword/phrase matching of language related to homicides, Massachusetts locations
134 and neighborhoods, and article types. The algorithm is calibrated to identify all possible MH
135 candidates even though some of the articles initially identified may be false positives due to the use

136 of similar language, requiring further manual filtering to finalize the MH set. We then test the final
137 algorithm against a randomly selected set of articles labeled as MH or not MH by a human reader,
138 demonstrating 100% sensitivity and 98.2% specificity. The algorithm identified 22,640 of the 1.2
139 million articles as being potential MH, and subsequent manual inspection reduced the potential set
140 down to a final confirmed set of 5,042 MH articles.

141 Using the demographic and crime detail information in the articles, we then manually match
142 each one of the confirmed MH articles to one or more MHs from the SHR. 835 articles that referred
143 to pre-1976 MH are removed from further consideration. Details of the algorithm used for filtering
144 are included in Appendix: Article filtering algorithm. Of the remaining articles, 988 are unable to
145 have any homicide record assigned to them because they represent homicides that were clearly not
146 in the SHR database, or the article is so vague that there is not enough information to assign it to a
147 particular homicide. The final count of articles positively matched and assigned to one or more
148 SHR entries is 3,219, with a total of 1,071 homicides matched to one or more articles.

149 **Humanizing coverage**

150 Our analysis takes a broad view of how media coverage can humanize—or fail to humanize—a
151 homicide victim by considering both the quantity of coverage (reflective of, e.g., the journalistic
152 process at the editorial level) as well as the nature of the language employed in that coverage
153 (reflective of, e.g., the journalistic process at the level of the individual author). While the former is
154 straight-forward to measure (see below), the latter warrants further elaboration. We consider the
155 language used in an article to be depicting a homicide victim in a humanizing manner when at least
156 one news article mentions additional information about the victim beyond age, race, location, and
157 circumstances of death, resulting in the victim being presented as a person, not just a statistic. These
158 details usually include information about the victim’s occupation, family, background, interests,
159 and/or personal comments about the victim from family, friends and neighbors. When we determine
160 the presence or absence of humanizing language, we do not attempt to evaluate articles for evidence
161 of offensive, derogatory, or otherwise explicitly biased language. Therefore, we do not characterize
162 the absence of humanizing coverage as *dehumanizing* but rather *impersonal* coverage. Impersonal
163 (not humanizing) coverage of the homicide is when only the basic facts of the victim are presented
164 and no additional background is included.

165 GPT-3 is used to automate the determination of whether an article is humanizing or impersonal

166 in its coverage of the homicide victim. Our prompt strategy involves the use of three consecutive
167 prompts to the GPT-3 text completion interface (Figure 1). An initial prompt sends the full article
168 text to GPT-3 such that the completion text returned is an extract of the article containing exclusively
169 the information about the particular victim in question. The second prompt submits this extract back
170 with a request to summarize it by removing demographic information such as age and race of the
171 victim, as well as information about the offender and the circumstances of the crime itself. The
172 summary returned as completion text is more precisely focused on personalized victim information.
173 The third prompt sends the summary from the second prompt, requesting a final classification as
174 humanizing vs. impersonal using the few-shot technique. The actual prompts used, with examples
175 of the results, are included in Appendix: GPT-3 prompt design. To determine whether an article is
176 humanizing, we submit the three prompts to GPT-3 and use the first word (*humanizing* or *impersonal*)
177 of the final response as the result. For each homicide victim, we consecutively send each article on
178 that homicide to GPT-3 as outlined above, until at least one article is determined to be humanizing
179 or all articles are considered impersonal.

180 *FIGURE 1 ABOUT HERE.*

181 We test the accuracy of GPT-3's humanizing determination against a human reader's ground
182 truth labeling of 100 articles covering 30 randomly selected homicides, with GPT-3 agreeing with
183 the human reader for 28 of the homicides (93.3% accuracy). Out of the 1,071 homicides with news
184 articles evaluated for humanization, 602 homicides are classified as humanizing (56%).

185 **FBI-Boston Globe database**

186 The final combined dataset, along with Stata code to reproduce all of the statistical analysis, tables,
187 and figures in the remainder of the paper, has been made publicly available and is maintained at
188 <https://github.com/emily-ocasio/news>.

189 **Results**

190 Table 1 reports summary statistics for the combined dataset. Specifically, within victim age-sex-race
191 cells, the table reports (i) the share of victims described with humanizing language (Pr(Humanized)),

192 (ii) the average number of articles per homicide (Articles/homicide), and (iii) the number of
193 homicides (Homicides).

194 *TABLE 1 ABOUT HERE.*

195 The goal of our empirical analysis is to use the data described above to understand how
196 demographic characteristics influence the extent to which media coverage conveys homicide victims
197 in a humanizing light. Specifically, the detailed homicide data from the SHR enables us to study rich
198 *interactions* among demographic factors. It also allows us to control for demographic-correlated
199 factors related to the nature of the homicide that would otherwise be likely to confound an analysis
200 that only included demographic factors. The humanizing score that we construct with GPT-3 enables
201 us to study how the demographic factors interact to shape humanizing coverage.

202 The analysis that follows first describes our empirical strategy and then presents three sets
203 of results, each corresponding to a different measure of humanizing language: Our main results
204 are based on a *composite humanizing score*, which measures the use of humanizing language in
205 the full sample of homicides, including those receiving no coverage³ and those receiving coverage
206 (which we categorize as humanizing or non-humanizing as described above). Next, we consider
207 a *news prominence score*, which measures the quantity of coverage—i.e., the number of articles
208 written about a homicide, irrespective of the nature of the language used in those articles. Finally,
209 we consider a *conditional humanizing score*, which instead measures the quality of coverage—i.e.,
210 the likelihood of humanizing language, conditional on a homicide receiving coverage.

211 **Empirical strategy**

212 Formally, let h_v denote any of the three measures described above of the extent to which coverage
213 of victim v is humanizing, and let X_v denote the corresponding vector of demographic and other
214 controls from the FBI’s SHR database for victim v , including:

- 215 1. Victim demographic controls: Race, sex, and age;
216 2. Offender demographic controls: Race, sex, and age;⁴

³An absence of coverage is deemed to be non-humanizing.

⁴We include an indicator for whether or not a murder was solved, which is implicitly an offender demographic control because, if a homicide is unsolved, no offender demographic information is available.

217 3. Homicide controls: Weapon, relationship between victim/offender, specific circumstances
 218 of the homicide (e.g., gang-related), whether there were multiple victims, and whether there
 219 were multiple offenders;

220 4. Time (month-by-year) and county fixed effects.

221 For all three measures of humanizing coverage, we estimate the relationship between h_v and X_v
 222 within a generalized linear regression framework, i.e.

$$E[h_v|X_v] = f^{-1}(X'_v\beta)$$

223 where β is the vector of coefficients we seek to estimate and f is the link function mapping X_v to
 224 the conditional mean of h_v .

225 **Composite humanizing score**

226 Our main results are based on the *composite humanizing score* ($h_v = h_v^{\text{comp}}$), which we define as an
 227 indicator variable that takes on a value of one if at least one article using humanizing language is
 228 written about homicide victim v and zero otherwise:

$$h_v^{\text{comp}} = \begin{cases} 0 & \text{No humanizing articles (incl. no articles)} \\ 1 & \text{At least one humanizing article} \end{cases}$$

229 Thus, there are two ways in which a homicide victim can fail to be humanized: Either no
 230 article is written, or at least one article is written but none uses humanizing language.

231 We begin with a simple analysis in which we only consider a subset of demographic controls—
 232 victim/offender race and sex—thus neglecting age as well as potentially important interactions
 233 among these variables, to which we return below. Table 2 reports the results from this analysis. The
 234 first two columns correspond to estimation of a simple linear probability model ($f^{-1}(X'_v\beta) = X'_v\beta$)
 235 to facilitate interpretation of parameters, and the second two columns correspond to estimation of a
 236 Probit model ($f^{-1}(X'_v\beta) = \Phi(X'_v\beta)$, where Φ is the cumulative distribution function of the standard
 237 normal distribution). Within each functional form specification, we report results when the model is
 238 estimated only with demographic controls as well as with the full set of controls listed above.

239

TABLE 2 ABOUT HERE.

240 The results reveal three important insights into how humanizing coverage varies across
241 demographic groups. First, referring to the results from the linear probability model, on average
242 black victims are between 15 and 18 percentage points less likely to receive humanizing coverage
243 than white victims (first row). Second, on average female victims are between 13 and 17 percentage
244 points more likely to receive humanizing coverage than male victims (second row). Both results are
245 significant at all conventional levels and robust to the inclusion of a rich set of controls and fixed
246 effects. Third, in the models that fail to control for non-demographic factors (first and third columns),
247 there is a significant positive relationship between an offender being black and humanizing coverage
248 of the victim—homicides committed by black offenders appear to be roughly eight percentage points
249 more likely to be described in a way that humanizes the victim than those committed by white
250 offenders. However, when we account for the rich set of homicide-specific controls included in the
251 FBI data, as well as time and locality fixed effects, this effect shrinks dramatically in magnitude and
252 ceases to be distinguishable from zero. This observation highlights the importance of controlling for
253 non-demographic factors that are likely to be correlated across demographic groups when analyzing
254 the determinants of media coverage of crime.

255 While the analysis in Table 2 sheds important light on some demographic determinants
256 of humanizing coverage of homicide, it fails to account for the rich ways in which intersectional
257 demographic factors *interact* to shape that coverage.⁵ Accordingly, we re-estimate the fourth
258 specification in Table 2, augmented with fully-interacted indicators for race, age, and sex.⁶ To
259 facilitate exposition of the results, Figure 2 plots the predicted probabilities of humanizing coverage
260 disaggregated by race, age, and sex. Specifically, the left panel plots the predicted probability of
261 humanizing coverage for white and black male victims by age, and the right panel does the same for
262 female victims.⁷ Table 3 reports formal tests of equality of the predicted probabilities for black and
263 white victims within each age-sex cell.

⁵Indeed, including age and demographic interactions in the regression in Table 2 increases the adjusted- R^2 by more than 15%.

⁶Throughout the analysis, we focus on black and white victims under the age of 70 and partition victims and offenders into four age groups: children (< 18), young-adult (18-29), middle-aged (30-49), and older (50-69). Other groups do not contain enough observations for identification. Additionally, an interesting extension of our emphasis on demographic interactions would be to study whether there are important interactions between demographic characteristics of victim and offender. Unfortunately, our sample size is too small to meaningfully identify such effects, but we view this as a fruitful avenue for future work.

⁷Predictive margins are computed by integrating over all covariates distributed as observed in the data.

264

FIGURE 2 ABOUT HERE.

265

TABLE 3 ABOUT HERE.

266

267

268

269

270

271

272

273

274

275

276

277

278

279

280

281

282

283

284

285

286

287

288

The results in Figure 2 reveal a much richer and more nuanced picture of how demographics interact to shape media coverage of crime compared with Table 2. First, among male victims, the probability of humanizing coverage is significantly higher for white victims than for black victims in all age categories, with the greatest difference among juveniles (30 percentage points) and an average difference of between 12 and 16 percentage points for all other age groups. As may be seen in Table 3, all of these differences within age groups are statistically significant at the 5% level, and all but the difference within the 50-69 year old category are statistically significant at the 1% level.⁸ Second, the same pattern does not hold for females: in general, there is not a statistically significant difference between white females and black females. However, a difference does emerge among young-adult females, in which case white victims are 23 percentage points more likely to receive humanizing coverage than black females. Third, females generally receive more coverage than their male counterparts, regardless of age. Finally, young males (18-29) of both races are less likely to be covered in a humanizing light than their older or younger counterparts.

In Appendix: Robustness, we include results for various robustness tests of our main specification in Figure 2. Specifically, Figure A1 replaces our relatively coarse county fixed effects with much more granular FBI agency fixed effects, thus providing a better control for geographic variation; Figure A2 trims our sample to exclude outlying homicides—those about which more than 20 articles were written—which could potentially be driving our results; and Figure A3 excludes homicides that occurred during 1984, since we are unable to obtain newspaper articles written after 1984 which could potentially concern homicides occurring during 1984.⁹ Reassuringly, none of these alternative specifications materially change our results.

The dependent variable in our analysis to this point—the composite humanizing score—implicitly reflects two ways in which media coverage can humanize a victim: through the quantity

⁸Note that overlapping confidence intervals in Figure 2 need not indicate a statistically insignificant difference. See the case of 50-69 year old males for a case in which the 95% confidence intervals overlap but in which the difference is significant at the 5% level.

⁹In principle, articles may have been written after 1984 about homicides that occurred in 1983, 1982, etc., although this is unlikely to be a significant issue since the vast majority of articles written about homicides were written shortly after the homicide occurred. We thus view dropping all homicides that occurred during 1984 from the sample in this robustness check as a conservative test.

289 of coverage (e.g., whether or not an article is written about a homicide irrespective of the language
290 of the reporting) and through the quality of that coverage conditional on an article being written.
291 Below, we study these two sources of humanizing coverage in greater detail.

292 **News prominence score**

293 Because we have linked homicides from the FBI's SHR database to *all* corresponding *Boston Globe*
294 articles covering those homicides between 1976 and 1984, we are able to measure the number
295 of articles written about any given homicide. This represents a natural alternative measure of
296 humanizing media coverage—over the course of multiple articles, a richer and more complex
297 picture can be painted of a victim than would be possible in a single article (or, indeed, no
298 articles)—reflective of the editorial decision-making process. Accordingly, we define the *news*
299 *prominence score* ($h_v = h_v^{\text{prom}}$) as the number of articles written about a given homicide:

$$h_v^{\text{prom}} = \# \text{ of articles}$$

300 and estimate a Poisson regression ($f^{-1}(X'_v\beta) = \exp(X'_v\beta)$) where the dependent variable is now
301 h_v^{prom} . The set of explanatory variables X_v is identical to those used in the analysis in Figure 2.
302 Figure 3 summarizes the predicted number of articles written, broken down again by race, sex, and
303 age, and Table 4 reports formal tests of equality of the predicted number of articles for black and
304 white victims within age-sex cells.

305 *FIGURE 3 ABOUT HERE.*

306 *TABLE 4 ABOUT HERE.*

307 With this alternative measure of humanizing coverage, we observe some similarities as well as
308 some important differences relative to the results described above. Among the important similarities
309 are systematically more coverage of white victims than black victims among males, virtually no
310 racial discrepancy at any age among females, and the tendency for young black males (18-29) to
311 receive less coverage than all other groups. However, when the dependent variable is the number of
312 articles, the gap between males and females is substantially larger than previously. Moreover, there

313 is no gap between white and black coverage of females at any age group, whereas the difference
 314 between white and black males is only robustly significant among young (18-29) and middle-aged
 315 (30-49) males (second and third rows of Table 4). The gap between black and white male children is
 316 once again the most pronounced, but the difference in this case is not significant (first row of Table
 317 4).

318 **Conditional humanizing score**

319 Our third and final approach to understanding what determines humanizing language focuses
 320 strictly on the language used in the articles rather than the probability of coverage or the number of
 321 articles. Accordingly, we define the *conditional humanizing score* ($h_v = h_v^{\text{cond}}$) as the probability of
 322 humanizing language conditional on at least one article being written:

$$h_v^{\text{cond}} = \begin{cases} 0 & \text{No humanizing articles — } \geq 1 \text{ article} \\ 1 & \text{At least one humanizing article — } \geq 1 \text{ article} \end{cases}$$

323 and estimate a Probit model ($f^{-1}(X'_v\beta) = \Phi(X'_v\beta)$) where now the dependent variable is h_v^{cond}
 324 and we thus restrict attention to only those homicides about which at least one article was written.
 325 Furthermore, in addition to the standard set of controls used in Figures 2 and 3, we also control for
 326 the total number of articles written about a homicide, which will tend to be mechanically correlated
 327 with our dependent variable due to our approach to coding humanizing articles (see Materials and
 328 Methods section for details). Figure 4 reports the results from this analysis (analogously to Figures
 329 2 and 3) and Table 5 reports formal tests of equality of the predicted probabilities for black and
 330 white victims within age-sex cells (analogously to Tables 3 and 4).

331 *FIGURE 4 ABOUT HERE.*

332 *TABLE 5 ABOUT HERE.*

333 The results in Figure 4 and Table 5 add yet another layer of nuance to our understanding of the
 334 determinants of humanizing coverage. Specifically, we now observe patterns that are qualitatively
 335 quite similar to the patterns that we observed in Figure 2: A marked gap between humanizing

336 coverage of black and white male victims with young males (18-29) receiving less coverage than
337 other age groups and no universal difference between the language used to describe black and white
338 female victims. Notwithstanding this second observation, and also similar to Figure 2, we see a
339 pronounced racial gap among young (18-29) females, although in this case the gap is not statistically
340 significant whereas it was in Figure 2. On the other hand, the gap between males and females of all
341 ages is somewhat diminished relative to what we see in Figure 2, and the pervasive black-white
342 gap among males that was significant for all age groups in Figure 2 is now only robustly significant
343 among young (18-29) and middle-aged (30-49) males, similar to the results in Figure 3 and Table 4.

344 **Discussion**

345 Our study provides robust empirical support for the hypothesis that newspaper depictions of
346 homicides disproportionately humanize victims of certain demographic groups more than others.
347 Our data does not support the idea that newsworthiness is mostly explained by an event's rarity.
348 Although the observed greater degree of humanizing coverage of juvenile victims can be understood
349 as a natural consequence of the relative infrequency of juvenile homicides, this explanation does not
350 account for what we observe when comparing juvenile homicide coverage of young males by race:
351 white male juveniles are much more likely to be humanized than black male juveniles despite the
352 fact that there were almost three times as many white juvenile victims as there were black juvenile
353 victims. In this way, our intersectional analysis allows us to discriminate between theories in a way
354 that would not otherwise be possible.

355 Differences across demographic groups are brought into sharp contrast when measured
356 intersectionally. Across age groups, we observe significantly less humanizing coverage for black
357 male victims as compared to white males, while white female victims overall are not more likely to
358 be humanized than black female victims. However, when we account for differential effects across
359 age groups, significant differences are seen between certain groups: black females between 18 and
360 29 are significantly less likely to be humanized than white females of the same age. Moreover, while
361 black males are humanized less across all age groups, the gap is largest for juveniles.

362 These findings are consistent with ideal victim theory: young people are seen as more
363 vulnerable, while being a black victim counteracts this by being seen as aggressive or strong. While
364 age, sex, and race all influence societal perceptions of innocence and vulnerability, they act in
365 different directions across groups. Our results can be understood as reflecting identified stereotypes

366 towards black females and males; younger black females, for example, are masculinized and seen as
367 strong, falling into the “Jezebel” or “strong black woman” trope, while older black females instead
368 fall into the “mammy” stereotype, which is associated with caring and motherhood. Similarly, for
369 males, the largest difference in the humanizing probability is between black and white juveniles.
370 This is the most vulnerable and innocent age group, so the observed difference is consistent with
371 dehumanizing perceptions of black boys in particular as older, stronger, and less innocent (Goff,
372 Jackson, Di Leone, Culotta, and DiTomasso 2014).

373 Ideal victimhood thus provides the most consistent explanation of the humanization patterns
374 we observe in our data. Children are the most ideal victims and thus they are the most likely age
375 group to be humanized, followed by the elderly. Females are perceived as weaker, more innocent,
376 and subsequently more ideal victims than males and thus they are more likely to be humanized.
377 When additionally considering race, we see specific interactions that are also explained by ideal
378 victim theory—younger adult black females and black males of all ages, but most particularly young
379 black males, are less likely to be humanized because they are perceived as stronger, less vulnerable,
380 and less innocent.

381 The results also demonstrate the importance of a methodological approach to media distortion
382 analysis that combines the presence/extent of coverage (quantitative measures) with determination
383 of humanizing language (qualitative measures). When we use the composite humanizing score,
384 we detect significant differences between the humanizing probabilities of black and white females
385 between 18 and 29, whereas in the news prominence score regression, which instead considers
386 article count, we fail to detect this difference. The conditional humanizing score regression, which
387 considers the humanizing categorization only among the subset of homicides about which at least
388 one article is written, predicts a small difference which is not statistically significant. In addition,
389 our use of GPT-3, a generative LLM, to perform content analysis from articles with varied and
390 complex text, shows the viability and scalability of this automation technique.

391 **Limitations**

392 One potential limitation of our study is inherent in the lack of a source of crime reports that reliably
393 includes 100% of the homicides in the analyzed location and time period. We thoroughly search all
394 newspaper articles for MH-related stories independent of the crime reports, and subsequently find a
395 number of cases not included in the SHR. This gap has not usually been identified or measured in

396 media distortion analysis since the articles are usually found from a search by victim and/or offender
397 name based on the available crime reports. Nevertheless, studies of SHR accuracy have validated
398 their reliability when compared with aggregate data from the National Vital Statistics System, with
399 the exception of cities with very low numbers of black homicides (Loftin, McDowall, Curtis, and
400 Fetzer 2015; Rokaw, Mercy, and Smith 1990). Specifically, some infrequent kinds of homicides,
401 in particular justifiable homicides (e.g., law enforcement killings of felons during commission of
402 a crime) show very low rates of inclusion in SHR (Finch, Thomas, Beck, Burghart, Klinger, and
403 Johnson 2022).

404 The use of automated content analysis via GPT-3 allows for more scalable and consistently
405 reproducible determination of humanizing article language, especially when analyzing large numbers
406 of articles. Although it eliminates human variance and coding error, it may introduce biases arising
407 from the language model itself. This potential for bias has been recognized by the developers of
408 GPT-3 (Brown, Mann, Ryder, Subbiah, Kaplan, Dhariwal, Neelakantan, Shyam, Sastry, Askell,
409 et al. 2020). When GPT-3 was provided racial cues during prompting, the developers noticed the
410 resulting text differed in sentiment based on the prompted race. This was most likely due to the
411 inherent discrepancy in depictions of race among the millions of source articles used to train the
412 model. Our use of GPT-3 does not explicitly provide race in the prompts and also removes whatever
413 references to demographic characteristics may be in the article during the second prompt (Figure 1),
414 so the final determination of humanization is as race- and age-agnostic as possible. Gender is harder
415 to remove completely from the prompts because of the associated pronouns and the presence of
416 gender-specific names (this may also be an issue with race to a much smaller extent). In addition to
417 making the prompts neutral, we also do not ask GPT-3 to generate new language, only to summarize
418 and extract information from existing articles, as well as to make a final determination. Future
419 studies should explore the effects of artificially adding, removing, or changing racial depictions in
420 prompts on GPT-3 humanization classification accuracy.

421 Finally, the usual qualifications regarding internal and external validity apply to our study.
422 In terms of internal validity, we emphasize that our estimates do not necessarily imply a causal
423 relationship between a victim's demographics and humanizing coverage. While inclusion of the
424 rich set of interacted demographic controls and non-demographic fixed effects made available in
425 the SHR goes a long way towards accounting for the possibility of correlated confounding factors
426 explaining the demographic patterns we observe, we cannot rule out that other such factors may
427 be unobserved and driving the relationships we find in the data. In terms of external validity, our

428 study focuses on a specific region and time period—Massachusetts in the 1970s and 1980s. This
429 particular choice was driven by the requirement that the entire universe of published articles from a
430 source which covers crimes statewide be considered in the analysis so that even brief or secondary
431 mentions of each homicide are accounted for. This could only be done using comprehensively
432 organized and converted to textual form newspaper articles from a time period where the SHR
433 database also exists. As discussed above, while the fact that the patterns of disparate coverage that
434 we document are consistent with various identified stereotypes is suggestive of the broad relevance
435 of our results, it is nevertheless possible that different social forces may be operative in different
436 time periods and regions. This study provides a road map for future research, which we hope will
437 apply our methodology to study disparate media coverage in different localities and time periods as
438 their text becomes available in full for data mining.

439 **Conclusion**

440 We document significant differences in both the quantity and quality of media coverage of homicide
441 across victims from different demographic groups, with stereotypical ideal victims more likely to be
442 humanized than those considered non-ideal victims. Average differences across races are large and
443 pronounced, with black victims receiving significantly less coverage than white victims, consistent
444 with historical stereotypes surrounding innocence and blackness. Yet once we allow for differential
445 effects of race across age groups and sex, we find that the average effects of race conceal substantial
446 heterogeneity across different sub-populations: For example, the fact that black victims receive less
447 humanizing coverage than white victims on average is driven almost entirely by males; the only
448 difference across races among females appears among 18-29 year-olds; and the racial gap among
449 males is twice as large among juveniles as it is among older age groups.

450 Our findings thus demonstrate that failing to account for intersectionality in analyses of
451 media coverage of crime both risks the erasure of important variation in human experience across
452 groups and makes it difficult to discriminate between sociological theories. In addition, our findings
453 demonstrate how content analysis (facilitated by the use of a state-of-the-art large language model,
454 GPT-3) can be applied to produce holistic measures of the extent to which media language is
455 humanizing and how that language is applied unequally. We hope that future researchers will learn
456 from these insights and continue to work towards deepening our understanding of how the media
457 both reflects and influences how society perceives members of different demographic groups.

458 **Appendix: Methodological Details**

459 **Homicides**

460 The SHR data was downloaded from the Murder Accountability Project website (Hargrove 2022).
461 It includes individual entries for each homicide, with year and month of the homicide, victim
462 demographics, offender demographics if there is a known offender, the police department that
463 reported the case, and other information about the homicide such as the method of death (shooting,
464 stabbing, strangulation, etc.), the offender relationship with the victim, the kind of weapon, and
465 circumstances of the case (whether it was related to argument, robbery, drug-related, etc.). The
466 SHR does not include personal information on the victims or offenders, such as names or addresses.

467 **News articles**

468 We obtained the text of newspaper articles from the ProQuest Historical Newspapers database,
469 which includes the digitized text of every *Boston Globe* article published through 1984, available
470 in bulk for analysis via Text Data Mining Studio. For our analysis, we extracted the text of every
471 article published in the morning edition for the years 1976 through 1984. These text transcripts were
472 originally produced via optical character recognition from scanned paper copies of the newspapers,
473 and as such they contained some errors such as misspelled words, occasional extraneous spaces
474 within the words, uneven sentence formatting, and lack of paragraph separation. Despite these
475 issues the article texts were readable and understandable. The article database also includes
476 additional metadata about each article, including the date, title, and article type (e.g., news article,
477 advertisement, editorial, stock quote).

478 **Article filtering**

479 Our automated filtering algorithm included three main criteria for determining whether an article
480 may be a MH: (1) article type, (2) presence of keywords or word combinations related to homicides,
481 and (3) reference to Massachusetts locations. The article type criteria removed from consideration
482 database entries such as classified ads and paid birth/death notices. The algorithm looked for the
483 presence of homicide-related keywords (e.g., “murder,” “slain”) as well as word combinations
484 (e.g., “shot to death”). It also ensured the articles referred to homicides in Massachusetts by

485 ensuring reference to a location from a list of towns, cities, and counties in Massachusetts as well as
486 neighborhoods and places in Boston.

487 Our goal for the algorithm was to ensure it did not miss any MH articles (false negatives)
488 while reasonably minimizing the number of non-MH articles it erroneously marked as MH (false
489 positives). We randomly selected article subsets to measure the algorithm’s accuracy during the
490 initial development round, three rounds of algorithm optimization validation, and a final testing
491 round. Articles for each subset were assigned by randomly picking dates within the 1976-1984
492 time span and including in the subset all the articles for the selected dates. Every article in all the
493 subsets was ground truth labeled MH/not MH by a human reader, and sensitivity and specificity of
494 the algorithm results against the ground truth were computed at the end of each round. The final
495 algorithm specifications are described in Appendix: Article filtering algorithm

496 Once the algorithm was finalized it was used to perform the first pass filter of the full article
497 database. Articles identified as possible MH by the algorithm were then manually inspected by a
498 human reader who further filtered the group down to the final number of validated MH articles used
499 in the analysis.

500 **Homicide matching to articles**

501 Because names of victims and offenders are not included in the SHR, we used information in the
502 news articles including date and location of homicide, victim and offender demographics, and other
503 facts such as type of weapon, circumstances and victim-offender relationship, to identify each SHR
504 homicide entry by victim name. Based on this we were able to determine whether each MH article
505 matched to zero, one, or more SHR entries. The final multiplicity relationship between articles and
506 homicides was zero to many in both directions.

507 **Humanizing determination via NLP**

508 Like other similar LLMs, GPT-3 is a pre-trained AI model and thus it does not require large training
509 sets for classification. We used the few-shot prompt technique (Brown, Mann, Ryder, Subbiah,
510 Kaplan, Dhariwal, Neelakantan, Shyam, Sastry, Askell, et al. 2020) in which, using natural English
511 text, a small number of examples are provided as a way to “prompt” the model to follow the
512 pattern. To optimize the accuracy of our prompt design we manually adjudicated a random subset
513 of homicides with their respective articles as humanizing or impersonal ground truths and used this

514 validation set to tweak the prompts and improve the accuracy. Once all three prompts were deemed
515 to be optimized, we measured the prompt design’s performance against a randomly selected test set
516 of homicides for which we manually determined ground truths.

517 The prompts with the accompanying article or extract text were submitted via the GPT-3 text
518 completion interface (<https://api.openai.com/v1/completions>) using the text-davinci-002 model with
519 the following interface parameters: temperature = 0, max_tokens = 256, frequency_penalty = 0,
520 presence_penalty = 0. The temperature parameter in particular removes the element of randomness,
521 guaranteeing the same response after repeated identical prompts. The Davinci GPT-3 model has an
522 inherent limit of the length of the text it can analyze. This limit of 4,000 tokens (each token is a
523 meaningful part of a word; usually 750 words would average out to 1,000 tokens) includes both
524 the prompt and the response. There were three homicides in the total set for which one of their
525 respective articles was so long that we could not evaluate via GPT-3. For these three homicides we
526 determined their humanizing classification by human inspection.

527 **Regression analysis**

528 The empirical framework for our regression analysis is described in the Results section above.
529 Probit regressions are estimated in Stata 17 using the *probit* command and Poisson regressions are
530 estimated in Stata 17 using the *ppmlhdfe* package (Correia, Guimaraes, and Zylkin 2019).

531 **Data Availability**

532 **Appendix: Article filtering algorithm**

533 **Description**

534 In summary, the algorithm has three main filters. An article must pass all three filters in order to be
535 classified as a MH. Filter 1 is that the article must be one of the allowed article types (types are
536 determined within the ProQuest database). Standard types such as “article,” feature,” “cover story,”
537 and also “editorial” and “letters to the editor” are allowed for consideration as possible MH-related
538 types. Types such as “front matter,” and “table of contents” are filtered out. Although we initially
539 also fully filtered out the “advertisement” type, we noticed during validation phases that a small
540 number of actual news articles, including homicides, were erroneously typed as “advertisement” in

541 ProQuest so that rule had to be modified to make sure the title itself contains a reference to “ad”
542 before it is filtered out. Table A1 contains the full list of types that were allowed and filtered out.

543 *TABLE A1 ABOUT HERE.*

544 Filter 2 uses partial word text matching to determine whether certain words or combinations
545 of words refer to homicide deaths. For the homicide words, certain word roots such as “murder”
546 or “homicide” are considered *absolute*, such that any single inclusion of those roots (beginning
547 of the word matching the root) marks the article as homicide-related. Other word roots such as
548 “shot,” “wounded,” and “stabb” were considered *relative*, meaning that they needed to be present in
549 proximity to a *death-related* root such as “death” or “fatal.” The result of this relative rule is that
550 the phrases “shot to death” and “fatally wounded” would mark an article as a homicide but not
551 “fatal attraction” or “took a shot.” The rule defined proximity as a relative and death-related root
552 being within 300 characters of each other. Filter 2 removed from consideration articles that had no
553 words with absolute roots and also no words with conditional roots in proximity to a word with a
554 death-related root. Table A2 enumerates the roots used for each of the three categories. This filter
555 was implemented in Python using regular expressions.

556 *TABLE A2 ABOUT HERE.*

557 Filter 3 ensures that the article refers to a Massachusetts location. We use a list of all cities,
558 towns and counties in the state, as well as all the known Boston neighborhood names. Articles
559 with no mention of a location name anywhere in the text are filtered out. There are three very
560 small localities in the state whose names are very commonly used in a different context, namely
561 “Lee,” “Florida,” and “Washington” and thus we removed them from the localities list because they
562 produced a very large number of false positives during our validation, while there was not a single
563 case of a death in one of those localities where the only reference was one of those three names.
564 This filtering was implemented in Python using the *flashtext* library.

565 **Validation and Testing**

566 To validate the accuracy of the automated algorithm, we compared its performance with an
567 independent human determination (after reading the full article) of whether it was MH-related or

568 not. After initial selection of the algorithm’s parameters, we defined a subset of articles by randomly
569 selecting a number of dates within our date range and including every article for those days. We
570 applied the algorithm filters and simultaneously read each article and marked it as MH (positive) or
571 non-MH (negative). The goal was to ensure no false negatives occur (algorithm determined the
572 article was non-MH where human detected it as MH) while reasonably reducing the false positives
573 (algorithm marked a non-MH article as potential MH). Since a human subsequently examined
574 all articles which the algorithm marked as MH, false positives were eventually removed from
575 consideration, while any potential false negatives would represent an article that actually referred to
576 a MH article but would have been missed from our analysis. After the initial round of comparison,
577 we adjusted the algorithm to try to make it stricter while still not producing false negatives and
578 then repeated validation using a different random article sample. We performed three rounds of
579 adjustment and validation, and once satisfied with the accuracy, performed final testing on another
580 random sample never used for algorithm adjustment or validation. Accuracy results for the initial,
581 validation, and final testing subsets are shown in Table A3 based on the final algorithm logic. Ground
582 truth for all subsets was determined before the algorithm was run.

583

TABLE A3 ABOUT HERE.

584 Common examples of false positives which were filtered out by post-algorithm manual
585 processing included a number of sports-related articles that use metaphors such as “this team
586 murdered their opponents in the court last night,” as well as cases where the murders did not occur
587 in Massachusetts but there were references to places elsewhere in the world that share names with
588 Massachusetts places such as Cambridge, England. In addition, our analysis specifically excluded
589 vehicular homicides (though they are sometimes be considered “homicides” in a legal sense, they
590 are not captured in the FBI database).

591 **Appendix: GPT-3 prompt design**

592 Using a validation set of articles for which the ground truth was determined by a human reader,
593 we optimized the GPT-3 prompt design to improve its accuracy. The final strategy involves three
594 separate prompts, as illustrated in the main text of the study. The first prompt includes the full
595 text of the article and is meant to simplify the text to specifically focus on the information about

628 murder in Boston so far this year There Nere 71 murders in
629 the city all of last year It's very unusual to
630 have five killings in three days said Maillet For some reason
631 don't know why seem to pick up during the holiday season
632 Maybe there s more drinking There's lot of frustration any
633 There were developments reported yesterday in
634 three of the other four murders that have occurred in Boston
635 since last Friday

636 James Baumgardner 23, of 40 Linsky Barry court Silth Boston
637 walked into District 6 and surrendered himself in connection
638 with the shooting death early Saturday of 28-year-old James
639 Keenan of Adams street Dorchester Baumgardner was charged with
640 manslaughter Police said Keenan and friend
641 were driving van aid had stopped to talk to some women outside
642 Zimbee's bar at and West Fourth streets in South Boston when
643 fight broke out between the two men and the women s escorts
644 2 Police said Keenan and Lindberg drove off
645 but short time later Police said Keenan was armed with shotgun
646 Agroup of people subseqently converged on the van and scuffle
647 ensued between Keenan and Baumgardner shot rang out Keenan was
648 later pronouced dead at the scone from single gunshot wound
649 to the head

650 The shot that killed Keenan was from the gun he had iAi his
651 possession said District 6 detective Jack Parlon Whether that
652 was his gun or not we don't know yet Baumgardner contends that
653 he was struggling for the gun with Keenan when shot went off
654 Police said Baumgardner was arraigned at South
655 Bostg District Court and later released on ball probable
656 hearing was set for Dec 12.

657 'Meanwhile 28-year-old Jamaica Plain man was yesterday in Roxbury
658 District Court on murder carges stemming from the shooting
659 death early Saturdij morning of Miguel Santana 20, of Highland

660 avenue Bosbury
661 'Ramon Cartagena of Haverford street Jamaica Plain wiss ordered
662 held without bail in Charles Street Jail cause hearing was set
663 for Nov 28 at Roxbury Dlitric Court 'In the
664 murder of Margaret Hammond the 69-year-old who was found dead
665 Friday in her Brighton aptment police issued an arrest
666 warrant for Roert McKay 57, who was said to have lived with
667 Hamat 3 Lothian rd
668 McKay was still at large last night autopsy performed on Hammond
669 showed she died massive skull fracture caused by blow from
670 instrument Police said she was found wearing red print dress
671 covered with bedding She had been dead fai several days before
672 neighbors notified police after det tting decaying odor
673 'Tolice reported no arrests or new leads in
674 connection the murder of 28-year-old cab
675 driver Jeffery Bi of Swampscott Boyajian was robbed and killed
676 last Friday apparently by three passengers he 6 to dead end
677 street at the Archdale Housing in Roslindale

678 This article contained spelling and formatting errors mostly secondary to the OCR, but
679 is representative of the quality of trascribed text in the ProQuest database (some articles have
680 better OCR with fewer errors). This news story references five separate homicides, but the prompt
681 specifically requests the language model to extract information about one victim (James Santos).
682 The response from GPT-3 was:

683
684 James Santos , 28, of 17 Horan Way, Jamaica Plain , was killed
685 early Sunday morning after an argument with 20-year-old James
686 Abernathy of Roxbury. Santos was stabbed in the abdomen and
687 taken to Carney Hospital , where he was pronounced dead at 1:58
688 a.m. His sister Nina Santos was treated for bruises at the
689 same hospital and later swore out a complaint against
690 Abernathy for assault and battery. Abernathy was taken to

691 Charles Street Jail , where he was ordered held without bail. A
692 hearing was set for Nov. 28 at Roxbury District Court.

693 The language model clearly extracted the relevant information about the particular victim
694 and removed much of the other extraneous information in the article. It also produced text written in
695 correct English with no spelling or transcription errors. This extract can now be used in a subsequent
696 prompt to determine whether the language is humanizing. Although it was possible to construct a
697 prompt that directly requested a humanizing determination for a particular victim from the initial raw
698 article, we found that it resulted in more errors, particularly in complex articles like this one, where
699 it picked up humanizing characteristics from portions of the article that were not directly related to
700 the victim in question. In addition, the final humanizing prompt also attributed humanizing factors
701 to basic demographic information such as the victim’s age or humanizing details about the offender.
702 Thus, we added an additional second prompt which slightly modified the extract specifically to
703 remove age, gender, address, and information about the suspect. In this example the streamlined
704 extract returned by the second prompt was as follows:

705 James Santos was killed early Sunday morning after an argument
706 with another person. Santos was stabbed in the abdomen and
707 taken to Carney Hospital , where he was pronounced dead at 1:58
708 a.m. His sister Nina Santos was treated for bruises at the
709 same hospital and later swore out a complaint against the
710 other person for assault and battery. The other person was
711 taken to Charles Street Jail , where they were ordered held
712 without bail. A hearing was set for Nov. 28 at Roxbury
713 District Court.

714 In this case, the victim’s age and address were removed from the text, and the offender is now
715 referred to as “another person.” As a result, an article that originally had lots of detail about many
716 different things is now distilled down to only the facts (other than demographics) about the victim.
717 In this case, it turns out that nothing other than the circumstances of death (and the fact that his
718 sister was also a victim) are included, and so the response to the final prompt declares the article as
719 impersonal about this specific victim.

720 The third and final prompt uses the few-shot technique, where we provide several examples
721 of how to respond. This is not considered training—the NLP logic is embedded in GPT-3’s large

722 pretrained model. What the few-shot technique accomplishes is guiding the model’s answer towards
723 a specific structure, in this case a sentence starting with the word “Humanizing” or “Impersonal”
724 followed by a short explanation of the reasoning behind the decision (we used this explanation in the
725 process of tweaking the prompt design to infer the reasons for incorrect responses). The few-shot
726 prompt is not “teaching” the model how to classify the article, but rather showing the model in what
727 way we would like the answer to be formatted.

728 **Appendix: Robustness**

729 Figures A1, A2 and A3 provide robustness tests of the results using the composite humanizing score
730 from the main text. See the main text for further discussion.

731 *FIGURE A1 ABOUT HERE.*

732 *FIGURE A2 ABOUT HERE.*

733 *FIGURE A3 ABOUT HERE.*

734 **References**

- 735 Barlow, Melissa Hickman, David E Barlow, and Theodore G Chiricos. 1995. “Economic
736 conditions and ideologies of crime in the media: A content analysis of crime news.” *Crime &*
737 *Delinquency* 41:3–19.
- 738 Bouchard, Jessica, Jennifer S Wong, and Kelsey Gushue. 2020. “The differential portrayal of
739 ‘sympathetic’ homicide victims in the media.” *International Review of Victimology* 26:314–
740 331.
- 741 Boulahanis, John G and Martha J Heltsley. 2004. “Perceived fears: The reporting patterns of
742 juvenile homicide in Chicago newspapers.” *Criminal Justice Policy Review* 15:132–160.
- 743 Brown, Christopher Henry. 2018. *Analyzing group behavior from language use with natural*
744 *language processing and experimental methods: three applications in political science and*
745 *sociology*. Ph.D. thesis, The University of Texas at Austin.
- 746 Brown, Tom, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal,
747 Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. “Language
748 models are few-shot learners.” *Advances in neural information processing systems* 33:1877–
749 1901.
- 750 Chermak, Steven M. 1995. *Victims in the news: Crime and the American news media*. Westview
751 Press Boulder, CO.
- 752 Christie, Nils. 1986. “The Ideal Victim.” In *From Crime Policy to Victim Policy: Reorienting the*
753 *Justice System*, edited by Ezzat A. Fattah, chapter 1, pp. 17–30. London: McMillan.
- 754 Clayman, Steven E and Ann Reisner. 1998. “Gatekeeping in action: Editorial conferences and
755 assessments of newsworthiness.” *American Sociological Review* pp. 178–199.
- 756 Correia, Sergio, Paulo Guimaraes, and Thomas Zylkin. 2019. “PPMLHDFE: Stata module for
757 Poisson pseudo-likelihood regression with multiple levels of fixed effects.” .
- 758 Dukes, Kristin Nicole and Sarah E Gaither. 2017. “Black racial stereotypes and victim blaming:
759 Implications for media coverage and criminal proceedings in cases of police violence against
760 racial and ethnic minorities.” *Journal of Social Issues* 73:789–807.
- 761 Entman, Robert M. 1992. “Blacks in the news: Television, modern racism and cultural change.”
762 *Journalism quarterly* 69:341–361.
- 763 Finch, Brian Karl, Kyla Thomas, Audrey N Beck, D Brian Burghart, David Klinger, and Richard R
764 Johnson. 2022. “Assessing data completeness, quality, and representativeness of justifiable
765 homicides in the FBI’s Supplementary Homicide Reports: A research note.” *Journal of*
766 *Quantitative Criminology* 38:267–293.

- 767 Goff, Phillip Atiba, Matthew Christian Jackson, Brooke Allison Lewis Di Leone, Carmen Marie
768 Culotta, and Natalie Ann DiTomasso. 2014. "The essence of innocence: consequences of
769 dehumanizing Black children." *Journal of personality and social psychology* 106:526.
- 770 Graham, Sandra and Brian S Lowery. 2004. "Priming unconscious racial stereotypes about
771 adolescent offenders." *Law and human behavior* 28:483–504.
- 772 Gruenewald, Jeff, Jesenia Pizarro, and Steven M Chermak. 2009. "Race, gender, and the
773 newsworthiness of homicide incidents." *Journal of criminal justice* 37:262–272.
- 774 Hargrove, Thomas. 2022. "Murder accountability project."
- 775 Humphries, Drew. 1981. "Serious crime, news coverage, and ideology: A content analysis of
776 crime coverage in a metropolitan paper." *Crime & Delinquency* 27:191–205.
- 777 Jensen, Jeffrey L., Daniel Karell, Cole Tanigawa-Lau, Nizar Habash, Mai Oudah, and Dhia
778 Fairus Shofia Fani. 2022. "Language Models in Sociological Research: An Application to
779 Classifying Large Administrative Data and Measuring Religiosity." *Sociological Methodology*
780 52:30–52.
- 781 Johnstone, John WC, Darnell F Hawkins, and Arthur Michener. 1994. "Homicide reporting in
782 Chicago dailies." *Journalism Quarterly* 71:860–872.
- 783 Lewis, Jerome A, James C Hamilton, and J Dean Elmore. 2021. "Describing the ideal victim: A
784 linguistic analysis of victim descriptions." *Current Psychology* 40:4324–4332.
- 785 Loftin, Colin, David McDowall, Karise Curtis, and Matthew D Fetzer. 2015. "The accuracy of
786 Supplementary Homicide Report rates for large US cities." *Homicide studies* 19:6–27.
- 787 Long, Lisa J. 2021. "The ideal victim: A critical race theory (CRT) approach." *International
788 Review of Victimology* 27:344–362.
- 789 Lowry, Dennis T, Tarn Ching Josephine Nio, and Dennis W Leitner. 2003. "Setting the public
790 fear agenda: A longitudinal analysis of network TV crime reporting, public perceptions of
791 crime, and FBI crime statistics." *Journal of communication* 53:61–73.
- 792 Lundman, Richard J. 2003. "The newsworthiness and selection bias in news about murder:
793 Comparative and relative effects of novelty and race and gender typifications on newspaper
794 coverage of homicide." In *Sociological Forum*, volume 18(3), pp. 357–386. Springer.
- 795 Meyers, Marian. 1996. *News coverage of violence against women: Engendering blame*. Sage
796 Publications.
- 797 Paulsen, Derek J. 2003. "Murder in black and white: The newspaper coverage of homicide in
798 Houston." *Homicide Studies* 7:289–317.
- 799 Pritchard, David. 1985. "Race, homicide and newspapers." *Journalism Quarterly* 62:500–507.
- 800 Regoeczi, Wendy C and John P Jarvis. 2021. "FBI Supplementary Homicide Reports." *The
801 Encyclopedia of Research Methods in Criminology and Criminal Justice* 1:43–46.

- 802 Rezapour, Rezvaneh. 2021. *From user-generated text to insight context-aware measurement of*
803 *social impacts and interactions using natural language processing*. Ph.D. thesis, University
804 of Illinois at Urbana-Champaign.
- 805 Rokaw, William M, James A Mercy, and Jack C Smith. 1990. "Comparing death certificate data
806 with FBI crime reporting statistics on US homicides." *Public Health Reports* 105:447.
- 807 Schildkraut, Jaclyn and Amy M Donley. 2012. "Murder in black: A media distortion analysis of
808 homicides in Baltimore in 2010." *Homicide Studies* 16:175–196.
- 809 Sorenson, Susan B, Julie G Manz, and Richard A Berk. 1998. "News media coverage and the
810 epidemiology of homicide." *American Journal of Public Health* 88:1510–1514.
- 811 Steffensmeier, Darrell, Jeffery Ulmer, and John Kramer. 1998. "The interaction of race, gender,
812 and age in criminal sentencing: The punishment cost of being young, black, and male."
813 *Criminology* 36:763–798.
- 814 Tans, Mary Dee and Steven H Chaffee. 1966. "Pretrial publicity and juror prejudice." *Journalism*
815 *Quarterly* 43:647–654.
- 816 Wang, Yilei, Jingyuan Tian, Deniz S Ones, Richard N Landers, et al. 2022. "Using natural
817 language processing and machine learning to replace human content coders." *Psychological*
818 *Methods* .
- 819 Weiss, Alexander and Steven M Chermak. 1998. "The news value of African-American victims:
820 An examination of the media's presentation of homicide." *Journal of Crime and Justice*
821 21:71–88.
- 822 White, Kailey, Forrest Stuart, and Shannon L Morrissey. 2021. "Whose lives matter? Race, space,
823 and the devaluation of homicide victims in minority communities." *Sociology of Race and*
824 *Ethnicity* 7:333–349.
- 825 Wilson, Michael and Erin O'Brien. 2016. "Constructing the ideal victim in the United States of
826 America's annual trafficking in persons reports." *Crime, Law and Social Change* 65:29–45.

827 **Figures**

Extract Summary of Victim	Modify Extract	Classify Article
<p><u>Initial Prompt:</u></p> <p>The following is a newspaper article that may include unrelated stories mixed together:</p> <p>< Original news article text ></p> <p>Create a text that only includes every piece of information in the article that refers to the victim <Victim name>, especially any facts about <Victim name>'s life, occupation, family, and relationship before the incident:</p> <p><u>GPT-3 Initial Output:</u></p> <div data-bbox="219 737 586 1083" style="background-color: black; color: white; padding: 5px;"> <p>James Santos, 28, of 17 Horan Way, Jamaica Plain, was killed early Sunday morning after an argument with 20-year-old James Abernathy of Roxbury. Santos was stabbed in the abdomen and taken to Carney Hospital, where he was pronounced dead at 1:58 a.m. His sister Nina Santos was treated for bruises at the same hospital and later swore out a complaint against Abernathy for assault and battery. Abernathy was taken to Charles Street Jail, where he was ordered held without bail. A hearing was set for Nov. 28 at Roxbury District Court.</p> </div>	<p><u>Second Prompt:</u></p> <p>The following text includes information about homicide victim <Victim Name>:</p> <p><Initial Output Text ></p> <p>Rewrite text after removing age, gender, address, and information about the suspect:</p> <p><u>GPT-3 Second Output:</u></p> <div data-bbox="625 737 992 1062" style="background-color: black; color: white; padding: 5px;"> <p>James Santos was killed early Sunday morning after an argument with another person. Santos was stabbed in the abdomen and taken to Carney Hospital, where he was pronounced dead at 1:58 a.m. His sister Nina Santos was treated for bruises at the same hospital and later swore out a complaint against the other person for assault and battery. The other person was taken to Charles Street Jail, where they were ordered held without bail. A hearing was set for Nov. 28 at Roxbury District Court.</p> </div>	<p><u>Third Prompt:</u></p> <p>Articles about homicide victims are humanizing when they focus on their lives, such as their occupations, relationships and families. Articles are impersonal when they focus only on what happened and the criminal's background. An expert editor is asked to evaluate this article based on these definitions and provide an opinion on whether the article is humanizing or impersonal.</p> <p>Summary: "Raymond Johnson was a man who lived with his companion, Carol Ross. On the day of his death, he was shot once in the chest and was found by police on the floor of his apartment."</p> <p>Response: Impersonal, because it only mentions what happened</p> <p>Summary: "George Healey was a homicide victim who had been boarding in the home for 15 years. The 100-pound Healey's favorite pastime was fishing in the nearby Connecticut River."</p> <p>Response: Humanizing, because it mentions the victim's pastime</p> <p>Summary: <Second Output Text ></p> <p>Response:</p> <p><u>GPT-3 Final Output:</u></p> <div data-bbox="1031 1052 1398 1104" style="background-color: black; color: white; padding: 5px;"> <p>Impersonal, because it only mentions what happened</p> </div>

Figure 1: GPT-3 prompt text and sample responses. Text in black background is actual response from GPT-3 for each of the three prompts. <Original news article text> is replaced with the full text of the news article. <Victim name> is replaced by the victim's full name. <Initial Output Text> in the second prompt is replaced by the full text response of the first prompt. <Second Output Text> in the third prompt is replaced by the full text response of the second prompt.

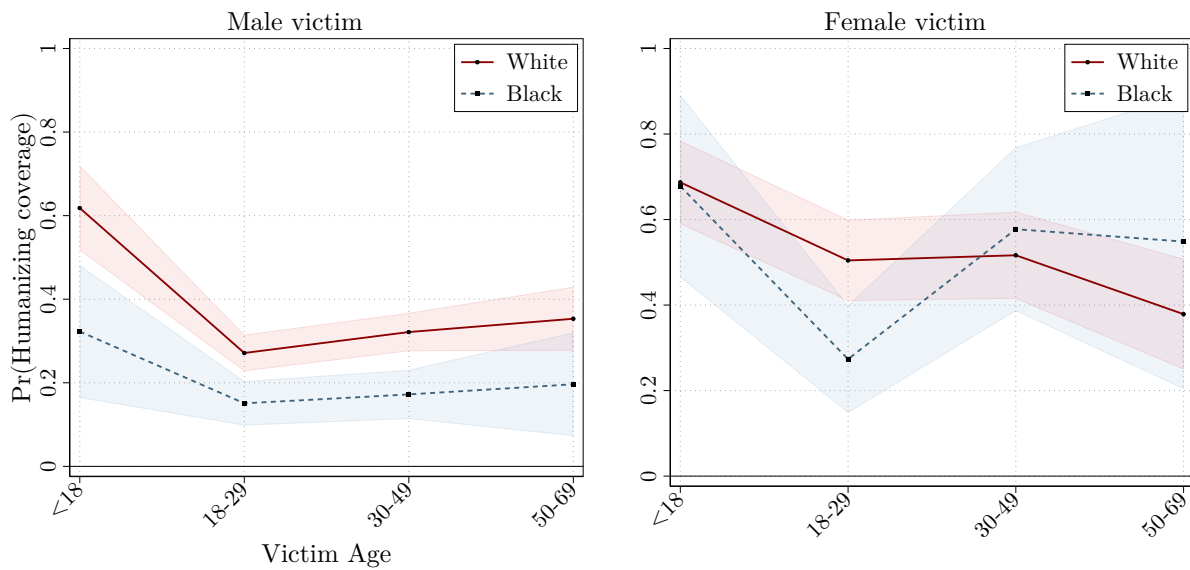


Figure 2: Predicted probabilities of humanizing coverage by race, age, and sex with 95% confidence intervals. Results are computed from a Probit regression model including fully interacted demographic indicators as well as weapon, circumstance, relationship, county, and time fixed effects.

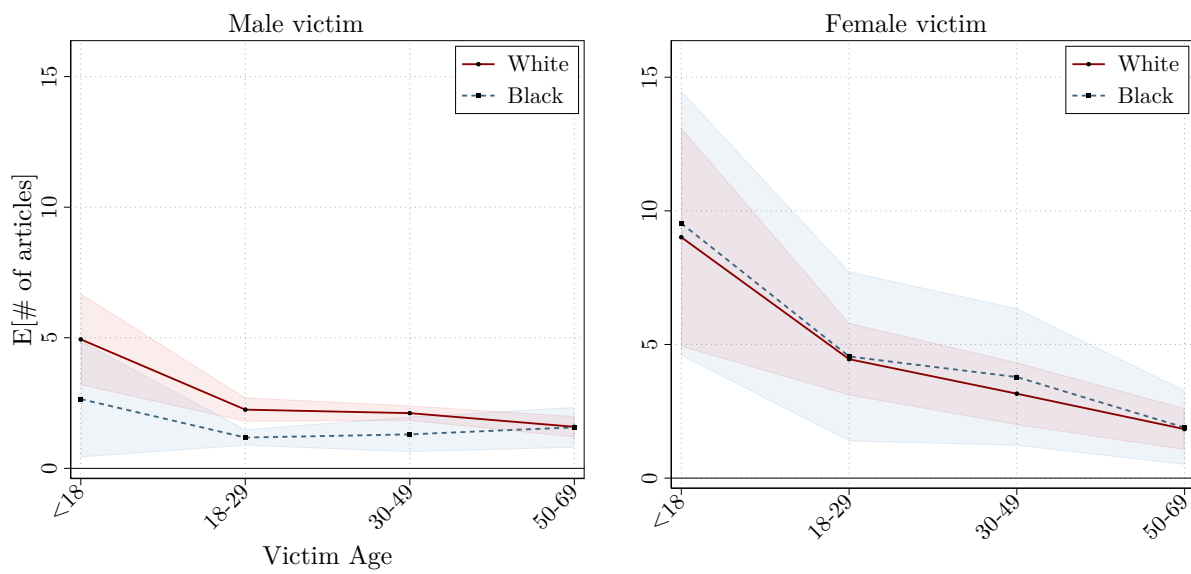


Figure 3: Predicted number of humanizing articles by race, age, and sex with 95% confidence intervals. Results are computed from a Poisson regression model including fully interacted demographic indicators as well as weapon, circumstance, relationship, county, and time fixed effects.

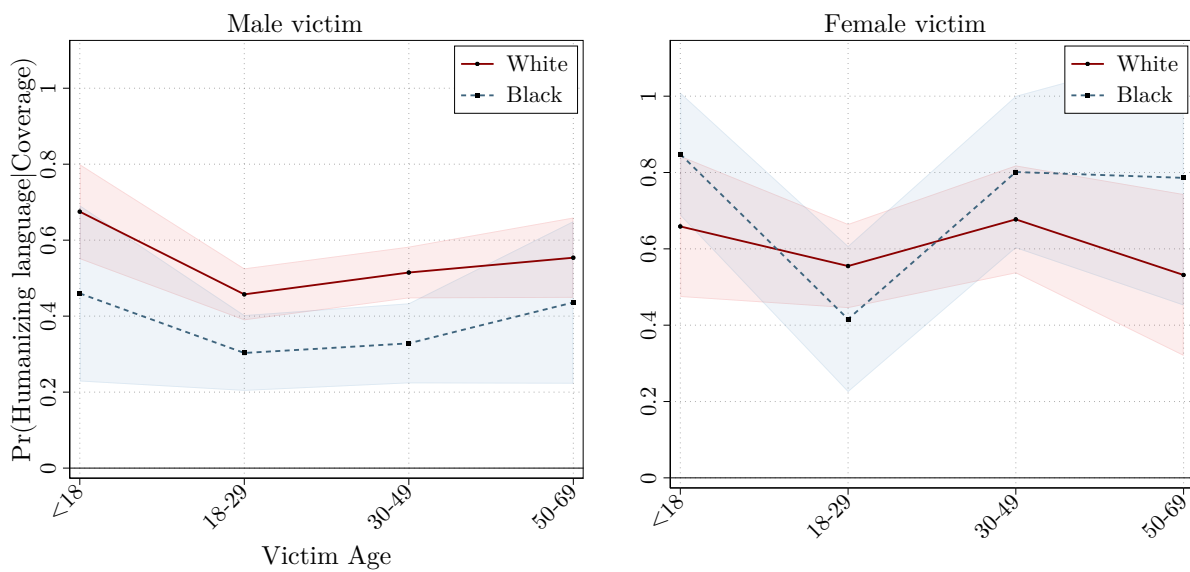


Figure 4: Predicted probability of at least one humanizing article, conditional on at least one article being written, by race, age, and sex with 95% confidence intervals. Results are computed from a Probit regression model including fully interacted demographic indicators as well as weapon, circumstance, relationship, county, and time fixed effects.

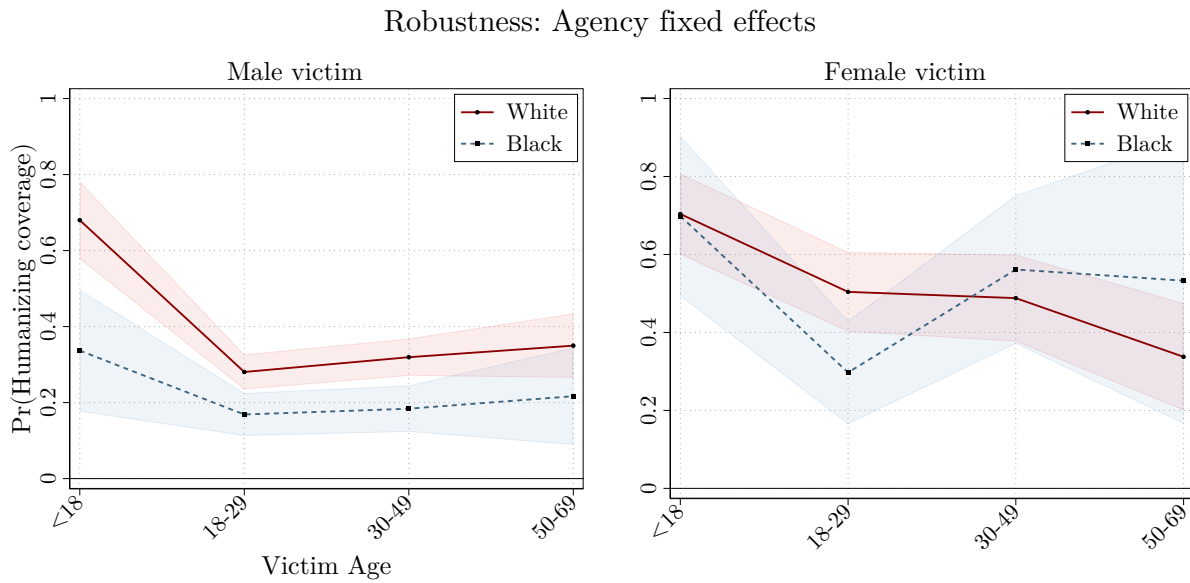


Figure A1: Baseline regression (composite humanizing score) replacing county fixed effects with agency fixed effects.

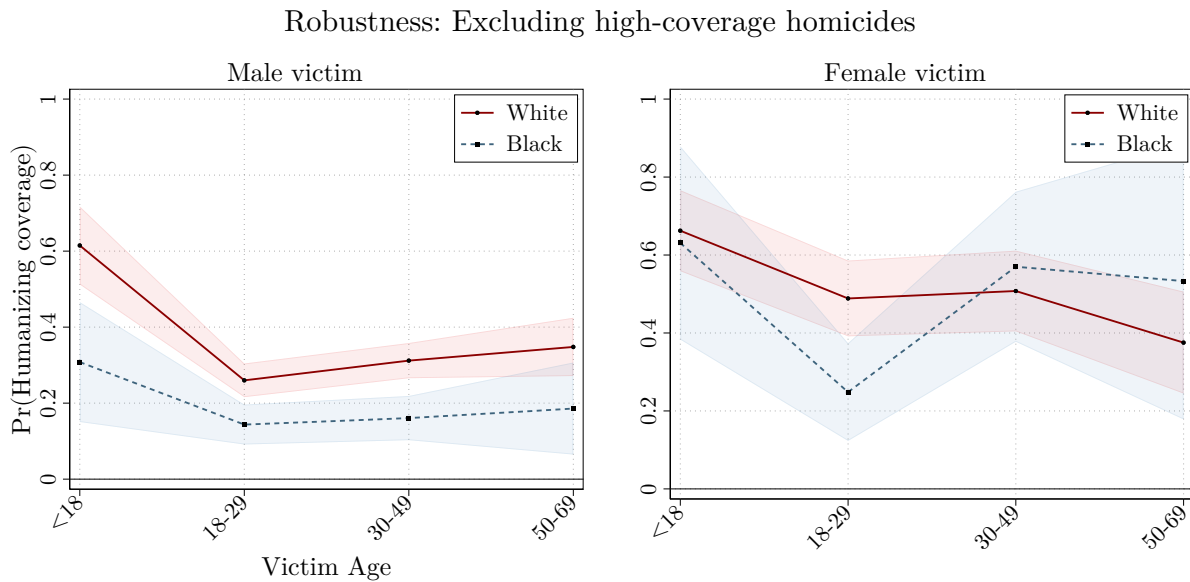


Figure A2: Baseline regression (composite humanizing score) excluding homicides about which more than 20 articles were written.

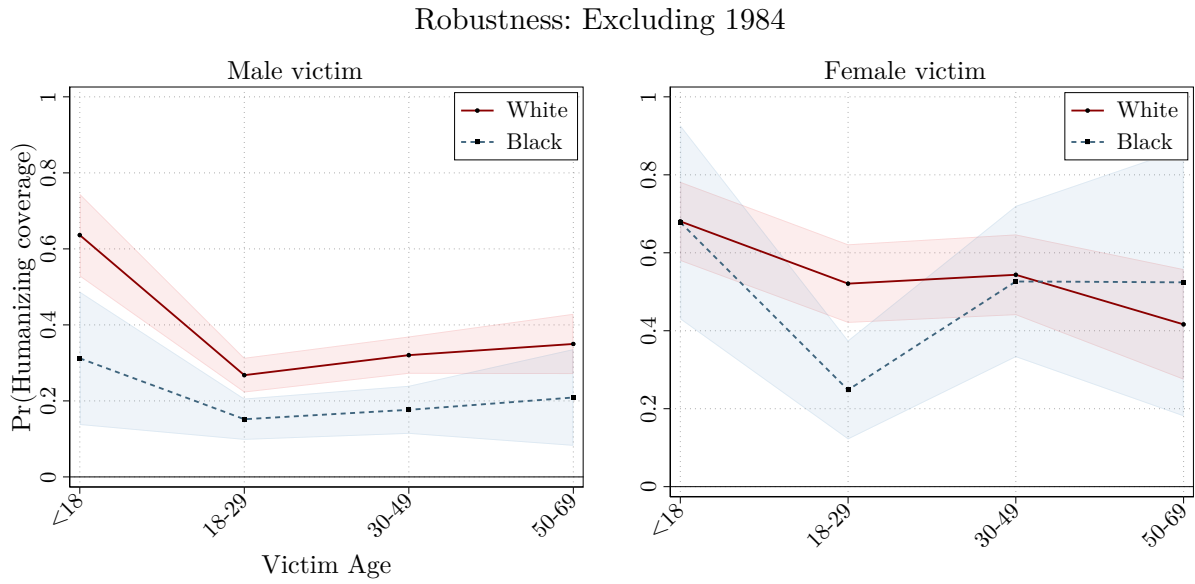


Figure A3: Baseline regression (composite humanizing score) excluding homicides that occurred during 1984.

828 **Tables**

Table 1: Summary statistics by victim age, sex, and race

	Male			Female		
	White	Black	Other	White	Black	Other
<18						
Pr(Humanized)	0.70	0.56		0.85	0.91	0.67
Articles/homicide	3.14	1.67		5.76	10.00	1.67
Homicides	84	33	0	68	15	3
18-30						
Pr(Humanized)	0.47	0.27	0.67	0.68	0.45	1.00
Articles/homicide	2.29	1.19	2.12	3.77	2.16	3.50
Homicides	339	168	8	114	44	2
30-49						
Pr(Humanized)	0.59	0.34	1.00	0.78	0.80	0.75
Articles/homicide	3.23	1.10	7.38	2.52	2.16	1.75
Homicides	337	126	8	79	25	4
50-69						
Pr(Humanized)	0.61	0.38	1.00	0.63	0.75	1.00
Articles/homicide	1.45	1.22	4.00	1.33	0.86	7.00
Homicides	143	36	2	64	7	1

Demographic Correlates of Humanizing Media Coverage of Homicide

Table 2: Determinants of Humanizing Coverage of Homicide

	<i>Lin. Prob. Model</i>		<i>Probit Model</i>	
	Dem.	Full	Dem.	Full
<u>Victim</u>				
Black	-0.18*** (0.03)	-0.15*** (0.03)	-0.53*** (0.10)	-0.53*** (0.11)
Female	0.13*** (0.03)	0.17*** (0.03)	0.35*** (0.07)	0.61*** (0.10)
<u>Offender</u>				
Black	0.08** (0.04)	0.04 (0.04)	0.23** (0.11)	0.16 (0.13)
Female	-0.04 (0.04)	0.03 (0.04)	-0.13 (0.13)	0.10 (0.16)
<u>Controls/FE</u>				
Unsolved	0.12 (0.15)	-0.01 (0.17)	0.44 (0.59)	-0.05 (0.61)
Mult. victims		0.25*** (0.05)		0.86*** (0.17)
Mult. offenders		0.04 (0.04)		0.17 (0.14)
Weapon (FE)		×		×
Circum. (FE)		×		×
Relation (FE)		×		×
County (FE)		×		×
Time (FE)		×		×
Observations	1682	1682	1680	1629
Adjusted R^2	0.033	0.169		

Robust standard errors in parentheses. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

FE indicates inclusion of fixed effects for the corresponding variable.

The sample consists of all homicides with a black or white victim under 70 years of age. Sample sizes differ across specifications because, as is well known, perfect predictors can lead to numerical problems in maximum likelihood estimation. As a result, these predictors and their associated observations are automatically dropped in the estimation procedure (as is the default behavior of most statistical packages).

Table 3: Testing significance of black-white differential in composite humanizing score

Δ (Black-White)	Male	Female
<18	-0.30*** (0.10)	-0.01 (0.12)
18-29	-0.12*** (0.04)	-0.23*** (0.08)
30-49	-0.15*** (0.04)	0.06 (0.11)
50-69	-0.16** (0.07)	0.17 (0.18)

Standard errors in parentheses.

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Table 4: Testing significance of black-white differential in news prominence score

Δ (Black-White)	Male	Female
<18	-2.29 (1.58)	0.52 (3.34)
18-29	-1.07*** (0.31)	0.10 (1.74)
30-49	-0.81** (0.38)	0.62 (1.50)
50-69	-0.02 (0.45)	0.06 (0.80)

Standard errors in parentheses.

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Table 5: Testing significance of black-white differential in conditional humanizing score

Δ (Black-White)	Male	Female
<18	-0.22 (0.13)	0.19 (0.12)
18-29	-0.15** (0.06)	-0.14 (0.11)
30-49	-0.19*** (0.07)	0.12 (0.13)
50-69	-0.12 (0.12)	0.25 (0.20)

Standard errors in parentheses.

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Table A1: Filtering algorithm article types

Accepted	Filtered out
Feature	Advertisement
Front Page/Cover Story	Classified Advertisement
Article	Front Matter
Commentary	General Information
Editorial	Birth Notice
News	Illustration
Legal Notice	Table of Contents
Marriage Announcement	
Stock quote	
Obituary	
Review	
Letter to the Editor	
Correspondence	
Editorial Cartoon/Comic	
Credit/Acknowledgement	
Image/Photograph	
Undefined	
Military/War News	

Table A2: Filtering algorithm word roots

Absolute	Relative	Death-related
slain	shot	fatal
slaying	shooting	dead
murder	wounded	dies
homicide	stabb	killed
manslaughter	kidnap	death
	strangl	died
	beat	
	bludgeon	
	asphyxiat	

Table A3: Filtering algorithm validation and test results

Sample	N	FP	FN	TP	TN	Specificity (%)	Sensitivity (%)
Initial	9416	112	0	31	9273	98.81	100
Validation 1	8658	126	0	89	8493	98.54	100
Validation 2	7200	129	0	25	7046	98.20	100
Validation 3	6235	95	0	16	6124	98.47	100
Final Test	6425	116	0	28	6281	98.19	100

N = Number of articles, FP = False positives, FN = True negatives, TP = True positives, TN = True negatives