

# Demographic correlates of humanizing media coverage of homicide: evidence from the *Boston Globe*, 1976-84

Emily Ocasio<sup>1</sup>, Tristan Potter<sup>2\*</sup>,

**1** Stanford University, Stanford, CA

**2** LeBow College of Business, Drexel University, Philadelphia, PA

\* tristan.l.potter@drexel.edu

## Abstract

We merge FBI data on Massachusetts homicides from 1976-1984 with the corresponding *Boston Globe* articles to study how the 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 measure the extent to which media coverage was humanizing through three conceptually distinct measures, two of which are derived from natural language processing analysis of article content. Our findings demonstrate patterns of differential coverage across different demographic groups: (i) Among male victims, whites were significantly more likely to receive humanizing coverage than blacks; (ii) there was no significant difference in coverage across races for female victims, except those aged 18-29, among whom blacks received less humanizing coverage; (iii) female and juvenile victims received more humanizing coverage; and (iv) race/sex of the offender does not correlate with coverage once we control for circumstantial factors of the homicide. We argue that the rich patterns of differential coverage that we identify across race, age, and sex in our sample are consistent with ideal victim theory. Our analysis highlights the importance of viewing differential media coverage of violent crime through an intersectional lens and how content analysis can be applied to produce comprehensive measures of humanizing language.

## Introduction

Crime stories make up 10-50% of all newspaper articles [1], with disproportionate attention devoted to stories about homicide in particular despite its relatively low prevalence [2]. Among homicides, the quality and quantity of coverage varies substantially, and this variation in coverage may plausibly relate systematically to the characteristics of the individuals involved in nuanced ways that have important social implications [1,3,4]. For example, depictions of crime in newspapers can provoke public anxieties and direct those anxieties towards certain groups [5,6], change jurors' decisions and trial outcomes [7], and affect attitudes of law enforcement officials towards both victims and perpetrators of crime [8,9]. Furthermore, in addition to influencing public perceptions and treatment of different groups, discrepancies in coverage of homicide also *reflect* the complex social forces from which they emanate.

We are not the first to study differential media coverage across demographic or socioeconomic groups. An important existing literature has sought to analyze media distortions by measuring, e.g., differences in news coverage of crime based on

characteristics of victims and offenders [10–15]. The studies that comprise this literature can be understood as varying principally along three dimensions: (i) How media coverage is measured, (ii) which demographic groups are considered in the analysis, and (iii) data availability, including which other homicide-specific factors the analysis is able to control for (and thus the plausibility of a causal interpretation), which homicides are chosen for analysis (and how they are matched with media coverage), and the geographic scope of the analysis.

Regarding the first dimension: Media coverage is traditionally measured through readily available quantitative metrics, such as whether or not any articles are written, the number of articles written, or an article’s “prominence,” (defined variously as the total number of words, column length, or where the article appears within the paper). However, several important recent papers have instead applied content analysis to understand the qualitative nature of homicide coverage [16], focusing on whether offenders and victims are described in positive or negative terms [17–20] or coding articles according to a strict standardized criterion, such as the number of close relations of the victim (e.g., family, friends, etc.) mentioned in the article [20]. Regarding the second dimension, demographic groups considered, the most common finding in this literature has been reduced coverage of black victims, usually without a distinction between sexes or age groups. Regarding the third dimension, data availability, existing studies usually attempt to match homicide case reports to specific news articles by searching for victim or offender names based on a single large metropolitan area and its corresponding major newspaper; what geographic, temporal, and homicide-specific controls are able to be included in the analysis varies significantly across studies. As we describe below, we build on this literature along each of these three dimensions.

A central contribution of our analysis concerns our approach to measuring coverage, and specifically how we determine whether a news article humanizes a victim. Our approach takes a comprehensive view of the textual content of the articles in our sample rather than focusing on a single dimension of humanizing language. We do this by applying natural language processing (NLP) techniques to automate the classification of each article as humanizing or impersonal with respect to a particular victim. NLP has shown utility in automating and standardizing content analysis [21–23]. Specifically, our analysis uses GPT-3, a large language model (LLM) that has been pre-trained with 300 billion tokens on 45 terabytes of data obtained from a variety of sources including books and public web pages [24]. A language model predicts the most likely next words to follow given a previous text, based on examples given during the model’s training phase. The value of language models in text classification lies in their ability to derive meaning out of precise textual context [25], thus affording a great deal of flexibility in determining whether or not language is humanizing. It provides a standardized and reproducible approach to classification that replaces the more time-consuming method of using human coders. Since the language model uses an algorithm that can be calibrated to remove any source of randomness, the analysis can be repeated with identical results, which may not be the case when human coders are involved. To limit the effect of such individual subjectivity, human coding approaches can apply very strict criteria (e.g., counting the number of references to family members). Such criteria limits the analysis to a narrow set of characteristics. The use of automated NLP classification allows for a more holistic analysis of the text, i.e., one that takes into account the language of the text in its entirety without restricting the criteria used, while at the same time allowing for fully reproducible results. In our use case, the automation of news article textual content analysis incurs additional complexity. We wish to analyze the specific language used to characterize homicide victims, and as such, the automation technique must recognize and extract the precise portions of the text that refer to the victim (and in the case of articles which include more than one homicide or more than

one victim, the language referring to one particular victim at a time). We selected GPT-3, a newer generation *generative* LLM, meaning that it is specifically designed and pre-trained to generate natural language as well as encode and analyze text input. The model can use its inherent knowledge of natural language to analyze the entire article and generate an extract of only the portions referring to a particular victim, then from the extract analyze the specific language used to describe the victim and determine whether it humanizes them.

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

Finally, our analysis is built around a new dataset that we construct from the FBI's Supplementary Homicide Reports (SHR) database and the universe of *Boston Globe* articles between 1976 and 1984. By using this specific time period, we are able to include nearly all homicides in an entire state, and to match those homicides to nearly all corresponding articles in a major regional newspaper. Furthermore, in addition to the demographic information described above (both for victim and offender), the FBI's SHR data includes detailed information on the circumstances of each homicide, as well as when and where it occurred. This allows us to control for a wide range of factors that might otherwise confound our analysis of how demographics relate to media coverage of crime. While we are able to control for many factors, we emphasize that our results are specific to the region and time period on which our analysis is focused.

## Data and methods

### Overview

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

### Homicide data and news articles

The universe of homicides available for our study includes all homicides reported anywhere in Massachusetts in the years 1976-1984. Rather than relying on police records from individual localities, we utilize data from the SHR [26], which captures homicide data from all local and state jurisdictions. The location and particular years

were selected because they represent the confluence of dates between when case-level information is published in the SHR (starting in 1976) and when the digitized textual transcripts of the article content of a major regional newspaper (*Boston Globe*) were available in bulk format for data mining (through 1984).

The SHR data was downloaded from the Murder Accountability Project website [27]. It includes individual entries for each homicide, with year and month of the homicide, victim demographics, offender demographics if there is a known offender, the police department that reported the case, and other information about the homicide such as the method of death (shooting, stabbing, strangulation, etc.), the offender relationship with the victim, the kind of weapon, and circumstances of the case (whether it was related to argument, robbery, drug-related, etc.). The SHR does not include personal information on the victims or offenders, such as names or addresses.

Our analysis includes each of the 1,800 Massachusetts homicides reported in the SHR as well as every newspaper article published in the *Boston Globe* for that time frame (1,215,056 articles, including non-news items such as advertisements and stock quotes). We obtain the text of newspaper articles from the ProQuest Historical Newspapers database, which includes the fully digitized text of every *Boston Globe* article published through 1984, available in bulk for analysis via Text Data Mining Studio. For our analysis, we extract the text of every article published in the morning edition for the years 1976 through 1984. These text transcripts were originally produced via optical character recognition from scanned paper copies of the newspapers, and as such they contained some errors such as misspelled words, occasional extraneous spaces within the words, uneven sentence formatting, and lack of paragraph separation. Despite these issues the article texts were readable and understandable. The article database also includes additional metadata about each article, including the date, title, and article type (e.g., news article, advertisement, editorial, stock quote).

## Identification of homicide articles

We develop and validate a filtering algorithm to automate the initial identification of which of the entire universe of articles published in the time period are likely related to Massachusetts homicides (MH), based on keyword/phrase matching of language related to homicides, Massachusetts locations and neighborhoods, and article types. The algorithm is calibrated to identify all possible MH candidates even though some of the articles initially identified may be false positives due to the use of similar language, requiring further manual filtering to finalize the MH set.

### Article filtering algorithm specification

Our goal for the algorithm is to ensure it does not miss any MH articles (false negatives) while reasonably minimizing the number of non-MH articles it erroneously marks as MH (false positives). The algorithm has three filters: an article must pass all three in order to be classified as an MH. Filter 1 requires that the article be one of the allowed types (types are determined within the ProQuest database). Standard types such as “article,” “feature,” “cover story,” and also “editorial” and “letters to the editor” are allowed for consideration as possible MH-related types. Types such as “front matter,” and “table of contents” are filtered out. Although we initially filtered out the “advertisement” type, we noticed during validation phases that a small number of actual news articles, including homicide-related articles, were erroneously typed as “advertisement” in ProQuest. We subsequently modified the rule to ensure the title itself contains a reference to “ad” before it is filtered out. Table 1 contains the full list of types that were allowed and filtered out.

**Table 1. 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	

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

**Table 2. 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	

Filter 3 ensures that the article refers to a Massachusetts location. We use a list of all cities, towns and counties in the state, as well as all the known Boston neighborhood

names. Articles with no mention of a location name anywhere in the text are filtered out. Note that there are three very small localities in the state whose names are very commonly used in a different context, namely “Lee,” “Florida,” and “Washington.” Accordingly, we remove these localities from the localities list because they produced a very large number of false positives during our validation, while there was not a single case of a death in one of those localities in which the only reference was to one of those three names. This filtering is implemented in Python using the *flashtext* library.

### Validation and testing of filtering algorithm accuracy

To validate the accuracy of the automated algorithm, we compared its performance with an independent human determination (after reading the full article) of whether it was MH-related or not. After initial selection of the algorithm’s parameters, we defined a subset of articles by randomly selecting a number of dates within our date range and including every article for those days. We applied the algorithm filters and simultaneously read each article and marked it as MH (positive) or non-MH (negative). The goal was to ensure no false negatives occur (algorithm determined the article was non-MH where human detected it as MH) while reasonably reducing the false positives (algorithm marked a non-MH article as potential MH). Since a human would subsequently examine all articles which the algorithm marks as MH, false positives would eventually be removed from consideration, while any potential false negatives would represent an article that actually referred to a MH but would have been missed from our analysis. After the initial round of comparison, we adjusted the algorithm to try to make it stricter while still not producing false negatives and then repeated validation using a different random article sample. We performed three rounds of adjustment and validation, and once satisfied with the accuracy, performed final testing on another random sample never used for algorithm adjustment or validation. Accuracy results for the initial, validation, and final testing subsets are shown in Table 3 based on the final algorithm logic. Ground truths for all subsets were determined before the algorithm was run. The final algorithm tests demonstrate a 100% sensitivity and 98.2% specificity.

**Table 3. 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 = False negatives, TP = True positives, TN = True negatives

### Post-algorithm manual article filtering

Since the automated filtering algorithm produces some false positives they must be removed by manual inspection by a human reader. Common examples of these false positives include a number of sports-related articles that use metaphors such as “this team murdered their opponents in the court last night,” as well as cases where the homicides did not occur in Massachusetts but instead in other places that share names with Massachusetts localities such as Cambridge, England. In addition, our analysis

specifically excludes vehicular homicides (though they may be considered homicides in a legal sense, they are not captured in the FBI database).

Once the filtering algorithm is finalized it is used to perform the first-pass filter of the full article database. Articles identified as possible MH by the algorithm are then manually inspected to reduce the group down to the final number of validated MH articles used in the analysis. The algorithm identified 22,640 of the 1.2 million articles as being potential MH; we subsequently read all of the 22,640 identified articles and manually filtered the potential set down to a final confirmed set of 5,042 MH articles.

### Matching of homicide reports with newspaper articles

Using the demographic and crime detail information in the articles, we then manually match each one of the confirmed MH articles to one or more MHs from the SHR. Since the names of victims and offenders are not included in the SHR, we use information in the news articles including date and location of homicide, victim and offender demographics, and other facts such as type of weapon, circumstances and victim-offender relationship, to identify each SHR homicide entry by victim name. Based on this we are able to determine whether each MH article matches to zero, one, or more SHR entries. The final multiplicity relationship between articles and homicides can be zero to many in both directions.

835 articles that refer to pre-1976 MH are removed from further consideration. Of the remaining articles, 988 are unable to have any homicide record assigned to them because they represent homicides that were clearly not in the SHR database, or the article is so vague that there is not enough information to assign it to a particular homicide. The final count of articles positively matched and assigned to one or more SHR entries is 3,219, with a total of 1,071 homicides matching to one or more articles.

### Humanizing coverage

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

GPT-3 is used to automate the determination of whether an article is humanizing or impersonal in its coverage of the homicide victim. Like other similar LLMs, GPT-3 is a pre-trained AI model and thus it does not require large training sets for classification. We use the few-shot prompt technique [24] in which, using natural English text, a small number of examples are provided as a way to “prompt” the model to follow the pattern.

## Prompt design

Using a validation set of articles for which the ground truth was determined by a human reader, we optimized the GPT-3 prompt design to improve its accuracy. Our prompt strategy involves the use of three consecutive prompts to the GPT-3 text completion interface (Figure 1). An initial prompt sends the full article text to GPT-3 such that the completion text returned is an extract of the article containing exclusively the information about the particular victim in question. This first prompt includes the full text of the article and is meant to simplify the text to specifically focus on the information about the particular victim in question. This step is necessary because many articles include substantial information not specifically related to the victim. This can occur in a variety of circumstances, such as: (a) the article refers to multiple homicide victims, (b) the article is an amalgamation of multiple news stories, only one of which refers to the homicide (e.g., the *Boston Globe* has a recurring news section known as “Massachusetts News in Brief” and the entire section is stored as a single article in ProQuest), (c) the article is not primarily about the homicide itself but refers to it secondarily. Without separating out this first prompt, we found that trying to classify the article directly led to many errors when the model tried to apply humanizing classifying logic to the entire article. The second prompt submits this extract back with a request to summarize it by removing demographic information such as age and race of the victim, as well as information about the offender and the circumstances of the crime itself. The summary returned as completion text is more precisely focused on personalized victim information. The third prompt sends the summary from the second prompt, requesting a final classification as humanizing vs. impersonal using the few-shot technique.

**Fig 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.

To illustrate in further detail how GPT-3 handles each prompt sequentially, below we provide the unredacted text from the ProQuest database for an article which is particularly complex since it covers multiple separate homicides:

oston th slaying in 3 days By  
Ben Bradlee Globe Staff  
20-year-old Roxbury man surrendered to Boston Pdlice yesterday  
morning and was charged with committIng the city s fifth  
murder in three days James Abernathy of 227  
Heath st turned himself into District 2 officers and was  
later arraigned at Roxsiyry District Court Abernathy was  
accused of killing James Santos 28, of 17 Horan way Jamaica  
Plain Police said witnesses told them that  
Abernathy and Santos got into an argument early Sunday  
morning after Abernathy allegedly kicked Santos’ sister  
Nina Santos as found stabbed in the abdomen at 141  
Intervale st Roxbury  
He was admitted to Carney Hospital where he was pronounced dead  
at 1:58 yesterday morning His sister Was treated at the  
hospital for bruises and released Poli’e said she later



swore out complaint against Abernothy for assault and battery  
Officers reported that Santos and Abernathy were acquaintances  
Abernathy was taken to Charles Street Jail here he was ordered held without bail hearing was set for Nov 28 at Roxbury District Court  
Detective Sgt John Maillet said the Santos killing was the 82d murder in Boston so far this year  
There Nere 71 murders in the city all of last year  
It's very unusual to have five killings in three days said Maillet  
For some reason don't know why seem to pick up during the holiday season  
Maybe there s more drinking There's lot of frustration any There were developments reported yesterday in three of the other four murders that have occurred in Boston since last Friday

James Baumgardner 23, of 40 Linsky Barry court Siltth Boston walked into District 6 and surrendered himself in connection with the shooting death early Saturday of 28-year-old James Keenan of Adams street Dorchester  
Baumgardner was charged with manslaughter  
Police said Keenan and friend were driving van aid had stopped to talk to some women outside Zimbee's bar at and West Fourth streets in South Boston when fight broke out between the two men and the women s escorts  
2 Police said Keenan and Lindberg drove off but short time later Police said Keenan was armed with shotgun

Agrouop of people subsequantly converged on the van and scuffle ensued between Keenan and Baumgardner shot rang out Keenan was later pronouced dead at the scone from single gunshot wound to the head

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

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

'Ramon Cartagena of Haverford street Jamaica Plain wiss ordered held without bail in Charles Street Jail cause hearing was set for Nov 28 at Roxbury Dlitric Court  
In the murder of Margaret Hammond the 69-year-old who was found dead Friday in her Brighton aptment police issued an arrest warrant for Roert McKay 57, who was said to have lived with Hamat 3 Lothian rd

McKay was still at large last night autopsy performed on Hammond showed she died massive skull fracture caused by blow from instrument  
Police said she was found wearing red

print dress covered with bedding She had been dead for  
several days before neighbors notified police after det  
tting decaying odor 'Tolice reported no  
arrests or new leads in connection the  
murder of 28-year-old cab driver Jeffery Bi of Swampscott  
Boyajian was robbed and killed last Friday apparently by  
three passengers he 6 to dead end street at the Archdale  
Housing in Roslindale

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

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 .

The language model clearly extracted the relevant information about the particular  
victim and removed much of the other extraneous information in the article. It also  
produced text written in correct English with no spelling or transcription errors. This  
extract can now be used in a subsequent prompt to determine whether the language is  
humanizing. Each prompt is separately sent to GPT-3 completion interface as individual  
requests with no memory of the prior prompts (as opposed to a chat-like approach in  
which the language model retains memory of the entire session). Although it was  
possible to construct a prompt that directly requested a humanizing determination for a  
particular victim from the initial raw article, we found that it resulted in more errors,  
particularly in complex articles like this one, where it picked up humanizing  
characteristics from portions of the article that were not directly related to the victim in  
question. In addition, the final humanizing prompt also attributed humanizing factors  
to basic demographic information such as the victim's age or humanizing details about  
the offender. Thus, we added an additional second prompt which slightly modified the  
extract specifically to remove age, gender, address, and information about the suspect.  
In this example the streamlined extract returned by the second prompt was as follows:

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 .

In this case, the victim's age and address were removed from the text, and the  
offender is now referred to as "another person." As a result, an article that originally

had much detail about many different things is now distilled down to only the facts (other than demographics) about the victim. In this case, it turns out that nothing other than the circumstances of death (and the fact that his sister was also a victim) are included, and so the response to the final prompt declares the article (correctly) as impersonal as it relates to this specific victim.

The third and final prompt uses the few-shot technique, where we provide several examples of how to respond. This is not considered training or fine-tuning—the NLP logic is embedded in GPT-3’s large pretrained model. What the few-shot technique accomplishes is guiding the model’s answer towards a specific structure, in this case a sentence starting with the word “Humanizing” or “Impersonal” followed by a short explanation of the reasoning behind the decision (we used this explanation in the process of tweaking the prompt design to infer the reasons for incorrect responses). The few-shot prompt is not “teaching” the model how to classify the article, but rather showing the model in what way we would like the answer to be formatted.

The prompts with the accompanying article or extract text were submitted via the GPT-3 text completion interface (<https://api.openai.com/v1/completions>) using the text-davinci-003 model with the following interface parameters: temperature = 0, max\_tokens = 256, frequency\_penalty = 0, presence\_penalty = 0. The completion interface utilizes the base capabilities of GPT-3 as a raw language model—the input to the interface is the prompt and the output is the model’s interpretation of the most likely set of words that would follow. GPT-3 and its successor versions underlie commonly known user applications such as ChatGPT, which have additional chat session capabilities that “remember” the contents of the chat or that search the internet for other sources of information beyond the works used for model training. We do not use these apps but instead rely exclusively on the base GPT-3 language model in completion mode. The temperature parameter (0) in particular removes any element of randomness, guaranteeing the same textual response after repeated identical prompts, and thus allowing for the results to be replicated by anyone using the same prompts and article texts. The text-davinci-003 model has an inherent limit of the length of the text it can analyze. This limit of 4,000 tokens (each token is a meaningful part of a word; usually 750 words would average out to 1,000 tokens) includes both the prompt and the response. There were three homicides in the total set for which one of their respective articles was so long that we could not evaluate via GPT-3. For these three homicides we determined their humanizing classification by human inspection.

To determine whether an article is humanizing, we submit the three successive prompts to GPT-3 and use the first word (*humanizing* or *impersonal*) of the final response as the result. For each homicide victim, we consecutively send each article on that homicide to GPT-3 as outlined above, until at least one article is determined to be humanizing or all articles are considered impersonal.

To optimize the accuracy of our prompt design we manually adjudicate a random subset of homicides with their respective articles as humanizing or impersonal ground truths and used this validation set to tweak the prompts and improve the accuracy. For the ground truth, a human reader, blind to the GPT-3 results, is given the original text of the articles from ProQuest and the name of the victim and asked to label each of the articles corresponding to our given criteria, where humanizing is where details of the victim other than demographics and circumstances of the crime are given, and impersonal otherwise. The ground truth for coverage of that homicide victim is considered to be humanizing if at least one of the articles was labeled as humanizing by the reader, and impersonal if all the articles are labeled impersonal. These ground truths were used to compare to the results from GPT-3 and prompt adjustments were then made for further optimization.

Once all three prompts were deemed to be optimized, we measure the prompt

design’s performance against a newly randomly selected final test set of homicides for which we also manually determine ground truths in a blind fashion. This final test of accuracy of GPT-3’s humanizing determination against a human reader’s ground truth labeling consists of 30 randomly selected homicides spanning 100 newspaper articles, with GPT-3 agreeing with the human reader for 28 of the homicides (93.3% accuracy).

Using the prompt strategy outlined here, we then determine the nature of the coverage for 1,071 homicides having at least one related news article, and we find that 602 homicides (56%) are classified as humanizing. This classification is used for our statistical analysis.

## Results

Table 4 reports summary statistics for the combined dataset. Specifically, within victim age-sex-race cells, the table reports (i) the share of victims described with humanizing language ( $\text{Pr}(\text{Humanized})$ ), (ii) the average number of articles per homicide (Articles/homicide), and (iii) the number of homicides (Homicides). Throughout 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.

**Table 4. Summary statistics by victim age, sex, and race**

	Male			Female		
	White	Black	Other	White	Black	Other
<b>&lt;18</b>						
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
<b>18-30</b>						
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
<b>30-49</b>						
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
<b>50-69</b>						
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

The statistics provided in the table illustrate the rich heterogeneity in coverage when sub-populations are broken out not just by race, but also by age and sex. For example, black male children are more than twice as likely to receive humanizing coverage than black male young adults (0.56 vs. 0.27), whereas this gap is substantially smaller for white males (0.70 vs 0.47); black males of all ages are substantially less likely to receive humanizing coverage than black females, whereas these gaps are typically smaller among white homicide victims. This rich heterogeneity in news coverage will bear out in our regression analysis and, as we will argue, can be used to inform theory.

The goal of our empirical analysis is to use the data described above to understand how demographic characteristics relate to the extent to which media coverage conveys homicide victims in a humanizing light. Specifically, the detailed homicide data from the SHR enables us to study rich *interactions* among demographic factors. It also

allows us to control for demographic-correlated factors related to the nature of the homicide that would otherwise be likely to confound an analysis that only included demographic factors. The humanizing score that we construct with GPT-3 enables us to study how the interaction of demographic factors relate to humanizing coverage.

The analysis that follows first describes our empirical strategy and then presents three sets of results, each corresponding to a different measure of humanizing language: Our main results are based on a *composite humanizing score*, which measures the use of humanizing language in the full sample of homicides, including those receiving no coverage (an absence of coverage is deemed to be impersonal) and those receiving coverage, which we categorize as humanizing or impersonal as described above. Next, we consider a *news prominence score*, which measures the quantity of coverage—i.e., the number of articles written about a homicide, irrespective of the nature of the language used in those articles. Finally, we consider a *conditional humanizing score*, which instead measures the quality of coverage—i.e., the likelihood of humanizing language, conditional on a homicide receiving coverage. Because we study all three measures, our analysis does not hinge on a particular approach to measuring media coverage and accounts for the fact that many homicides receive no coverage at all.

## Empirical strategy

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

1. Victim demographic controls: Race, sex, and age;
2. Offender demographic controls: Race, sex, and age [We 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];
3. Homicide controls: Weapon, relationship between victim/offender, specific circumstances of the homicide (e.g., gang-related), whether there were multiple victims, and whether there were multiple offenders;
4. Time (month-by-year) and county fixed effects.

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

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

where  $\beta$  is the vector of coefficients we seek to estimate and  $f$  is the link function mapping  $X_v$  to the conditional mean of  $h_v$ .

## Composite humanizing score

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

$$h_v^{\text{comp}} = \begin{cases} 0 & \text{No humanizing articles (incl. no articles)} \\ 1 & \text{At least one humanizing article} \end{cases} \quad (2)$$

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

We begin with a simple analysis in which we only consider a subset of demographic controls—victim/offender race and sex—thus neglecting age as well as potentially important interactions among these variables, to which we return below. Table 5 reports the results from this analysis. The first two columns correspond to estimation of a simple linear probability model ( $f^{-1}(X'_v\beta) = X'_v\beta$ ) and the second two columns correspond to estimation of a Probit model ( $f^{-1}(X'_v\beta) = \Phi(X'_v\beta)$ , where  $\Phi$  is the cumulative distribution function of the standard normal distribution). We estimate a linear probability model to facilitate interpretation of parameters and a Probit model to be consistent with the more in-depth interacted analysis that follows (results here and throughout the paper are not materially affected if we instead use a Logit model). Within each functional form specification, we report results when the model is estimated only with demographic controls as well as with the full set of controls listed above.

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

While the analysis in Table 5 sheds important light on some demographic determinants of humanizing coverage of homicide, it fails to account for the rich ways in which intersectional demographic factors *interact* to shape that coverage. Indeed, including age and demographic interactions in the regression in Table 5 increases the adjusted- $R^2$  by more than 15%. Accordingly, we re-estimate the fourth specification in Table 5, augmented with fully-interacted indicators for race, age, and sex. To facilitate exposition of the results, Figure 2 plots the predicted probabilities of humanizing coverage disaggregated by race, age, and sex. Specifically, the left panel plots the predicted probability of humanizing coverage for white and black male victims by age, and the right panel does the same for female victims. Predictive margins are computed by integrating over all covariates distributed as observed in the data. Table 6 reports formal tests of equality of the predicted probabilities for black and white victims within each age-sex cell. 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.

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 5. 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

**Table 5. 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).

**Fig 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.

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 6, 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. Note that overlapping confidence intervals in Figure 2 need not indicate a statistically insignificant difference—for example, in the case of 50-69 year old males the 95% confidence intervals overlap but the difference is significant at the 5% level. Second, the same pattern does not hold for females: in general, there is not a statistically significant

581  
582  
583  
584  
585  
586  
587  
588

**Table 6.** Testing significance of black-white differential in composite humanizing score

$\Delta$ (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.

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

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.

We include in the Supporting Information results for various robustness tests of our main specification in Figure 2. Specifically, Figure S1 Fig replaces our relatively coarse county fixed effects with much more granular FBI agency fixed effects, thus providing a better control for geographic variation (we do not report the estimates of the fixed effects themselves because there are so many geographic regions); Figure S2 Fig trims our sample to exclude outlying homicides—those about which more than 20 articles were written—which could potentially be driving our results (results are similar with different cutoffs); and Figure S3 Fig 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.

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

## News prominence score

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

$$h_v^{\text{prom}} = \# \text{ of articles} \quad (3)$$

and estimate a Poisson regression ( $f^{-1}(X_v'\beta) = \exp(X_v'\beta)$ ) where the dependent variable is now  $h_v^{\text{prom}}$ . The set of explanatory variables  $X_v$  is identical to those used in the analysis in Figure 2. Figure 3 summarizes the predicted number of articles written,



broken down again by race, sex, and age, and Table 7 reports formal tests of equality of the predicted number of articles for black and white victims within age-sex cells.

**Fig 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.

**Table 7.** Testing significance of black-white differential in news prominence score

$\Delta$ (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.

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

With this alternative measure of humanizing coverage, we observe some similarities as well as some important differences relative to the results described above. Among the important similarities are systematically more coverage of white victims than black victims among males, virtually no racial discrepancy at any age among females, and the tendency for young black males (18-29) to receive less coverage than all other groups. However, when the dependent variable is the number of articles, the gap between males and females is substantially larger than previously. Moreover, there is no gap between white and black coverage of females at any age group, whereas the difference between white and black males is only robustly significant among young (18-29) and middle-aged (30-49) males (second and third rows of Table 7). The gap between black and white male children is once again the most pronounced, but the difference in this case is not significant (first row of Table 7).

### Conditional humanizing score

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

$$h_v^{\text{cond}} = \begin{cases} 0 & \text{No humanizing articles} \mid \geq 1 \text{ article} \\ 1 & \text{At least one humanizing article} \mid \geq 1 \text{ article} \end{cases} \quad (4)$$

and estimate a Probit model ( $f^{-1}(X'_v\beta) = \Phi(X'_v\beta)$ ) where now the dependent variable is  $h_v^{\text{cond}}$  and we thus restrict attention to only those homicides about which at least one article was written. Furthermore, in addition to the standard set of controls used in Figures 2 and 3, we also control for the total number of articles written about a homicide, which will tend to be mechanically correlated with our dependent variable due to our approach to coding humanizing articles. Figure 4 reports the results from

this analysis (analogously to Figures 2 and 3) and Table 8 reports formal tests of equality of the predicted probabilities for black and white victims within age-sex cells (analogously to Tables 6 and 7).

**Fig 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.

**Table 8.** Testing significance of black-white differential in conditional humanizing score

$\Delta$ (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.

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

The results in Figure 4 and Table 8 add yet another layer of nuance to our understanding of the factors that relate to humanizing coverage. Specifically, we now observe patterns that are qualitatively quite similar to the patterns that we observed in Figure 2: A marked gap between humanizing coverage of black and white male victims with young males (18-29) receiving less coverage than other age groups and no universal difference between the language used to describe black and white female victims. Notwithstanding this second observation, and also similar to Figure 2, we see a pronounced racial gap among young (18-29) females, although in this case the gap is not statistically significant whereas it was in Figure 2. On the other hand, the gap between males and females of all ages is somewhat diminished relative to what we see in Figure 2, and the pervasive black-white gap among males that was significant for all age groups in Figure 2 is now only robustly significant among young (18-29) and middle-aged (30-49) males, similar to the results in Figure 3 and Table 7.

## Discussion

Our analysis—both the inclusion of rich demographic interactions and our use of different measures of humanizing coverage—allows us to discriminate between theories of media coverage in a way that would not otherwise be possible. We first argue that our results are not consistent with several prominent theories of media coverage, and then argue that they are consistent with ideal victim theory.

First, our results are not consistent with the theory that newsworthiness is mostly explained by an event’s rarity. Although the observed greater degree of humanizing coverage of juvenile victims can be understood as a natural consequence of the relative infrequency of juvenile homicides, this explanation does not account for what we observe when comparing juvenile homicide coverage of young males by race: white male juveniles are much more likely to be humanized than black male juveniles despite the fact that

there were almost three times as many white juvenile victims as there were black juvenile victims. Second, we think it is unlikely that our results are driven by certain communities not sharing humanizing details following a homicide because of a lack of trust in reporters. If this were the case, we should expect to see significant differences across demographic groups in only the portions of our analyses using our NLP-based measure of humanizing language, but not in our news prominence score (which reflects the number of articles written rather than the provision of humanizing details in an article). Moreover, it is unclear why such a theory would predict, e.g., that there is not a consistent difference in coverage across black and white females whereas there is a pronounced difference across black and white males—a lack of trust in journalists may plausibly exist in racially defined communities but why this lack of trust would differentially affect male versus female victims is less clear. Finally, the possibility that reporters at the *Boston Globe* were simply unaware of certain homicides that occurred disproportionately among certain groups is also unlikely to explain the data. To the extent that differential availability of information about homicides is geographically determined, this will be (at least partially) picked up in our geographic fixed effects.

Instead, any theory consistent with our results must account for the fact that observed differences across demographic groups are brought into sharp contrast when measured intersectionally. Across age groups, we observe significantly less humanizing coverage for black male victims as compared to white males, while white female victims overall are not more likely to be humanized than black female victims. However, when we account for differential effects across age groups, significant differences are seen between certain groups: black females between 18 and 29 are significantly less likely to be humanized than white females of the same age. Moreover, while black males are humanized less across all age groups, the gap is largest for juveniles.

These findings are consistent with ideal victim theory: young people are seen as more vulnerable, while being a black victim counteracts this by being seen as aggressive or strong. While age, sex, and race all influence societal perceptions of innocence and vulnerability, they act in different directions across groups. Our results can be understood as reflecting identified stereotypes towards black females and males; younger black females, for example, are masculinized and seen as strong, falling into the “Jezebel” or “strong black woman” trope, while older black females instead fall into the “mammy” stereotype, which is associated with caring and motherhood. Similarly, for males, the largest difference in the humanizing probability is between black and white juveniles. This is the most vulnerable and innocent age group, so the observed difference is consistent with a lack of humanizing perceptions of black boys who are instead depicted as older, stronger, and less innocent [28].

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

The results also demonstrate the importance of a methodological approach to media distortion analysis that combines the presence/extent of coverage (quantitative measures) with determination of humanizing language (qualitative measures). When we use the composite humanizing score, we detect significant differences between the humanizing probabilities of black and white females between 18 and 29, whereas in the news prominence score regression, which instead considers article count, we fail to

detect this difference. The conditional humanizing score regression, which considers the humanizing categorization only among the subset of homicides about which at least one article is written, predicts a small difference which is not statistically significant. In addition, our use of GPT-3, a generative LLM, to perform content analysis from articles with varied and complex text, shows the viability and scalability of this automation technique.

## Limitations

One potential limitation of our study is inherent in the lack of a source of crime reports that reliably includes 100% of the homicides in the analyzed location and time period. We thoroughly search all newspaper articles for MH-related stories independent of the crime reports, and subsequently find a number of cases not included in the SHR. This gap has not usually been identified or measured in media distortion analysis since the articles are usually found from a search by victim and/or offender name based on the available crime reports. Nevertheless, studies of SHR accuracy have validated their reliability when compared with aggregate data from the National Vital Statistics System, with the exception of cities with very low numbers of black homicides [29,30]. Specifically, some infrequent kinds of homicides, in particular justifiable homicides (e.g., law enforcement killings of felons during commission of a crime) show very low rates of inclusion in SHR [31].

A second limitation relates to the use of automated content analysis. The use of automated content analysis via GPT-3 allows for more scalable and consistently reproducible determination of humanizing article language, especially when analyzing large numbers of articles. Although it eliminates human variance and coding error, it may introduce biases arising from the language model itself. This potential for bias has been recognized by the developers of GPT-3 [24]. When GPT-3 was provided racial cues during prompting, the developers noticed the resulting text differed in sentiment based on the prompted race. This was most likely due to the inherent discrepancy in depictions of race among the millions of source articles used to train the model. Our use of GPT-3 does not explicitly provide race in the prompts and also removes whatever references to demographic characteristics may be in the article during the second prompt (Figure 1), so the final determination of humanization is as race- and age-agnostic as possible. Gender is harder to remove completely from the prompts because of the associated pronouns and the presence of gender-specific names (this may also be an issue with race to a much smaller extent). In addition to making the prompts neutral, we also do not ask GPT-3 to generate new language, only to summarize and extract information from existing articles, as well as to make a final determination. Future studies should explore the effects of artificially adding, removing, or changing racial depictions in prompts on GPT-3 humanization classification accuracy.

An additional limitation relates to our data set, and in particular the fact that our sample size is too small to allow us to meaningfully identify how interactions between the demographic characteristics of the victim and those of the offender shape media coverage. Future research employing more temporally and geographically comprehensive datasets should consider the role of such interactions in accounting for news coverage of crime.

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

relationships we find in the data. In terms of external validity, our study focuses on a specific region and time period—Massachusetts in the 1970s and 1980s—as well as a particular news source with its own political leanings—the *Boston Globe*—all of which is relevant context for our results. These particular choices were driven by the requirement that the entire universe of published articles from a source which covers crimes statewide be considered in the analysis so that even brief or secondary mentions of each homicide are accounted for. This could only be done using newspaper articles that have been comprehensively organized and converted to textual form, from a time period during which the SHR database also exists. As discussed above, while the fact that the patterns of disparate coverage that we document are consistent with various identified stereotypes is suggestive of the broad relevance of our results, different social forces operate in different time periods and regions, and thus there remains a pressing need to study these questions in different and more modern contexts. Our study provides a road map for such future research, which we hope will apply our methodology to study disparate media coverage in different localities and time periods as their text becomes available in full for data mining.

## Conclusion

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

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

## Supporting information

**S1 Fig.** Baseline regression (composite humanizing score) replacing county fixed effects with agency fixed effects.

**S2 Fig.** Baseline regression (composite humanizing score) excluding homicides about which more than 20 articles were written.

**S3 Fig.** Baseline regression (composite humanizing score) excluding homicides that occurred during 1984.

## Acknowledgments

825

Thanks to seminar participants at Drexel University for helpful comments and suggestions.

826

827

## References

1. Chermak SM. *Victims in the news: Crime and the American news media*. Westview Press Boulder, CO; 1995.
2. Paulsen DJ. Murder in black and white: The newspaper coverage of homicide in Houston. *Homicide Studies*. 2003;7(3):289–317.
3. Meyers M. *News coverage of violence against women: Engendering blame*. Sage Publications; 1996.
4. Clayman SE, Reisner A. Gatekeeping in action: Editorial conferences and assessments of newsworthiness. *American Sociological Review*. 1998; p. 178–199.
5. Lowry DT, Nio TCJ, Leitner DW. Setting the public fear agenda: A longitudinal analysis of network TV crime reporting, public perceptions of crime, and FBI crime statistics. *Journal of communication*. 2003;53(1):61–73.
6. Boulahanis JG, Heltsley MJ. Perceived fears: The reporting patterns of juvenile homicide in Chicago newspapers. *Criminal Justice Policy Review*. 2004;15(2):132–160.
7. Tans MD, Chaffee SH. Pretrial publicity and juror prejudice. *Journalism Quarterly*. 1966;43(4):647–654.
8. Graham S, Lowery BS. Priming unconscious racial stereotypes about adolescent offenders. *Law and human behavior*. 2004;28(5):483–504.
9. Dukes KN, Gaither SE. Black racial stereotypes and victim blaming: Implications for media coverage and criminal proceedings in cases of police violence against racial and ethnic minorities. *Journal of Social Issues*. 2017;73(4):789–807.
10. Johnstone JW, Hawkins DF, Michener A. Homicide reporting in Chicago dailies. *Journalism Quarterly*. 1994;71(4):860–872.
11. Weiss A, Chermak SM. The news value of African-American victims: An examination of the media's presentation of homicide. *Journal of Crime and Justice*. 1998;21(2):71–88.
12. Sorenson SB, Manz JG, Berk RA. News media coverage and the epidemiology of homicide. *American Journal of Public Health*. 1998;88(10):1510–1514.
13. Lundman RJ. The newsworthiness and selection bias in news about murder: Comparative and relative effects of novelty and race and gender typifications on newspaper coverage of homicide. In: *Sociological Forum*. vol. 18(3). Springer; 2003. p. 357–386.
14. Gruenewald J, Pizarro J, Chermak SM. Race, gender, and the newsworthiness of homicide incidents. *Journal of criminal justice*. 2009;37(3):262–272.
15. Schildkraut J, Donley AM. Murder in black: A media distortion analysis of homicides in Baltimore in 2010. *Homicide Studies*. 2012;16(2):175–196.

16. Humphries D. Serious crime, news coverage, and ideology: A content analysis of crime coverage in a metropolitan paper. *Crime & Delinquency*. 1981;27(2):191–205.
17. Barlow MH, Barlow DE, Chiricos TG. Economic conditions and ideologies of crime in the media: A content analysis of crime news. *Crime & Delinquency*. 1995;41(1):3–19.
18. Pritchard D. Race, homicide and newspapers. *Journalism Quarterly*. 1985;62(3):500–507.
19. Entman RM. Blacks in the news: Television, modern racism and cultural change. *Journalism quarterly*. 1992;69(2):341–361.
20. White K, Stuart F, Morrissey SL. Whose lives matter? Race, space, and the devaluation of homicide victims in minority communities. *Sociology of Race and Ethnicity*. 2021;7(3):333–349.
21. Wang Y, Tian J, Ones DS, Landers RN, et al. Using natural language processing and machine learning to replace human content coders. *Psychological Methods*. 2022;.
22. Brown CH. Analyzing group behavior from language use with natural language processing and experimental methods: Three applications in political science and sociology [PhD thesis]. The University of Texas at Austin; 2018.
23. Rezapour R. From user-generated text to insight context-aware measurement of social impacts and interactions using natural language processing [PhD thesis]. University of Illinois at Urbana-Champaign; 2021.
24. Brown T, Mann B, Ryder N, Subbiah M, Kaplan JD, Dhariwal P, et al. Language models are few-shot learners. *Advances in neural information processing systems*. 2020;33:1877–1901.
25. Jensen JL, Karell D, Tanigawa-Lau C, Habash N, Oudah M, Fani DFS. Language Models in Sociological Research: An Application to Classifying Large Administrative Data and Measuring Religiosity. *Sociological Methodology*. 2022;52(1):30–52. doi:10.1177/00811750211053370.
26. Regoeczi WC, Jarvis JP. FBI Supplementary Homicide Reports. *The Encyclopedia of Research Methods in Criminology and Criminal Justice*. 2021;1:43–46.
27. Hargrove T. Murder accountability project; 2022. Available from: <https://www.murderdata.org/>.
28. Goff PA, Jackson MC, Di Leone BAL, Culotta CM, DiTomasso NA. The essence of innocence: consequences of dehumanizing Black children. *Journal of personality and social psychology*. 2014;106(4):526.
29. Loftin C, McDowall D, Curtis K, Fetzer MD. The accuracy of Supplementary Homicide Report rates for large US cities. *Homicide studies*. 2015;19(1):6–27.
30. Rokaw WM, Mercy JA, Smith JC. Comparing death certificate data with FBI crime reporting statistics on US homicides. *Public Health Reports*. 1990;105(5):447.

31. Finch BK, Thomas K, Beck AN, Burghart DB, Klinger D, Johnson RR. Assessing data completeness, quality, and representativeness of justifiable homicides in the FBI's Supplementary Homicide Reports: A research note. *Journal of Quantitative Criminology*. 2022;38(1):267–293.