

NBER WORKING PAPER SERIES

OFFICER-INVOLVED:
THE MEDIA LANGUAGE OF POLICE KILLINGS

Jonathan Moreno-Medina (r)

Aurelie Ouss (r)

Patrick Bayer (r)

Bocar A. Ba (r)

Working Paper 30209
<http://www.nber.org/papers/w30209>




NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2022

We are grateful to Mackenzie Alston, Julia Cagé, Tony Cheng, Kareem Haggag, Anjelica Hendricks, Justin Holz, Andrea Kiss, Dean Knox, Rachel Mariman, Arnaud Philippe, John Rappaport, Elliot Ash, Roman Rivera, Martin Salzmänn and Jennifer Tamas for helpful comments. We also want to thank seminar and workshop participants at Stanford/Hoover, NYC Media Seminar, Text-as-Data Workshop, ASSA, APPAM, ALEA, SEA, ETH Zurich, Sciences Po, Brock University and UTSA. We thank Arianna Ornaghi for sharing data on the dates of Sinclair acquisitions of local television stations. We thank Meghna Baskar, Sanjit Beriwal, Jasmine Carter, Sofia Casado, Jack Duhon, Alicia Gong, JoonYup Park, Caroline Milgram, Corey Rudman, Melissa Sandoval, Adam Soliman, Audrey Wang, and Alex Whitefield for providing excellent research assistance. The author order has been randomized and recorded on the AEA Author Randomization Tool, with confirmation ID “LPFDU1546yLy.” The randomization of order is indicated by the symbol (r). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2022 by Jonathan Moreno-Medina (r) Aurelie Ouss (r) Patrick Bayer (r) Bocar A. Ba. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Officer-Involved: The Media Language of Police Killings

Jonathan Moreno-Medina  Aurelie Ouss  Patrick Bayer  Bocar A. Ba


NBER Working Paper No. 30209


July 2022


JEL No. K14,K42,L82


ABSTRACT

This paper examines language patterns in US television news coverage of police killings. We first document that the media employ semantic structures that obfuscate responsibility—such as passive voice, nominalizations, and intransitive verbs—more frequently for police killings than for civilian killings. Using field variation and an online experiment, we demonstrate that these language differences matter. In the field, we find that people who happened to have taken a survey just after more obfuscatory coverage of a police killing are more likely to support police funding. In our online experiment, participants are less likely to hold police officers morally responsible and demand penalties when exposed to obfuscatory language, especially when the victim is unarmed. Returning to the news data, we find higher use of obfuscatory language when victims are unarmed, when video footage is available, or when the suspect is not fleeing—in other words, situations when obfuscation matters most. Turning to the causes of this differential obfuscation, our evidence is not consistent with either demand-side drivers or supply-side factors associated with TV station ownership and political leaning. Instead, our results point to original narratives crafted by police departments as a more likely driver of obfuscation. Our study emphasizes the importance of considering semantic language structures in understanding how media shapes perceptions, extending beyond coverage quantity and slant.

Jonathan Moreno-Medina 
Department of Economics
University of Texas at San Antonio
Business Building One UTSA Circle
San Antonio, TX 78249
jonathan.moreno-medina@utsa.edu

Aurelie Ouss 
Department of Criminology
University of Pennsylvania
571 McNeil Building
3718 Locust Walk
Philadelphia, PA 19104
and NBER
aouss@upenn.edu

Patrick Bayer 
Department of Economics
Duke University
213 Social Sciences
Durham, NC 27708
and NBER
patrick.bayer@duke.edu

Bocar A. Ba 
Department of Economics
Duke University
213 Social Sciences
Durham, NC 27708
and NBER
bocar.ba@duke.edu

As officers contacted the suspect an OIS [officer-involved shooting] occurred, one of the officer's rounds penetrated a wall that was behind the suspect. Beyond that wall was a dressing room. Officers searched the dressing room and found a 14 year old female victim who was struck by gunfire.

Tweet from Los Angeles Police Department Media Relations following the police killings of Valentina Orellana-Peralta and Daniel Elena Lopez, 2021

1 Introduction

News outlets have many choices in whether and how to describe any event. These choices could affect how people understand and imagine what happened, their perceptions of causality and moral responsibility, and, ultimately, their broader beliefs and judgments about the world around them (Pinker, 2007). Given its central role in society, many dimensions of media coverage have received considerable academic attention, including which events are covered in the first place (Eisensee & Strömberg, 2007; Enikolopov *et al.*, 2011) and what words are used to describe them—e.g., whether language is politically slanted, biased or gendered (Gentzkow & Shapiro, 2010; Chiang & Knight, 2011; Martin & Yurukoglu, 2017; Jakiela & Ozier, 2018; Gay *et al.*, 2018).

In this paper, we examine another critical aspect of media language—semantics—which is broadly concerned with how the structure of language affects its understood meaning. Specifically, we study the use of particular sentence structures, such as the active versus passive voice or the inclusion versus omission of a subject, which systematically work to either clarify or obfuscate the actor and/or actions taken during an incident. We do this in the context of media coverage of police killings of civilians. Over 1,000 individuals are killed by police annually in the US, with these deaths accounting for approximately four percent of all homicides.¹ Media watchdogs have called attention to the tendency of news reports and police department press releases to describe police killings using language structures specifically designed to diminish the central, active role of police officers in the killings. Journalist Radley Balko has coined the term “exonerative tense” for these language structures to highlight their apparent aim of dampening negative judgments about the appropriateness of the officers’ actions.²

Our paper proceeds in several steps. First, using data on the universe of American television news broadcasts (both local and national) from 2013 to 2019, we examine whether there is greater use of obfuscatory language structures in coverage of police killings versus reporting on homicides in general. We then use a field survey and an online experiment to causally test whether obfuscatory language matters for how people understand a news story about a police killing, how they assign agency and responsibility, and how they perceive the issue of police reform. We return to the news data to examine whether obfuscatory language is used more frequently by the media in circumstances when our experiment suggests it would have the greatest impact on a viewer’s perception. Finally, we offer a conceptual framework to analyze the underlying factors contributing to obfuscation. Our findings suggest that obfuscation primarily stems from the news media relying on official police reports in their narrative construction.

¹<https://www.cdc.gov/nchs/fastats/homicide.htm>

²See, for example, Balko (2014) and Blachor (2020) for related discussions in the popular press.

To characterize how language structure can clarify or obfuscate actions, actors, and the relationships between them, we draw heavily on the linguistics literature (Toolan, 2013; Pinker, 2007). We capture four dimensions of potential obfuscation of meaning relative to that of a sentence with an active-voice verb of which a police officer is the subject and the victim is the object (“police officer killed man”). The first is the use of the *passive voice* (‘man was killed by police officer’), which pushes the role of the police officer to the background of the sentence, potentially decreasing its salience to the reader. The second is a further transformation of the sentence to remove any reference to the police as the cause of the killing (‘man was killed’). We refer to this structure as: *no agent*. A third obfuscatory structure is the use of *nominalization*, which involves transforming the action of the police killing into a noun, leaving important information of the event ambiguous (‘deadly officer-involved shooting’). The final dimension we consider is the use of the intransitive – e.g., transforming the transitive verb ‘kill’, which requires an agent who generates the action, into the *intransitive* ‘die’, which does not require a cause (‘man dies [in shooting]’).³

Our primary dataset combines closed-caption texts (i.e., verbatim transcripts of audio) from the universe of American television news broadcasts from 2013 to 2019 (covering national and local stations) with data on the universes of (i) police killings of civilians drawn from the Mapping Police Violence (MPV) database and (ii) civilian gun homicides drawn from the Gun Violence Archive (GVA). We capture the four dimensions of obfuscation described above using recently developed natural language processing (NLP) algorithms for coreferencing (i.e., finding all expressions that refer to the same actors) and semantic role labeling tasks (Lee et al., 2018; Shi & Lin, 2019).⁴ A characterization of media reliance on obfuscation in absolute levels would be difficult to interpret, so as a benchmark, we compare obfuscation in coverage of police shootings with that in coverage of civilian homicide shootings. We restrict the comparison to homicides in which a suspect is identified somewhere in the story to ensure that the media’s potential language choices are comparable.⁵ We conduct a number of robustness checks, such as dropping sentences in which the suspect’s name appears, to ensure that our results hold beyond our main sample selection criteria or specification choices. Our sample includes a total of approximately 6,000 police killings, 8,000 nonpolice homicides, 200,000 stories, and 470,000 sentences that describe the killings.

Our main findings are as follows: We find higher obfuscation in the descriptions of police killings than in those of other homicides across all four dimensions. Overall, there is some obfuscation in 35.6 percent of stories about police killings, compared to 28.8 percent for civilian homicides, a 25 percent increase. The estimated effects are even greater when we include media market or station fixed effects and, importantly, when we restrict attention to the first sentence of the story. In first sentences, obfuscation is 40 percent greater for coverage of police killings than for reporting on civilian homicides, suggesting that the media employ more obfuscation in the most salient part of the story.

³To be more precise, the definitions of obfuscation works as follows: *passive voice*, by making the victim the subject of the sentence, deemphasizes the agent of the killing by moving it into the predicate; *no agent* is a syntax which entirely elides any reference to the police as agent of the killing; *nominalization*, whereby a verb (“to kill”) is transformed into a freestanding noun (“killing”), with no grammatical need for completion with a subject or object, potentially leaving both the agent and receiver of the action ambiguous (e.g., “deadly officer-involved shooting,” where “officer” might be understood to be either the agent or the receiver of the action); *intransitive*: cases in which the verb requires no grammatical receiver of the action, which can obscure any relationship between subject and object that might otherwise be implied by a transitive verb (compare the intransitive “to die,” whose subject is the victim but the agent of whose death is omitted, with the transitive “to kill,” which must be completed by an object that receives the subject’s action (“officer killed man”).

⁴We explain the essence of these tasks in Section 3.3 and describe our implementation in further detail in Appendix A.2 and A.3.

⁵As we explain in more detail below, if the shooter is not known (as is often the case in the immediate aftermath of an incident), it is natural for the media to use the passive voice with or without elision of the causal agent (“a victim was killed [by an unknown person]”) versus the active (“an unknown person killed a man”) in reporting the event.

We then test whether obfuscatory language changes how people understand and process the information in a news story. Although such an effect has been hypothesized in the linguistics literature (Toolan, 2013), there is scant experimental or empirical evidence of its operation to date. We proceed in two steps, first exploiting field variation and then an online experiment. For field variation, we use a large, nationally representative survey, the Cooperative Election Survey, to examine how stated support for police funding varies with the language used in recent media coverage of police killings, conditional on a police killing having recently occurred. Given its scale, the survey was administered over several days, and we exploit variation in local media coverage around the exact date that the respondent took the survey. We find that more obfuscation is correlated with more support for local police funding, providing motivating evidence from the field that obfuscation affects attitudes about policing.

We conducted the online experiment on Prolific in March 2022 with 2,402 participants. The experiment explores how sentence structures affect responses to a news story about a police killing. We focus on three main outcomes: judgment about the police officer’s moral responsibility in the incident, demand for penalties (departmental and legal) for the officer, and financial support for an organization supporting police reform. We evaluate how the outcomes vary with the degree of obfuscation: a passive-voice verb, a passive-voice verb with no agent (using “officer-involved” instead), and an intransitive verb (again with “officer-involved”), compared to an active-voice structure. We test two main hypotheses: i) that obfuscatory sentence structures decrease the perceived moral responsibility of the police, the demand for penalties, and support for police reform, and ii) that the degree of obfuscation (with an intransitive verb being the most and a passive-voice transitive verb the least obfuscatory) is reflected in the order of the effect sizes. We also vary each treatment arm by whether the person killed by the police is described as holding a weapon.

In line with our pre-registered hypotheses, we find strong evidence that obfuscatory language matters and that, directionally, the effect sizes increase in the degree of obfuscation. As hypothesized, we also find that the effects are smaller when the story specifies that the victim was armed.⁶ When no weapon is reported to have been present, our findings imply that obfuscatory reporting decreases the responsibility assigned to and desired level of accountability for the officer and increases the extent to which the officer is judged to have been justified in committing the shooting. When a weapon is reported to have been present, we find smaller, but still statistically significant, effects of sentence structures that do not explicitly identify the police officer as the grammatical subject (*no agent* and *intransitive*) but statistically insignificant effects for the use of obfuscatory language overall, as the effect of *passive voice* is weaker in this case. Overall, the experimental results indicate strong effects of syntactical structures in news coverage on participants’ judgments about the police officer’s actions in the specific incident described in the reporting. We also find that obfuscating the direct role of the police in the killing reduces donations for reform by a modest amount.

With the experimental results in hand, we return to our news data to examine whether the media are more likely to use obfuscatory language in incidents when language might matter more for perceptions. Media obfuscation is indeed more prevalent for police killings in which the victim was unarmed—that is, precisely when our experimental results imply that such language works to soften judgments about the moral responsibility of the police officer for the killing. We also find more obfuscation in police killings for which body-worn camera footage is available, when the suspect was not fleeing, or when the incident went viral—again suggesting that such language is used in cases where viewers are potentially more likely to form

⁶As expected, we also find evidence that the reported presence of a weapon matters directly, decreasing responsibility and the desire for accountability for the officer.

harsher judgments against the police.

We close the paper by providing a conceptual framework exploring the drivers of differential obfuscation. Following [Gentzkow *et al.* \(2015\)](#), we consider both demand- and supply-side drivers. Among the supply-side mechanisms, we distinguish the role of the desire to shape how readers understand the news on the part of either the press, which we proxy by conservative media ownership and by slant in the news; or, local police departments as primary sources of information for media coverage. We provide empirical analyses that examine each of these channels. We find little heterogeneity in the identified effects by local voting patterns, suggesting a small role of demand in obfuscatory reporting, and no heterogeneity by TV station ownership or slant. However, we find that how police departments describe events matters. For six police departments, we collect officer-involved shooting (OIS) reports, which are statements released by departments after an officer discharges a firearm. We compute the degree of obfuscation in these statements, match them to related events to the coverage in our TV data, and compare the degree of obfuscation in both sources. We find a positive correlation in both sources of obfuscation even after we include police department fixed effects, suggesting that obfuscatory language likely originates from media reliance on police narratives.

Our paper relates to several strands of prior literature. First, our work contributes to a growing literature that employs NLP and computational linguistics to analyze text data: for example, financial reports, newspaper articles, press releases, opinion pieces, social media comments, or congressional transcripts (see [Gentzkow *et al.* 2019](#)). These approaches have been fruitful for our understanding of how the tone of a text or speech impacts political outcomes ([Grimmer, 2010](#); [Gentzkow & Shapiro, 2010](#)), firm performance, and firm exposure to political, social, and climate risks ([Baker *et al.*, 2016](#); [Hassan *et al.*, 2019](#); [Engle *et al.*, 2020](#); [Giglio *et al.*, 2021](#)). These studies used NLP methods to capture political slant, company executives' views, and market participants' sentiments. Thematically related, a few recent working papers study word choices in arrest reports ([Abdul-Razzak & Hallberg, 2022](#); [Campbell & Redpath, 2022](#)). In economics, prior work has also considered the role of narratives from both theoretical ([Bénabou *et al.*, 2020](#)) and empirical ([Widmer *et al.*, 2022](#); [Ash *et al.*, 2023](#)) angles. Prior research has, for example, documented how grammatical structures affect saving behaviors ([Chen, 2013](#)) or how phrasing choices in media reporting impact perceptions of immigration ([Djourelouva, 2023](#)). Our paper adds to this literature by exploring explicitly a key aspect of language—semantics—in a new context: the use of obfuscatory language in news media.

Our paper also contributes to the literature documenting the impact of media on a variety of economically and socially relevant outcomes.⁷ Research shows that the content and presentation of news affects health choices ([Bursztyn *et al.*, 2022b](#)), financial markets ([Baker *et al.*, 2016](#); [Engle *et al.*, 2020](#)), and attitudes toward immigration ([Gentzkow & Shapiro, 2004](#); [DellaVigna *et al.*, 2014](#); [Djourelouva, 2023](#)) and that prospective media coverage influences politicians' actions ([Durante & Zhuravskaya, 2018](#)). We add to this literature by documenting differences in the media's semantic choices in coverage of police killings and analyzing how these choices affect perceptions of these incidents. Closer to our specific research question, past research shows that news also influences people's voting behaviors ([DellaVigna & Kaplan, 2007](#); [Chiang & Knight, 2011](#); [Cagé, 2020](#); [Couttenier *et al.*, 2021](#)). Our field evidence and online experiment shows that information and syntactical structures influence people's perception of events, which in turn could affect political stances and support for policies ([Bursztyn *et al.*, 2022a](#); [Alesina *et al.*, 2021](#); [Andre *et al.*, 2021](#)).

⁷For brevity, we focus on the effects of news coverage but acknowledge that there is a long literature on other types of media, including entertainment TV ([Kearney & Levine, 2015](#)), movies ([Dahl & DellaVigna, 2009](#)), and educational programming ([Gentzkow & Shapiro, 2008](#)). See [DellaVigna & La Ferrara \(2015\)](#) for a fuller review of the literature.

Moreover, several papers relate news coverage or political speeches to perceptions of crime and police or jury behaviors (Mastrorocco & Minale, 2018; Mastrorocco & Ornaghi, 2020; Philippe & Ouss, 2018; Grosjean *et al.*, 2022; Ash & Poyker, 2023).

Research outside economics has also investigated how police departments' activities are covered in the news, mainly through the coverage of crime. For example, in early studies, Gilliam & Iyengar (2000) and Gilliam *et al.* (1996) show overreporting of crime when the suspect is Black, while later studies do not find this to be the case (Dixon & Williams, 2015). Duxbury *et al.* (2018) establish that media are more likely to emphasize mental illness if the perpetrator of a mass shooting is White than if he is Black or Hispanic. Grunwald *et al.* (2022) find that police agencies' Facebook posts overrepresent Black suspects relative to their proportion among local arrestees. A few papers have also explored variation in media coverage of crime by perpetrator or victim gender (Frazer & Miller, 2009; Henley *et al.*, 1995; Yasmin, 2021) and documented the high use of passive-voice structures in coverage of sexual assault (Bohner, 2001; Lussos & Fernandez, 2018).

Last, our paper builds on findings in cognitive science and linguistics. Closest to our work is the research showing that semantic choices affect perceptions of the moral responsibility of perpetrators (De Freitas *et al.*, 2017) and victims (Henley *et al.*, 1995; Niemi & Young, 2016; Northcutt Bohmert *et al.*, 2019). Our experimental results provide new evidence along these lines, demonstrating that the use of obfuscatory language decreases the assignment of moral responsibility and the desired level of accountability for police officers who kill civilians.

2 Linguistics Framework

In this section, we describe the linguistics framework that forms the basis of our empirical analysis. Our primary goal is to identify particular sentence structures that work to either clarify or obfuscate the agents and actions described in a news story, thereby affecting the viewer's (reader's) perception of what happened and who was responsible. Our framework draws heavily on Chapter 8 of Toolan (2013), which includes a detailed characterization of how different narrative structures affect perceptions of causal relationships and the assignment of causal agency.⁸

In psycholinguistics, the term *causative construction* refers to how language is used to depict causation from one subject (the *causal agent*) to another (the *causal patient*). The baseline construction for comparison throughout our paper is a sentence structure that clearly identifies the action with an active-voice verb, the causal agent as its subject, and the causal patient as its object—i.e., sentences of the form “A police officer killed a man.” Following Toolan (2013), we focus on four key syntaxes that can be used to obscure the action or roles of the causal agent and patient: (i) passive-voice rather than active-voice verbs, (ii) nominalizations, i.e., nouns created from verbs, (iii) the failure to identify a causal agent, and (iv) intransitive rather than transitive verbs. We present each of these sentence structures in turn below. Table 1 provides simple examples (columns 1 and 3) and examples drawn directly from our news broadcast data (columns 2 and 4) to help illustrate how these sentence structures are deployed in the contexts of killings by police (columns 1 and 2)

⁸Toolan (2013) also discusses other narrative forms that can modify perceptions of causal relations, such as direct commentary and evaluation or editorial choices on how to name things (for example, choosing between the terms *rioter* or *demonstrator*). While also of potential interest regarding the media coverage of police killings, characterizing these aspects of media language is beyond the scope of this research because it is difficult to identify such strategies at scale and to compare their usage across different kinds of news stories, as these word choices are likely to be domain specific.

and killings by civilians (columns 3 and 4).

Passive versus active voice. The sentence “A police officer killed a man” uses a transitive verb in active voice and identifies the causal agent as subject, with no nominalization. A first way to diminish the responsibility of the causal agent for the action is to instead use the passive voice: “A man was killed by a police officer.” With the active voice, the subject acts upon its object through the act described by the verb. This sentence structure is considered to be strong, direct, and clear in tone. It also orients the causal agent in the place of rhetorical emphasis at the start of the sentence. In contrast, the passive voice turns the causal patient (the victim) into the verb’s subject and relegates the causal agent (police officer) to the sentence’s predicate, a position of lower salience to the reader.

Although the function of the passive is to deemphasize the agent, there is evidence that its use changes the perception of the reader or viewer. As noted by [Chestnut & Markman \(2018\)](#), “[S]tating ‘The woman was abused by the man’ rather than ‘The man abused the woman’ causes people to be more accepting of violence against women, because the passive voice distances perpetrators from their crimes and consequently makes the crimes seem less severe ([Henley et al., 1995](#)).” Recent work in linguistics and cognitive science suggests that the passive voice increases psychological distance with respect to the narrated event by making it seem more distant in time and space and more hypothetical ([Chan & Maglio, 2020](#)). Furthermore, as we will see below, the passive voice also makes it easier to omit a subject altogether, abstracting from the role of the causal agent to an even greater extent.⁹

Nominalization. Nominalization is the process of transforming an adjective or verb into a noun. It is a key linguistic resource in everyday language, as it allows one to refer to an event without fully narrating it. In news reporting, it helps shorten stories, but can also be a tool to obfuscate agency since it abstracts from the relationship between the person being killed and the person doing the killing and thus leaves some aspects of the narrative ambiguous. In the context of police killings, a common nominalization is the term “officer-involved shooting,” which can stand in for the less ambiguous “a police officer killed a civilian.” There are two things to note in this case. First, even though the participation of an officer is noted with this nominalization, the officer’s causal role in the shooting is not specified. Second, although the police officer might have killed someone (as in our data), the chosen verb for the nominalization is not “to kill” but “to shoot.”¹⁰ Thus, it also leaves ambiguous the fact that someone was killed. In the context of civilian killings, phrases such as “intimate partner killing” or “gang-related shooting” can also be used.

Failure to identify the causal agent. A third way to diminish the ability of a reader to assign responsibility for an action is to remove the causal agent from a passive-voice sentence altogether—e.g., “A man was killed following a police chase” instead of “A man was killed by a police officer.” In this case, the person responsible for the killing is not referenced at all. This construction contrasts with those in which the causal agent is the subject of a sentence (preceding an active-voice verb) or is explicitly mentioned after the causative preposition “by” (after a passive-voice verb), where a direct connection between the agent and the action is made.

⁹For a typological and functional overview of the passive, see [Kazenin \(2001\)](#).

¹⁰The comparable phrase “officer-involved killing” is not commonly used by the news media.

Table 1: Obfuscatory Sentence Structures – Simple Examples

	Police killing		Civilian killing	
Semantic structure	Simple example	News broadcast	Simple example	News broadcast
Active voice	A police officer killed a person	Investigators believe the police officer shot and killed the man just before nine o'clock	A suspect killed a person	They believe [name of suspect] stood over the boys as they slept on the couch and shot them
Passive voice	A person was killed by a police officer	New developments, a California man under arrest tonight accused of making a prank call that led to a victim being shot to death by the police	A person was killed by a suspect	The Goodhue County attorney says that the man was shot in the chest by the suspect early yesterday morning
Nominalization	A person was killed in an officer-involved shooting	An officer-involved shooting late Thursday night claimed the life of a Monroe man	A person was killed in a domestic-violence shooting	A 30-year-old man was shot and killed on Tuesday in a gang-related shooting
No agent	A person was killed	Officials say the 45-year old man was shot after he refused to drop a knife	A person was killed	Several shots were fired at the doorway into the apartment with several adults, a toddler and an infant inside
Intransitive	A person died	The man has died after a shootout with police officers in St. Louis	A person died	[Name of victim] died at the scene of that shooting

Intransitive versus transitive verbs. A fourth way to obfuscate causation and responsibility is to use an intransitive verb, which takes no object: “A man died following an incident on the North Side.” Wolff (2003) and Pinker (2007) categorize transitive verbs such as “to kill” and “to shoot” as causative verbs because they implicitly relay the idea that the causal agent in the sentence *acted upon* the receiver of the action directly, intentionally, and without an intervening actor. By contrast, intransitive verbs such as “to die” do not require a causal agent or even imply causation at all. Instead, intransitive constructions feature only the causal patient (as subject), here the deceased person. Thus, the use of intransitive verbs not only obfuscates responsibility for an event but implicitly abstracts from causality, directly increasing the ambiguity about what happened in the first place.¹¹

Putting it all together. For simplicity, we use the labels *Passive*, *Nominalization*, *No agent*, and *Intransitive* to refer to these four forms of obfuscatory sentence structure throughout the rest of the paper. Importantly, these sentence structures are not mutually exclusive and are often combined in practice. We use the term *Any obfuscation* to refer to the use of any of these structures and define a sentence as having *No explicit agent* if *No agent* or *Intransitive* applies. Finally, as the progression in Table 1 suggests, we use the order in which we introduced these sentence structures - *Active* >> *Passive* >> *Nominalization* >> *No agent* >> *Intransitive*—to define a hierarchy of causal clarity or, in reverse, a hierarchy of obfuscation.

Prior research in linguistics and moral philosophy in particular has demonstrated that variation in these sentence structures influences how readers interpret an event. For example, in early work, Trew (1979) argues that news writing uses narrative structures that reflect dominant social beliefs. Wolff (2003) and Pinker (2007), among other scholars, emphasize how perceptions of causation can be influenced by choices in sentence construction. Closely related to our work, De Freitas *et al.* (2017) show that there is a close relation between choices of causative verbs and the subsequent moral judgment of viewers/readers.

3 Data, Sample Construction, and Language Processing

3.1 Primary Data Sources

For our main analyses, we draw data from several sources: a comprehensive dataset on the universe of police killings in the US between 2013 and 2019; a database that includes the (near) universe of gun-related killings in the US from 2014 to 2018, and the closed captions (text transcriptions) of all televised news broadcasts on both local and national stations from 2013 to 2019.¹²

Police Killings. There is no official government record of police killings in the United States. As a result, in recent years, journalists, activists, and researchers have undertaken independent efforts to build a comprehensive database of all such killings. For our analysis, we use data from *Mapping Police Violence* (MPV). The MPV research group identifies and documents all police killings occurring in the US since 2013. The incidents are identified from other crowdsourced databases on police killings in the United States, including FatalEncounters.org. MPV processes each potential case and improves the quality and completeness of the

¹¹Note that you cannot say “The police officer *died* the victim” or “The victim was *died* by the police.” Pinker (2007) refers to this as the intransitive “resisting” a causative.

¹²As mentioned below, audits of our police killings data show coverage of 98.3 percent of all cases.

data by examining available information about the case from traditional and social media and obituaries. [Conner et al. \(2019\)](#) find that the MPV covers 98.3 percent of all police killings in 2015. The MPV dataset includes information on the victim’s name and race, the police department responsible for the killing, and the address and zip code where the incident took place; as well as some contextual features (whether the victim was allegedly armed, whether there was body-worn camera footage...).

Civilian Gun Homicides. There is also no official government database of civilian homicides in the United States. To build a database of civilian homicides comparable to that of the police killings identified by MPV, we draw on data from the *Gun Violence Archive* (GVA). This database is collected by a nonprofit organization aiming to register all known shootings in the country. Incidents in the GVA are collected daily from over 7,500 law enforcement, media, government, and commercial sources. Each incident is verified by an initial researcher and subsequently subjected to a secondary validation process. Like the MPV data, the GVA includes the victim’s name in the vast majority of cases. The GVA also includes information on the suspect, when available. Since the GVA does not include information on the race of the victim or suspect, we impute the posterior probability that each subject belongs to a particular racial or ethnic group with the information on the name and location where the incident took place ([Imai & Khanna, 2016](#); [Khanna et al., 2017](#)).¹³ Finally, to isolate only civilian homicides in the GVA, we drop all suicides and accidental deaths as well as deaths due to a police shooting.

Television News Broadcasts. Our media dataset contains the universe of closed caption text across all television news programs in the United States.¹⁴ The data were provided by *News Exposure* (NE), a data vendor that monitors and collect transcripts from over 948 distinct TV stations across the 210 media markets in the US. Both local and national stations are included in the database.¹⁵

All together, over 2 million station–days of news transcripts are available for our analysis. As we describe below, we search these comprehensive television news transcripts for stories about the police killings and civilian homicides recorded in the MPV and GVA data, respectively. In addition to the text of an associated news story, we obtain information on the broadcast station, network affiliation, media market, date and time, run time, publicity value, and ratings estimate.

We complement the previous data sources with information on the demographics of the tract and media market in which the killings took place from the American Community Survey and census. We also merge the designated media area (DMA) demographic data and electoral results from [Martin & McCrain \(2019\)](#).

3.2 Sample Construction

To measure the use of obfuscatory language structures in media coverage of police killings, we need a meaningful benchmark, as absolute levels of such usage would be difficult to interpret. To this end, our primary

¹³The imputation algorithm uses the probability of an individual’s being part of a racial or ethnic group based on the name and census tract, with a 50 percent probability threshold as in [Moreno-Medina \(2021\)](#) and [Humphries et al. \(2019\)](#). This posterior probability is estimated with the package `WRU` in R.

¹⁴The data encompass all available time slots. While the majority of news programs are typically broadcast in the morning (5 AM–8 AM), afternoon (4 PM–6 PM) and night (9 PM–11 PM) and at noon (12 PM), our sample also considers other time frames in case a TV station airs a story at a different time.

¹⁵These data are also used in [Moreno-Medina \(2021\)](#).

analysis compares media coverage of police killings to that of civilian homicides. We impose several sample restrictions to ensure that this comparison is as meaningful as possible.

First, since the GVA database includes only information on gun deaths (not other forms of homicide), we limit our sample to police killings caused by gunshot, which represent more than 90 percent of all police killings.¹⁶

Table 2: Descriptive statistics

	All (1)	Police killings (2)	Civilian killings (3)
Panel A: Subject Level			
Victim characteristics:			
Age	35.13	36.79	33.90
Male	0.81	0.95	0.72
Black	0.24	0.18	0.28
Hispanic	0.12	0.16	0.09
White	0.52	0.58	0.47
Other/Unknown	0.12	0.07	0.15
Incident characteristics:			
Body-worn camera	0.11	0.11	.
Victim not fleeing	0.66	0.66	.
Observations	13,702	5,759	7,943
Panel B: Sentence Level			
Dimensions of obfuscation:			
Passive	0.20	0.22	0.17
Nominalization	0.04	0.04	0.03
No Agent	0.16	0.16	0.15
Intransitive	0.11	0.12	0.11
Any Obfuscation	0.34	0.36	0.29
No Explicit Agent	0.26	0.27	0.25
Observations	466,636	320,042	146,594

Notes: This table presents means of different variables for the cases in our sample. *Data sources:* GVA, MPV and News Exposure.

Second, we impose a set of sample restrictions designed to isolate the circumstances in which media *could have used similar language* to describe police killings and civilian homicides. A key issue is that the available information is different for police and civilian killings. By definition, in most cases, when a police officer is the perpetrator, we can assert that a police officer killed a person. It is therefore feasible to construct active phrases of the form “a police officer killed a person.” However, the formulation of such phrases becomes more challenging in cases where a specific suspect for a civilian killing has not been identified. By contrast, it may not be possible for media to identify a causal agent in a civilian homicide when a suspect/perpetrator has not yet been identified. In the absence of such information, it is natural for media to instead focus the narrative on the victim and abstract from the agent—e.g., “a 40-year old man was shot last night.” To make

¹⁶In particular, of the 7,663 police killings documented by MPV, 7,299 are caused by gunshot, and our final sample consists of the 7,293 of these that could be geolocated.

the police and civilian killings data as comparable as possible, therefore, we limit our sample of civilian homicides to those in the GVA database in which the name of the suspect is known.¹⁷

For our baseline analysis, we further limit our sample to news stories in which the suspect’s name appears. To avoid concerns that this sample restriction biases our analysis toward finding a greater use of active sentence structure for civilian homicides, we consider a number of alternative specifications to ensure robustness, including dropping sentences that include the suspect’s name and focusing on the first sentence in the story.

Appendix Table D.1 shows how the case composition changes with our sample construction choices. For police killings, our sample is very similar on all observables. For civilian killings, the requirement that the suspect is known increases the share of domestic violence and murder-suicide cases in the data. In turn, our sample has more women, and the victims are a little older than the average shooting victim in the United States. The racial and geographic composition of the victims in our analysis sample is the same as that of the full sample.

To match police killings and civilian homicides to news coverage, we use a machine learning–based procedure that follows three sets of requirements intended to ensure a high-quality match, as in [Moreno-Medina \(2021\)](#). First, we subset the NE data to text transcriptions that include words related to a killing/homicide such as “shot,” “shoot,” or “killed,” which sharply increases the probability that a news story is about crime.¹⁸ We keep stories with a score above a certain threshold and manually check the accuracy of this threshold, finding that 99 percent of all the identified stories indeed cover a crime or police incident.¹⁹

Second, we require a story to contain either the name of the victim or the address (block and street) in which the event happened. Third, we consider only stories that aired within 7 days of the victim’s death. The goal of this last restriction is to limit misclassification of stories (especially for victims with common names) by essentially requiring a match on date and name or address. Finally, for our sentence-level analyses, we focus on sentences in which i) there are references either to the victim or the suspect and ii) the sentence is informing on the killing. See Section 3.3 for further details. Our final sample includes 192,944 stories and 466,636 sentences linked to 5,759 police gun killings and 7,943 civilian gun homicides for which we are able to find at least one broadcast news story.²⁰

Panel A of Table 2 presents descriptive statistics for the police killing and civilian homicide samples. Compared to victims of civilian gun homicides, victims of police shootings are much more likely to be male (95 percent compared to 72 percent), are slightly older (38 years old compared to 34 years old), more likely to be Hispanic (16 percent compared to 9 percent) and White (58 percent compared to 47 percent), and less likely to be Black (18 percent compared to 28 percent). We control for these demographic variables in all of our empirical analyses. Appendix Figure C.1 presents trends over time in the number of police killings and civilian homicides included in our final sample. Our sample includes more cases of civilian homicides but more stories about police killings.

¹⁷Starting from an original sample of 49,277 gun deaths, we drop all suicides and accidental deaths, drop deaths due to a police shooting, and restrict the sample to those civilian homicides in which the name of the suspect is available in the GVA. These sample restrictions yield a dataset of 19,325 civilian gun homicides, and our final sample consists of the 17,939 of these that could be geolocated.

¹⁸For civilian homicides, we use all forms of the following keywords: “shot,” “gunshot,” “kill,” and “homicide.” For the police shootings, we search for all forms of the keywords “shot,” “gunshot,” “kill,” “homicide,” “police,” and “officer.”

¹⁹In this way, our algorithm does a good job of ruling out unrelated stories that might uses similar words—e.g., a sports story in which the word “shot” describes a basketball or soccer play rather than the action of a gun.

²⁰Note that there are on average 56 stories per police killing and 19 per civilian killing. In Section 4.2, we provide evidence that our results are not driven by differences in coverage volume.

3.3 Language Processing

In this section, we explain how we process the text data once we have our primary sample of news stories linked to police killings and civilian homicides. Appendix Figure A.1 presents a flowchart of our process. We apply three language processing steps to construct the measures of obfuscation. Our implementation of these steps uses a series of NLP models, based on bidirectional encoder representations from transformers (*BERT*).²¹ As will be clear below, we need a language model like *BERT* that captures contextual embeddings for words. The three steps in our process are as follows:

1. Coreference resolution: identify all words that reference the same individual
2. Identify who did what to whom in sentences about the shooting
3. Encode our measures of obfuscation

Additional details can be found in Online Appendix A.

Coreference Resolution and Story Delimitation. First, we identify all words used to refer to the same individual (victim or suspect), including pronouns. We adopt the method proposed by Lee *et al.* (2018) and implemented by Gardner *et al.* (2017), which uses a BERT-like neural network called SpanBERT (see Appendix A.2 for more details). The model takes text as an input and outputs a list of clusters of tokens (or words) considered to refer to the same individual. We define a broadcast story about the killing as the span between the first and last sentence in which the victim or suspect appears in the caption text.

Semantic Role Labeling. Second, we need to identify for each sentence who (agent) does what (action) to whom (patient). This task is known in the NLP literature as semantic role labeling (Appendix A.3 provides more details). We implement another BERT-type model, this time the one proposed by Shi & Lin (2019). Given that we want to identify how the killing is being covered, we focus on sentences that include verbs informing on the killing (“kill,” “shoot,” etc.). For these sentences, the algorithm produces an analysis for each verb, detailing who is executing the action and who is being acted upon. For our purposes, we want to identify who is executing the action of killing or shooting (agent) and who is being killed or shot (patient). This same output allows us to see whether an individual is the subject of the intransitive “to die” as well. We classify each verb into the following categories:

- Transitive: verbs that start with any of the texts in the list “kill,” “shot,” “gun,” “murder,” “shoot,” “hit,” “fire,” “open,” or “strike” or that are passive voice conjugations of these verbs²² or of “declare,” “find” or “pronounce” followed by a past participle such as “shot” or “killed”
- Intransitive: the verb “die” or an auxiliary verb followed by the past participle “died”
- Irrelevant: all others

We focus on sentences in which the identified causal patient for these verbs is the victim in our data.

²¹*BERT* is a neural network model of language that has proven to be incredibly successful at a host of tasks in NLP. BERT has several technical features, but perhaps the most important is that it trains the model using not only previous words in a text but also future ones. The standard model allows up to 512 words (or tokens) in a text. The network has 7 layers, and it works with a type of word embedder model that captures the context in which the word is being used. Since 2019, *Google Search* has been applying BERT models for English language search queries within the US.

²²That is, the verb is a form of “to be” followed by the past participle of one of the above verbs (for example, “the man was killed”).

Obfuscation Classification. Last, we create our measures of obfuscation as described in Section 2; that is, we identify which of the following structures appear in each sentence: an active-voice verb, a passive-voice verb, a nominalization, no agent, or an intransitive verb. Appendix A.4 presents the exact phrases used.

4 Obfuscation in News Stories on Police Killings

We now present the results of our analysis of news broadcasts examining whether media are more likely to use obfuscatory sentence structures in stories about police killings than in a control sample of stories about civilian homicides. In the first subsection, we present two analyses using *all* the sentences and the *first* sentences within identified news stories that reference the killing. We focus on the first sentence or “lead” because it is likely the most salient to viewers. In the next subsection, we present results from a number of additional specifications to examine whether the main results are robust under alternative designs, primarily related to sample selection.

4.1 Main Results: TV News Broadcasts

To offer an initial sense of whether news media are more likely to use obfuscatory sentence structures in stories about police killings, Panel B in Table 2 presents summary statistics on the prevalence of different structures broken down by whether a police officer or civilian was responsible for the killing. Overall, in the raw data, the use of any obfuscatory sentence structure is 24 percent more likely when a police officer was responsible for the killing (36 percent versus 29 percent). This aggregate result reflects the greater use of all four obfuscatory sentence structures in stories about police killings, especially passive voice (22 percent versus 17 percent). Appendix Figure C.2 plots the fraction of sentences with obfuscation over time, revealing that the increased prevalence of obfuscation in stories about police killings is stable over the study period.

To control for other potential differences in stories about police and civilian killings (e.g., the age, sex, and race of the decedent), we estimate regressions of the following form:

$$Obfuscation_{eitsd} = \beta_1 Police_i + \beta_2 X_{itsd} + \epsilon_{eitsd} \quad (1)$$

where e indexes a sentence about incident i at time t on station s in media market d . $Police_i$ is a dummy equal to 1 if the news story is about a police killing, and X_{eitsd} includes controls for incident characteristics, date, television station, and media market. All of our analyses except the one with no controls include a linear time trend.

For our first analysis, we treat each sentence in the story that references a killing as an observation and cluster the standard errors at the individual subject level (that is, at the incident/victim level). Table 3 presents results for five specifications. The specification shown in Column 1 includes no controls. Column 2 adds story-level controls (age, sex, and race of the victim), while Columns 3 and 4 successively add media market and television station fixed effects. The specification shown in Column 4 is our preferred specification because it effectively compares the coverage of police killings with that of civilian killings for the same television station in the same media market. Column 5 replaces the linear time trend for the specification in Column 3 with month \times year fixed effects, primarily to check whether there are any nonlinear time effects for which the linear time trend does not adequately control.

Table 3: Obfuscation in News of Police Killings: Sentence-Level Analyses

Outcomes: Dimension of Obfuscation	Mean Civ. Shoot.	Police Killing				
		(1)	(2)	(3)	(4)	(5)
Panel A: Aggregate Dimensions						
Any obfuscation	0.2935	0.062*** (0.008)	0.070*** (0.008)	0.073*** (0.008)	0.074*** (0.008)	0.075*** (0.007)
No explicit agent	0.2481	0.024*** (0.007)	0.031*** (0.007)	0.034*** (0.007)	0.034*** (0.007)	0.035*** (0.006)
Panel B: Individual Dimensions						
Intransitive	0.1075	0.011** (0.005)	0.013*** (0.005)	0.017*** (0.005)	0.018*** (0.005)	0.024*** (0.004)
No agent	0.1497	0.011** (0.006)	0.016*** (0.006)	0.016*** (0.006)	0.015*** (0.006)	0.010* (0.005)
Nominalization	0.0282	0.010*** (0.002)	0.011*** (0.003)	0.012*** (0.003)	0.012*** (0.003)	0.013*** (0.003)
Passive	0.1692	0.047*** (0.006)	0.052*** (0.006)	0.051*** (0.006)	0.051*** (0.006)	0.046*** (0.006)
Story Controls			X	X	X	X
DMA FE				X	X	X
Station FE					X	X
Month-Year FE						X
Observations		466,639	466,639	466,639	466,639	466,639

Notes: This table presents the differential obfuscation in stories about police killings and stories about civilian killings from our estimation of Equation 1. Our analyses are at the sentence level. We vary what controls are included across columns. Each row presents a separate regression coefficient on a dummy equal to 1 if the story is about a police killing rather than a civilian killing for different measures of obfuscation, which are described in the first column. Our sample includes incidents and news stories where a suspect was identified for civilian killings. All sentences include some mention of either the victim or suspect. We define “Any obfuscation” as a sentence with a passive-voice verb, no agent, an intransitive verb, or a nominalization. We define “No explicit agent” as a sentence with no agent, an intransitive verb, or a nominalization. See Section 2 for more details. Source: News Exposure. Standard errors clustered by subject (*: $p < 0.1$, **: $p < 0.05$, ***: $p < 0.01$).

For each of these five specifications, we report results for six dependent variables in the table rows. The final four rows report the results for the four distinct obfuscatory sentence structures described above, while the first two rows report results for dependent variables that aggregate these outcomes. The second row reports results for the aggregate category *No explicit agent*, which combines *No Agent* and *Intransitive*, while the first row aggregates all four categories to report the use of *Any obfuscation* in the sentence. We find a consistent pattern across all the specification results shown in Table 3: sentences in stories about police killings are approximately 25 percent (7 percentage points) more likely to use some form of obfuscation than stories about civilian killings, and there is an increased propensity to use each distinct form of obfuscation (rows 3–6). Visually, the results for Column 3 are presented in Panel A of Figure 1 as a share of the mean obfuscation in civilian killings.

In our preferred specification (Column 4), *Passive* and *No explicit agent* sentence structures are 30 percent (5.1 percentage points) and 14 percent (3.4 percentage points) more likely to be employed in stories about police killings, respectively. *Nominalization* is generally the least common form of obfuscation used by media, but it too is much more prevalent (43 percent, 1.2 percentage points) in stories about police killings. Overall, our results indicate that news coverage of police killings is significantly more likely than coverage of civilian

homicides to use narrative structures that obscure responsibility for the killing.

Table 4: Obfuscation in News Coverage of Police Killings: First Sentence in a News Story

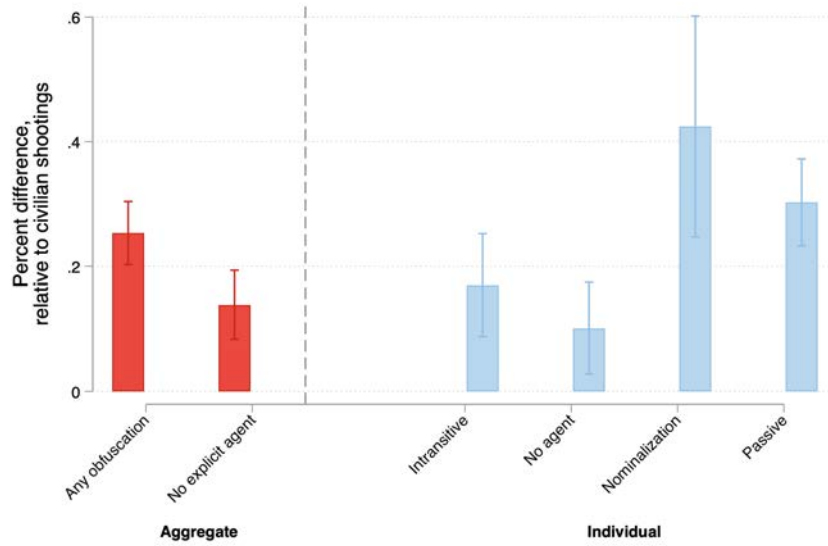
Outcomes: Dimension of Obfuscation	Mean Civ. Shoot.	Police Killing				
		(1)	(2)	(3)	(4)	(5)
Panel A: Aggregate Dimensions						
Any obfuscation	0.2815	0.105*** (0.010)	0.118*** (0.009)	0.120*** (0.009)	0.120*** (0.009)	0.120*** (0.009)
No explicit agent	0.2072	0.059*** (0.008)	0.070*** (0.009)	0.072*** (0.009)	0.071*** (0.009)	0.067*** (0.008)
Panel B: Individual Dimensions						
Intransitive	0.0748	0.035*** (0.005)	0.041*** (0.005)	0.046*** (0.005)	0.046*** (0.005)	0.051*** (0.005)
No agent	0.1405	0.022*** (0.006)	0.027*** (0.007)	0.025*** (0.007)	0.024*** (0.007)	0.014** (0.006)
Nominalization	0.0526	0.019*** (0.005)	0.020*** (0.005)	0.021*** (0.005)	0.022*** (0.005)	0.025*** (0.005)
Passive	0.1663	0.063*** (0.007)	0.069*** (0.008)	0.066*** (0.008)	0.065*** (0.007)	0.057*** (0.007)
Story Controls			X	X	X	X
DMA FE				X	X	X
Station FE					X	X
Month–Year FE						X
Observations		182,145	182,145	182,145	182,145	182,145

Notes: This table presents the differential obfuscation in stories about police killings and stories about civilian killings from our estimation of Equation 1, with our sample limited to the first sentence in a news story. We vary which controls are included across columns. Each row presents a separate regression coefficient on a dummy equal to 1 if the story is about a police killing rather than a civilian killing for different measures of obfuscation, which are described in the first column. Our sample includes incidents and news stories where a suspect was identified for civilian killings. All sentences include some mention of either the victim or suspect. We define “Any obfuscation” as a sentence with a passive-voice verb, no agent, an intransitive verb, or a nominalization. We define “No explicit agent” as a sentence with no agent, an intransitive verb, or a nominalization. See Section 2 for more details. Source: News Exposure. Standard errors clustered by subject (*: $p < 0.1$, **: $p < 0.05$, ***: $p < 0.01$).

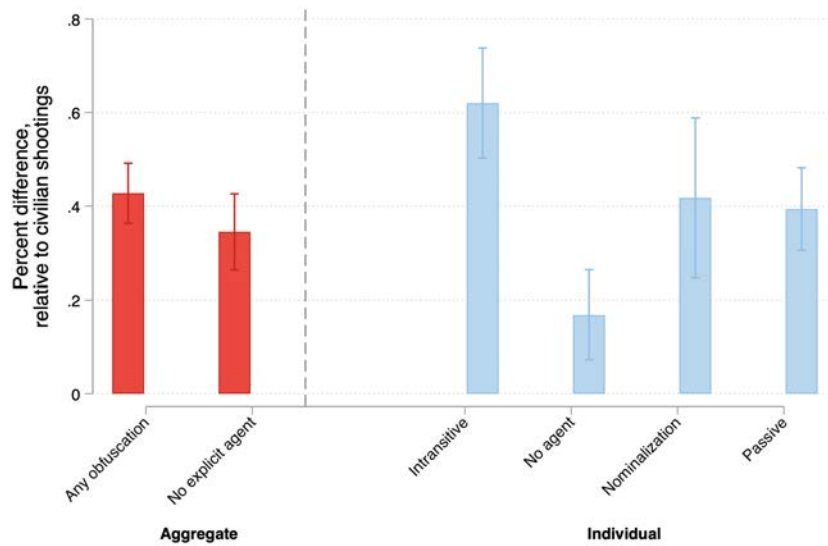
Table 4 presents a set of results analogous to those reported in Table 3 for specifications that include only the first sentence of the story. News organizations tend to present what they consider the most essential or attention-grabbing facts about a story in the first sentence, which is generally expected to be especially salient to viewers (AP, 2020). As a result, we expect any obfuscation in the lead to have an outsized effect on how viewers understand and respond to the incident.

The results of the analysis of first sentences are qualitatively similar to and quantitatively greater than those presented in Table 3, and in Panel B of Figure 1, where we plot coefficients as a share of the mean obfuscation in reporting on civilian killings. In this case, obfuscation is approximately 40 percent (12 percentage points) more likely for coverage of police killings than for reporting on civilian homicides. For first sentences, *Passive*, *No explicit agent*, and *Nominalization* are each approximately 35–50 percent more likely to be used in stories about police killings. Overall, the results presented in Table 4 suggest that the media’s use of obfuscation in coverage of police killings is especially likely in the first sentence—i.e., the most salient part of the story.

Figure 1: Obfuscation in News of Police Killings as Share of Mean Obfuscation in Civilian Killings



(a) All Sentences



(b) First Sentences

This figure plots the coefficients from Columns 4 in Tables 3 and 4 as a percentage of the mean obfuscation in the control group. Our sample includes incidents and news stories where a suspect was identified for civilian killings. Data sources: GVA, MPV and News Exposure.

4.2 Robustness Checks

Table 5 reports the results of a number of specifications designed to examine the robustness of our main findings to alternative study designs. For comparison, Panel A repeats the estimates from our preferred

specifications (Column 4) in Tables 3 and 4, while Panels B–D report analogous results for three alternative models.

Table 5: Robustness Tests

Dimension of Obfuscation:	All Sentences		1st Sentence	
	Any obfuscation (1)	No explicit agent (2)	Any obfuscation (3)	No explicit agent (4)
Panel A: Main results (Column 4 of Tables 2 and 3)				
Police Killing	0.074*** (0.008)	0.034*** (0.007)	0.120*** (0.009)	0.071*** (0.009)
Observations	466,639	466,639	182,145	182,145
Mean Civ.Shoot.	0.2935	0.2481	0.2815	0.2072
Panel B: Dropping sentences where suspect is named				
Police Killing	0.075*** (0.008)	0.035*** (0.007)	0.121*** (0.009)	0.072*** (0.009)
Observations	464,904	464,904	181,492	181,492
Mean Civ.Shoot.	0.2928	0.2474	0.2805	0.2061
Panel C: No domestic violence in civilian shooting				
Police Killing	0.069*** (0.008)	0.031*** (0.007)	0.117*** (0.009)	0.072*** (0.009)
Observations	434,362	434,362	170,012	170,012
Mean Civ.Shoot.	0.2983	0.2502	0.2831	0.2022
Panel D: Weighted by 1/# sentences				
Police Killing	0.073*** (0.005)	0.048*** (0.004)	0.111*** (0.007)	0.081*** (0.007)
Observations	466,639	466,639	182,145	182,145
Mean Civ.Shoot.	0.3237	0.2677	0.3062	0.2213
Controls	Story+DMA FE			

Notes: This table presents robustness tests for our main results for all sentences (Columns 1 and 2) and for the first sentence (Columns 3 and 4). Panel A presents our preferred specification from Tables 3 and 4 (Column 4). In Panel B, we drop sentences in the civilian killing sample where the suspect is named. In Panel C, we drop stories about domestic violence incidents. In Panel D, we reweight sentences by 1/number of sentences in a particular story. We define “Any obfuscation” as having a passive-voice verb, no agent, intransitive verb, or nominalization. We define “No explicit agent” as a sentence with no agent or with an intransitive verb. See Section 2 for more details. Source: News Exposure. Standard errors clustered by subject (*: $p < 0.1$, **: $p < 0.05$, ***: $p < 0.01$)

In our construction of the control sample of civilian homicides for the analysis, our goal was to identify situations in which the media faced a similar choice of language for both police and civilian killings. As a key sample selection criterion, we require the suspect’s name to appear in the story. This choice aims to ensure that an agent (police officer, civilian suspect) could possibly have been mentioned in the story—i.e., to exclude cases where a death occurred but nothing about a potential suspect (or even whether the incident was a homicide) was known at the time of the news report. One of our primary concerns with this sample selection criterion, however, is the possibility that requiring the suspect’s name to appear in the story might bias our sample toward including more active-voice sentence structures for civilian homicides. We are particularly concerned that the suspect’s name might appear commonly as the subject of a sentence

describing the murder.

To address this concern, Panel B reports the estimates for a specification that drops all sentences that include the suspect’s name. The results are virtually unchanged, suggesting that our concerns about greater use of active-voice sentence structures involving the suspect’s name are unfounded. Interestingly, only 1,735 sentences are dropped in this specification—a much smaller figure than the total number of stories (58,033) about civilian homicides, despite our criterion that the suspect’s name appears in the story. This implies that in most cases when the suspect’s name appears in a news story, it is not in a sentence included in our main analysis (which must directly describe the death/killing). Instead, the suspect is often named in a stand-alone sentence (e.g., “John Doe was identified as the suspect in the case”) and, thus, does not make it into our main analysis sample.

Panel C reports results from specifications that remove news stories about domestic violence. Our concern in this case is that news stories about domestic violence might be especially likely to center on the victim, resulting in the use of different sentence structures. The results reported in Panel C are largely unchanged.

We then provide two tests to see whether the difference in language is due to the difference in volume of coverage per story. As specified earlier, there are on average 56 stories per police killing but 19 stories per civilian killing. One may be concerned that a small number of viral incidents that receive a great deal of media attention drive our main results. First, Panel D reweights each sentence by the inverse of the total number of sentences per victim, thereby giving equal weight to all victims. The findings are again remarkably similar to our main results, implying that the increased use of obfuscatory sentence structures for police killings is not limited to high-profile cases. Second, in Appendix Figure C.4 we break down our sample by whether an incident led to “viral” media coverage, with viral coverage defined as coverage in more than 100 news segments. For nonviral incidents, there are on average 16 stories per civilian killing, compared to 27 per police killing (these numbers are 447 and 620 for viral stories). The results are very similar in these two kinds of stories, suggesting that volume of stories per incident does not drive the differential obfuscation.

Last, Appendix Figure C.3 provides further robustness tests. We rerun our main estimates for different subgroups: limiting our sample to years for which we have both GVA and MPV data (2014–2018); limiting our sample to killings covered on two or more days (these may be less likely to be updates related to having identified the perpetrator); separately considering news stories aired on the day of the shooting, on the next day, or on subsequent days; limiting our sample to stories that include in the text the words “shoot,” “kill,” both “shoot” and “kill,” or “die”; and dropping accidental shootings in the civilian stories, dropping police killings flagged as involving retired or off-duty officers; dropping police killings flagged in the GVA as “suicide by cop;” and splitting our sample by whether a news story is above or below the median length. The results are again unchanged in these additional robustness checks.

4.3 Heterogeneity by Victim Characteristics

As the discussion on policing, and in particular police killings, in the US is generally connected to societal treatment of Black civilians (Logan & Myers (2022); Mason *et al.* (2022)), it is crucial to evaluate which characteristics of the incidents predict higher levels of obfuscation. We begin by exploring how incident characteristics correlate with the use of obfuscation (Appendix Figure C.4). We focus on victim race and gender. We split the sample from our main analyses along these characteristics.

First, we find that the point estimates for the level of obfuscation are smallest for Black victims and largest

for White victims, suggesting that there is less obfuscation for cases involving the deaths of Black individuals. However, one needs to be cautious with this finding as the 95 percent confidence intervals indicate that we cannot reject that these coefficients are not statistically different from each other. Moreover, we find that incidents involving female victims show a level of obfuscation twice as large as that of their counterparts involving male victims.

4.4 Supporting Evidence from Newspapers

Is the obfuscation exclusive to television coverage of police killings, or does it manifest across various media platforms? To provide some insight into this question, we conduct a comprehensive analysis of the use of obfuscatory language in American newspapers.

We collected newspaper coverage of police and civilian killings in the United States from NexisUni, spanning from 2013 to 2019. The sample includes news stories sourced from 166 outlets, ranging from nationwide publications such as the *New York Times* to mid-sized local papers such as the *Chicago Daily Herald* and smaller sources such as the *Bakersfield Californian*. We apply the same sample restrictions as those described for the TV stories in Section 3.2. The newspaper sample contains approximately 49,000 stories on 4,915 incidents (2,792 police; 2,123 civilian). Appendix Table D.2 provides descriptive statistics for this sample; the characteristics of victims and incidents mirror those in our primary TV sample.

Our analytical approach mirrors the methodology used in our study of broadcast news reports. We measure different levels of obfuscation both overall and within the first sentences of newspaper stories, and we compute the differential obfuscation for police versus civilian killings. Our results are presented in Appendix Tables D.3 and D.4. Similarly to the case for TV broadcasts, our analysis shows that, across most measures, newspaper articles exhibit a significantly higher propensity to employ obfuscatory language when a police officer is the perpetrator than when the perpetrator is a civilian. There is one exception: the use of nominalizations, which is more prevalent in printed press narratives pertaining to civilian killings. This could stem from differences between oral and written language in commonly used phrases. If we set aside this difference, the magnitude of the outcomes remains notably consistent between newspapers and broadcast news reports. This analysis suggests that the use of obfuscatory language to describe police killings is not exclusive to television news but rather extends across the broader spectrum of news media.

5 Effects of Obfuscatory Language

The analysis presented in Section 4 reveals the systematic use of more obfuscatory language in broadcast news coverage of police killings than in reporting on civilian homicides. Does this matter in practice, however? It could be argued that the sentences “a police officer shot and killed a man” and “a man died in an officer-involved shooting” contain the same information. Does the difference in language really affect how viewers or readers understand and respond to a news story?

To answer this question, we pursue two complementary lines of analysis. First, we use a large, nationally representative survey to examine how stated support for police funding varies with the language used in recent media coverage of police killings. This analysis exploits variation in local media coverage around the exact date when the respondent took the survey. The results provide motivating evidence that obfuscation in

real-world media coverage affects attitudes about policing in the field.

Our second line of analysis is based on an online experiment that measures how the narrative structure used to describe a police killing affects a respondent’s assessment of the officer’s moral responsibility, demand for accountability for the officer, broader support for police reform, and subsequent retelling of the story. The strength of this experimental study is that the design can be tightly controlled to isolate the effect of language from the impacts of any other aspects of an event or its media coverage. The results of this analysis provide clear evidence that obfuscatory language affects respondents’ perceptions and judgment regarding the incident as well as their subsequent policy preferences.

5.1 Obfuscation and Support for Police Funding: Motivating Evidence from the Cooperative Election Study

Data Source and Sample Selection To study whether obfuscation influences perceptions of policing in a field setting, we link data from the Cooperative Election Study (CES), a biennial, nationally representative survey, with our primary broadcast news dataset. Specifically, we use data from the post-election waves of the CES conducted in November from 2014 to 2018 (Kuriwaki, 2023; Schaffner *et al.*, 2023). Respondents provide information on demographic characteristics, political orientation, education, and employment and marital statuses. The survey includes a question on law enforcement funding that asks whether they support state legislatures in “greatly increasing,” “slightly increasing,” “maintaining,” “slightly decreasing,” or “greatly decreasing” financial allocations to law enforcement.

We match the CES data from 2014, 2016, and 2018 with our measures of media coverage of police killings based on the media market (DMA) of the respondent’s residence at the time of survey participation. We limit our analysis sample to respondents who live in a DMA where there was at least one police killing covered on television on the day that the respondent took the survey. In this way, we isolate the effects of obfuscation, conditional on there having been media coverage of a police killing. Our main sample has 23,126 CES respondents; Appendix Table D.5 presents sociodemographic characteristics of the respondents in our sample. There are no notable differences in terms of sex, race, ethnicity, age or ideology across levels of obfuscation for stories about police killings on TV on the day before a story was aired.

Empirical Strategy Because of the survey’s extensive scope, the data collection occurred over multiple days. This results in quasi-random within-DMA variation in media content on the date when an individual happened to have responded to the survey.²³ For each respondent i residing in DMA d and exposed to the level of obfuscation on date t , support for more police funding is given by

$$Support_{idt} = \beta_t Obfuscation_{d,t} + \beta_2 X_i + \gamma_d + \epsilon_{idt} \quad (2)$$

In this equation, $Support_{idt}$ is an ordinal variable ranging from one to five, where one corresponds to “greatly decreasing” and five to “greatly increasing” financial allocations to law enforcement. We include in the analysis media market fixed effects, γ_d , and a vector of controls, X_i , including a survey year fixed effect and the respondent’s age, gender, race, political leaning, education level, and marital and employment

²³Similar approaches are used in Sharkey (2010) and in Philippe & Ouss (2018), which both exploit variation in the timing of when large-scale surveys are taken relative to local events and media stories.

statuses. Last, we control for the number of sentences in police-related stories and for whether a police incident occurred on the same day as a killing within the DMA on date t . The error term is captured by ϵ_{idt} , and standard errors are clustered at the DMA level. The variable of interest, $Obfuscation_{dt}$, captures the share of obfuscation in local broadcast news coverage on date t .

Results Panel A of Table 6 presents the impact of obfuscation on public sentiment toward police funding. Following the main analyses in Table 3, the first two columns use the share of sentences with any kind of obfuscation as the measure of media obfuscation; the last two columns use the share of sentences with no specific agent instead. All specifications include survey year fixed effects and controls for story length, whether the news story is about a killing that took place that day (versus a previous day), and media market fixed effects; the even columns add controls for respondent characteristics.

Table 6: Obfuscation and Support for Police Funding: Survey Evidence

	Outcome: Support for police funding			
	(1)	(2)	(3)	(4)
Panel A: Pre-survey police shooting				
Any Obfuscation	0.051*	0.042**		
	(0.029)	(0.020)		
No Explicit Agent			0.066**	0.060***
			(0.027)	(0.020)
Mean of Dep.	3.59	3.59	3.59	3.59
Observations	23,120	23,120	23,120	23,120
Panel B: Post-survey police shooting (placebo)				
Any Obfuscation	0.022	0.008		
	(0.015)	(0.006)		
No Explicit Agent			0.025	0.003
			(0.021)	(0.008)
Mean of Dep.	3.59	3.59	3.59	3.59
Observations	157,318	157,318	157,318	157,318
DMA FE	Yes	Yes	Yes	Yes
Controls	No	Yes	No	Yes

Notes: This table analyzes the effect of obfuscation on public support for police funding, focusing on changes in financial commitments to police departments. The dependent variable measures these changes on an ordinal scale ranging from 1 (“greatly decreasing”) to 5 (“greatly increasing”). Panel A correlates survey responses to obfuscation in stories on police killings aired before the survey was administered. Panels B and C and offer placebo tests: Panel B examines obfuscation in reporting on police killings after the survey, and Panel C examines obfuscation in reporting on civilian killings that occurred before the survey. Standard errors, presented in parentheses, are clustered at the DMA level. *** $p < 0.001$; ** $p < 0.05$; * $p < 0.1$. Source: News Exposure and Cooperative Election Survey.

The resulting estimates are consistent across specifications: conditional on the media’s covering a police

killing, more obfuscation is correlated with more support for police funding. Using our preferred estimates (Columns 2 and 4), we find that a ten-percentage-point increase in obfuscation increases public support for augmenting police funding by 11.6 percent (for *Any obfuscation*) and 16.7 percent (for *No explicit agent*). These results indicate that media obfuscation and ambiguous narratives in the coverage of police killings can shape public sentiment in a manner that favors continued support for law enforcement funding.

Placebo Test. To test our design’s validity, we conduct a placebo test examining obfuscation in media coverage of police killings within the 7 days *after* the date when the survey was taken. The idea is that post-survey obfuscation should not impact responses. Panel B of Table 6 presents these estimates. The results imply no systematic correlation between obfuscation in stories about police killings the week after the survey was taken and survey respondent support for policing.

Our analysis of the CES has the advantage of exploiting actual variation in obfuscation in media coverage. However, there are several limitations of this analysis. The CES was not designed with the primary intent of measuring people’s perceptions of policing, nor is the relevant subsample very large, leading to somewhat imprecise estimates. We also do not have information about whether survey participants watched TV on the days leading to the survey responses. Most importantly, these real-world incidents and corresponding stories may vary along other dimensions that are both correlated with the level of obfuscation and directly affect public perceptions (for example, variation in the nature of events or in the details included in the story). Motivated by these important potential concerns, the next section describes the results of an online survey experiment that we conducted. The key advantage of the experiment is that it allows us to precisely control exposure to the narrative structures themselves.

5.2 Obfuscation and Perceptions of Policing: Lab Experiment

5.2.1 Experimental Design

We conducted an online experiment with 2,402 participants in March 2022 using Prolific. Our hypotheses and research and analysis design were registered on the AEA registry (AEARCTR-0009052). Participants were required to reside in the United States and to be adults fluent in English.²⁴

We presented participants with a story (a headline sentence plus four sentences providing further detail) about a police killing.²⁵ Participants all read about the same incident but were randomly assigned to variations in how it was described using a 4×2 design. The first level of randomization was for narrative structure. Participants were randomized into one of four structures: (1) *Active*, (2) *Passive*, (3) *No Agent + Nominalization*, and (4) *Intransitive + Nominalization*.²⁶ We note that, unlike in our media analyses, we do

²⁴The survey took on average six minutes to complete, and participants were paid \$1.70 to participate. Appendix Table D.7 presents balance tables that confirm that randomization worked properly.

²⁵The story reflects a real incident, but we anonymized information about the person and city to make sure that participants were not influenced by any prior knowledge of the incident. We note that our main media analysis uses data on TV news stories but that, in our lab experiment, to isolate one causal channel and not capture differences in voice or tone, we use a written support. However, our results are similar when we analyze text used in the print press. In addition, we exploit the fact that the News Exposure database maintains an archive of video recordings for the three most recent months of news programs. We analyze the content of the video coverage of police killings for a convenience sample from April to June 2023. We compare the images in the videos based on whether there is obfuscation in the stories. Research assistants coded the videos using a rubric with the most frequent types of images present in news stories about police officers killing civilians. Appendix Table D.6 shows that there are no notable differences by level of obfuscation.

²⁶Note that we do not provide any information on the officer or on the victim’s race, since this was not part of our media analyses. This could be an interesting dimension to investigate in future research.

not have *No Agent* and *Intransitive* alone; instead, we also include the language “officer-involved shooting” in the story that we present to subjects. This is because, in the context of our experiment, readers would not be able to infer anything about police presence from pure “no agent” stories (“a person was killed”) or pure “intransitive” stories (“a person died”). We hypothesize that, if anything, our narrative structures reduce the contrasts that we study in our media analyses.

Table 7 provides the headline sentence used in each of these sentence structure treatment arms; the full prompts can be found in Appendix B.1. Following the definitions used earlier in the paper, we define *Any obfuscation* as sentences with grammatical structure (2), (3) or (4) and *No explicit agent* as sentences with narrative structure (3) or (4). The second level of randomization determined whether a clause stating that the man killed “was reportedly armed” was included in the story.

Table 7: Online Experiment: First Sentence of News Story for Each Narrative Treatment Arm

Narrative treatment arm	Headline
Active	A police officer killed a 52-year-old man on Friday night.
Passive	A 52-year-old man was killed by a police officer on Friday night.
No agent + Nominalization	A 52-year-old man was killed in an officer-involved shooting on Friday night.
Intransitive + Nominalization	A 52-year-old man died in an officer-involved shooting on Friday night.

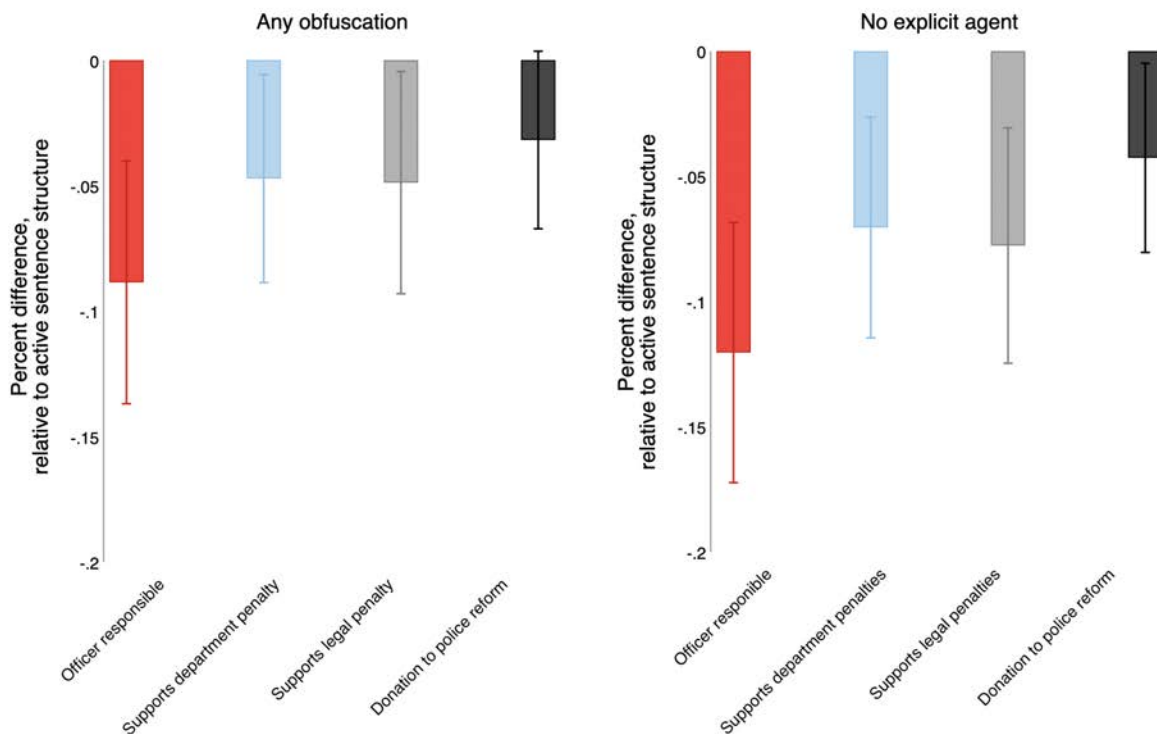
We are interested in understanding how the narrative structure influences three broad sets of outcomes. First, does it influence how someone understands and judges the specific event being described? We study this by asking participants questions on their perceptions of the officer’s moral responsibility for the civilian’s death and demand for penalties for the officer. Second, does it alter someone’s broader understanding of policing harms and their support for police reform? To get at this, we ask respondents how they would like to split a potential \$100 donation between two organizations: one supporting officer well-being and the other supporting police reform. We use donations to the latter as our primary measure of support for reform. Finally, does the narrative structure affect how respondents recall and retell the story? We study this by measuring both the information content (i.e., whether they report that a police officer was responsible for the killing) and narrative structure used by participants in their recounting of the story at the end of the experiment. Our exact questions can be found in Appendix B.2.

We hypothesize that obfuscation matters—that is, that respondents are less likely to consider the officer morally responsible, demand accountability, and support police reform when the news story presents some obfuscation than when the report is delivered in active voice. In addition, we hypothesize that the degree of obfuscation, as outlined in Section 2, is important: the greater the obfuscation, the less likely people are to assign responsibility, demand accountability, or support reform. Last, we hypothesize that these effects are strongest if we do not specify that the victim was armed.

5.2.2 Experimental Results

Judgments about Officer’s Actions. Table 8 and Figure 2 present the main results of the experiment. The table presents our primary outcomes of interest: whether the officer is morally responsible for the killing (Columns 1–3), support for penalties for the officer from their police department (Columns 4–6) and support for broader legal penalties for the officer (Columns 7–9). For each outcome, the three columns present the effect of (i) *Any obfuscation*, (ii) *No explicit agent*, (iii) *Passive, No Agent*, and *Intransitive*, using *Active* as the reference group for all three columns. Figure 2 presents these estimates as a share of the mean obfuscation in the control group.

Figure 2: Treatment Effect as Percentage of Control (Active) Mean



This figure plots the coefficients in Tables 8 and 11 as a percentage of the mean obfuscation in the control group. Our sample includes incidents and news stories where a suspect was identified for civilian killings. Data sources: GVA, MPV and News Exposure.

In line with our main hypothesis, how the story is told matters for perceptions of what happened. In particular, removing the mention of an explicit agent reduces perceptions of an officer’s responsibility and demand for penalties. The *No explicit agent* treatment decreases the share responding that the officer was morally responsible by 13 percent (9 percentage points, $P < 0.001$) and decreases the stated preference for departmental penalties by 7 percent ($P = 0.001$) and for legal penalties by 8 percent ($P < 0.001$). In contrast, we find no significant differences across the three measurements between the *Passive* and *Active* arms. As a

result, the statistical significance of *Any obfuscation* rises to only the 10 percent level for the outcomes related to demanding penalties for the officer.

We also hypothesize that specifying whether the victim was armed substantially influences respondents' perceptions of whether the officer's actions were justified. We find this to be the case, as shown in Table 9. Respondents are 13 percent (9 percentage points) less likely to say that the officer is morally responsible for the victim's death when the story specifies that the victim was armed. Participants are also 19 and 22 percent less likely to agree that the officer should face penalties within the department or legal penalties, respectively.

Table 8: Online Experiment: Narrative Structure and Judgment of Event

	Moral responsibility			Department penalty			Legal penalty		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Any obfuscation	-0.06*** (0.02)			-0.19* (0.10)			-0.18* (0.10)		
No explicit agent		-0.09*** (0.02)			-0.28*** (0.11)			-0.29*** (0.11)	
Passive		-0.02 (0.03)	-0.02 (0.03)		-0.00 (0.12)	-0.00 (0.12)		0.03 (0.12)	0.03 (0.12)
No agent + nominalization			-0.06** (0.03)			-0.28** (0.12)			-0.31** (0.12)
Intransitive + nominalization			-0.11*** (0.03)			-0.27** (0.12)			-0.26** (0.12)
Mean Dep. Var.	0.72	0.72	0.72	3.93	3.93	3.93	3.73	3.73	3.73
SD Dep. Var.	0.45	0.45	0.45	2.15	2.15	2.15	2.16	2.16	2.16
N	2402	2402	2402	2402	2402	2402	2402	2402	2402

Notes: "Any obfuscation" is equal to 1 if the sentence structure has a passive-voice verb, has no agent, or has an intransitive verb. "No explicit agent" is equal to 1 if the sentence structure has no agent or has an intransitive verb. The outcome in Columns 1–3 is a dummy equal to 1 if the respondents think that the police officer is morally responsible for the victim's death. The outcomes in the remaining columns capture support, on a scale from 1 to 7, for department penalties (Columns 4–6) and legal penalties (Columns 7–9), respectively. See Appendix B for the full questions. We report the mean and standard deviation of the dependent variable for the "Active" narrative structure. Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

In addition to the independent effects of narrative structure and the presence of a weapon on participants' responses, we propose a third hypothesis related to the interaction of the two treatments. In particular, we conjecture that obfuscatory narrative structures are especially impactful when the story does not mention a weapon. This hypothesis is based on the idea that the presence of a weapon would lead some participants to determine that the shooting was justified regardless of how the information is presented in the story. In economic terms, we hypothesize that the presence of a weapon and obfuscation may be substitutes when it comes to officer responsibility: perceptions of responsibility can be dampened by contextual features; or in the absence of these more favorable features, by how a story is told. To examine this hypothesis, Table 10 breaks down the analysis presented in Table 8 by whether the story includes (Panel B) or does not include (Panel A) a clause stating that the decedent "was reportedly armed."

Table 9: Online Experiment: Perceptions of Police Killing, Depending on Specification of Whether Victim Had a Weapon.

	Moral reponsibility (1)	Department penalty (2)	Legal penalty (3)
Weapon	-0.09*** (0.02)	-0.75*** (0.08)	-0.83*** (0.08)
Mean Dep. Var.	0.72	3.93	3.73
SD Dep. Var.	0.45	2.15	2.16
N	2402	2402	2402

Notes: The “Weapon” variable is equal to 1 if we specify that the victim was armed. The outcome in Column 1 is a dummy equal to 1 if the respondents think that the police officer is morally responsible for the victim’s death. The outcomes in Columns 2 and 3 capture support, on a scale from 1 to 7, for department and legal penalties, respectively. See Appendix B for the full questions. We report the mean and standard deviation of the dependent variable for stories that do not specify whether the victim had a weapon. Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

In line with our hypothesis, the point estimates are greater in magnitude for *all 15 effects* related to the use of obfuscatory language reported in each panel when the story omits any mention of a weapon, though we note that the differences are not significant at conventional levels. In the absence of information about a weapon, the estimated effects are especially large when the story does not explicitly identify an agent (*No Agent* and *Intransitive*). The estimated effects are also negative for the *Passive* treatment in this case but still mostly do not rise to statistical significance at conventional levels. Overall, the estimated effect of *Any obfuscation* and *No explicit agent* are negative and statistically significant for all three outcomes when the story does not mention the potential presence of a weapon. In contrast, the estimated effects for *Any obfuscation* are much smaller in magnitude and not statistically significant when the story states that the decedent was reportedly armed.

Appendix Table D.8 shows similar patterns for two additional outcomes: whether the respondent thinks that the officer was justified in shooting the person and whether the respondent thinks that the officer was depicted negatively in the story. We find that—especially for stories that do not mention a weapon—more obfuscation both increases the perception that the officer was justified in shooting and decreases feelings that the officer was negatively depicted.

Broader Perceptions of Policing Harms and Demand for Reform. We next investigate how narrative structure affects broader perceptions of policing harms and demand for reform beyond those related to the focal incident. Columns 1–3 of Table 11 report results for how narrative structure affects how much participants would donate to a nonprofit focused on police reform (vs. one focused on officer well-being), while Columns 4–6 report analogous results for participants’ estimate of the prevalence of police killings in the United States. The point estimates for donations are negative but generally smaller in magnitude than those related to the participants’ judgments about the specific event shown in Table 8. In this case, the narrative structure with *No explicit agent* reduces donations by approximately 4 percent (2.5 percentage points), while the use of *Passive* voice continues to have negligible effects. Interestingly, the use of more obfuscatory sentence structures also tends to reduce participants’ estimates of the number of annual police killings in the United States. This suggests that some of the decline in support for reform may arise from

Table 10: Online Experiment: Narrative Structure and Judgment of Event, by Presence or Absence of Weapon

	Moral responsibility			Department penalty			Legal penalty		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: No mention of victim weapon									
Any obfuscation	-0.10*** (0.03)			-0.35** (0.14)			-0.32** (0.14)		
No explicit agent		-0.12*** (0.03)			-0.40*** (0.15)			-0.40*** (0.15)	
Passive		-0.06* (0.03)	-0.06* (0.03)		-0.24 (0.16)	-0.24 (0.16)		-0.16 (0.16)	-0.16 (0.16)
No agent + nominalization			-0.08** (0.04)			-0.37** (0.17)			-0.40** (0.17)
Intransitive + nominalization			-0.15*** (0.04)			-0.44*** (0.17)			-0.41** (0.17)
Mean Dep. Var.	0.80	0.80	0.80	4.44	4.44	4.44	4.25	4.25	4.25
SD Dep. Var.	0.40	0.40	0.40	2.08	2.08	2.08	2.06	2.06	2.06
N	1201	1201	1201	1201	1201	1201	1201	1201	1201
Panel B: Victim has a weapon									
Any obfuscation	-0.04 (0.03)			-0.12 (0.14)			-0.15 (0.14)		
No explicit agent		-0.07** (0.03)			-0.23 (0.14)			-0.26* (0.15)	
Passive		0.01 (0.04)	0.01 (0.04)		0.11 (0.17)	0.11 (0.17)		0.08 (0.17)	0.08 (0.17)
No agent + nominalization			-0.05 (0.04)			-0.25 (0.17)			-0.28* (0.17)
Intransitive + nominalization			-0.08** (0.04)			-0.21 (0.17)			-0.24 (0.17)
Mean Dep. Var.	0.66	0.66	0.66	3.51	3.51	3.51	3.29	3.29	3.29
SD Dep. Var.	0.47	0.47	0.47	2.11	2.11	2.11	2.16	2.16	2.16
N	1201	1201	1201	1201	1201	1201	1201	1201	1201

Notes: Panel A presents the results when we do not specify and Panel B the results when we do specify that the victim was armed. “Any obfuscation” is equal to 1 if the sentence structure has a passive-voice verb, has no agent, or has an intransitive verb. “No explicit agent” is equal to 1 if the sentence structure has no agent or has an intransitive verb. The outcome in Columns 1–3 is a dummy equal to 1 if the respondents think that the police officer is morally responsible for the victim’s death. The outcomes in the remaining columns capture support, on a scale from 1 to 7, for department penalties (Columns 4–6) and legal penalties (Columns 7–9), respectively. See Appendix B for the full questions. We report the mean and standard deviation of the dependent variable for the “Active” narrative structure. Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

a reduced salience of police killings as a social issue when obfuscatory language is used in the story. These results also echo those from our field study, presented in Section 5.1, in which we find less support for police funding just after more obfuscatory media coverage of police killings.

Table 11: Online experiment: narrative structure and perceptions of policing, for all cases

	Donation reform			Yearly police killings		
	(1)	(2)	(3)	(4)	(5)	(6)
Any obfuscation	-2.1 (1.5)			-89.9* (53.9)		
No explicit agent		-2.9* (1.5)			-134.1** (56.3)	
Passive		-0.7 (1.7)	-0.7 (1.7)		-2.0 (67.6)	-2.0 (67.6)
No agent + nominalization			-2.1 (1.8)			-163.3*** (63.2)
Intransitive + nominalization			-3.6** (1.8)			-104.3 (65.5)
Mean Dep. Var.	67.47	67.47	67.47	1352.56	1352.56	1352.56
SD Dep. Var.	31.13	31.13	31.13	1164.26	1164.26	1164.26
N	2402	2402	2402	2402	2402	2402

Notes: The outcome in columns 1-3 is number of dollars out of their 100 dollar donation that respondents want to give to an organization supporting police reform. The outcome in columns 4-6 the estimated number of police killings each year. See Appendix B for the full questions. We report the mean and standard deviation of the dependent variable for the Active narrative structure. Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Informational Content. The use of obfuscatory sentence structures could potentially affect perceptions and judgments about the police killing described in our experiment in multiple ways. A natural distinction is whether the fundamental information about the incident that participants take away from the story is altered or whether the effects estimated above instead reflect the differential impact of the same information.²⁷ The latter could result, for example, if more active language leads readers to develop a more vivid picture of what happened or evokes a stronger emotional response.

To shed some light on these possible mechanisms, we explore whether differences in narrative structure affect how people recall and/or retell the story. At the end of the survey, we asked participants to retell the news story that they had read in their own words. For half of our sample, two research assistants, who were not informed about the treatment arms to which the responses corresponded, coded the content of the

²⁷Our conjecture that viewers may take away different information content is motivated by the fact that the information transmitted across the four treatment arms of the experiment is not equivalent. While the *Active* (“the police officer killed the man”) and *Passive* (“the man was killed by the police officer”) transmit clearly who the agent and the patient are and what the action is (albeit in a different order), the agent is not clear in the *No Agent* syntax (“man was killed in an officer-involved shooting”), and neither the agent nor the causal action are fully specified in the *Intransitive* syntax (“the man died in an officer-involved shooting”).

sentences to capture whether the person explicitly said that the police officer had killed or shot a person.²⁸

Table 12: Online experiment: participants retelling of the story.

	Explicit police shooting (1)	Active voice (2)	No explicit agent (3)
Passive	-0.01 (0.02)	-0.14*** (0.04)	0.02 (0.02)
No Agent + Nominalization	-0.02 (0.02)	-0.27*** (0.04)	0.12*** (0.03)
Intransitive + Nominalization	-0.16*** (0.03)	-0.35*** (0.04)	0.22*** (0.03)
Mean Dep. Var.	0.92	0.73	0.07
SD Dep. Var.	0.26	0.44	0.26
N	1198	1198	1198

Notes: We asked participants to write what they recalled of the news story. We hand-classified the text to capture whether the person explicitly said that at police officer shot or killed the person (column 1); whether they used an active voice (column 2); and whether there was no agent in the retell of the killing (column 3). We report the mean and standard deviation of the dependent variable for the Active narrative structure. Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

The results are presented in Column 1 in Table 12. In the *Active* treatment arm, the vast majority of the respondents explicitly mentioned the role of the police officer in the killing (92 percent) in their retelling of the story. Interestingly, this varied little across the *Active*, *Passive*, and *No Agent* treatments, all of which use a form of the word “kill,” but participants who read the *Intransitive* version of the story were 16 percentage points less likely to identify the police officer as having killed or shot someone. This suggests that the use of *Intransitive* verbs may fundamentally alter the information that viewers take away from the story. Moreover, it suggests that the impact of the *Passive* and *No Agent/Nominalization* sentence structures in the experiment is not likely related to participants’ understanding of the fundamental information in the story but is likely driven by the response to this information that the narrative structure evoked.

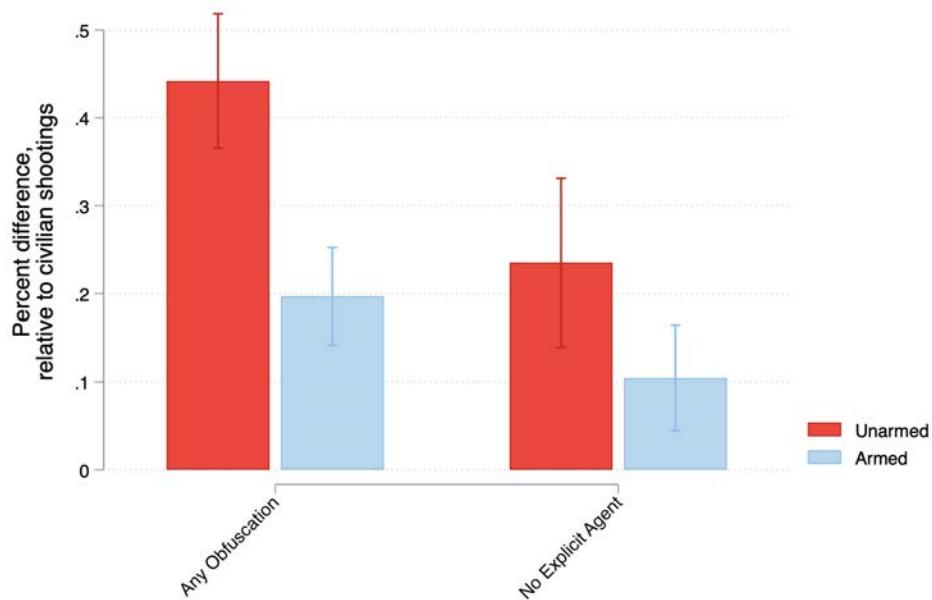
Language Transmission. We use the participants’ open-ended retelling of the story to explore a final question related to the impact of narrative structure: Is obfuscatory language getting retransmitted as information is shared with others? In evaluating the survey responses to answer this question, we measure the respondents’ use of obfuscatory sentence structures, presenting these results in Columns 2 and 3 of Table 12. Strikingly, the estimates reveal that people tend to mimic the sentence structures they read when retelling the story themselves several minutes later. As shown in Column 2, they are significantly less likely to use the active voice following treatments using any of the obfuscatory sentence structures. The effect sizes are increasing with our ex-ante expected gradient of obfuscation: i.e., the largest decrease in the use of active voice is generated in the *Intransitive* arm (-35 percentage points), followed by the *No Agent* (-27 percentage points) and *Passive* (-14 percentage points). Similarly, Column 3 shows that, when there was *No explicit*

²⁸The two coders agreed on the categorization of 85 percent of the sentences. Our team recoded those that they disagreed on.

agent in the story itself, respondents tend to also use sentence structures that do not explicitly acknowledge the police officer as a cause of the killing (+12 percentage points for *No Agent* and +22 percentage points for *Intransitive*). Notice that these effects arise despite the fact that the phrase “officer-involved shooting” occurs twice within each of these stories, suggesting that many participants do not interpret this phrase to automatically imply that the officer was the shooter. While the implications of this analysis are certainly limited by the short-term nature of the recall exercise in the experiment, the results suggest that there may be broader spillover effects of obfuscatory language structures. They suggest, in particular, that the media use of obfuscation may not only affect viewers directly, but may also shape the information they subsequently pass on to others.

Overall, our online experiment shows that obfuscation impacts the interpretation of the situation, shapes broader perceptions of policy issues, and influences the subsequent retelling of a story.²⁹ However, there is variation in the importance of obfuscation. Consistently across contexts, nominalization and not including an explicit agent influences all of our measures of perception of a police killing, and its social consequences. Using a passive instead of an active voice also matters, but less strongly, and not in all cases (for example, not when we specify that the victim was armed). In other words, the more tortured the language choice, the greater the impact it has on what people retain from a news story.

Figure 3: Heterogeneity in obfuscation: victim allegedly armed or not



This figure plots the coefficients for obfuscation in police killings with respect to civilian killings, depending on whether the victim was reportedly armed or unarmed, as measured in the MPV data. These estimates are based on the coefficients presented in Columns 1 and 2 of Tables 13. Our sample includes incidents and news stories where a suspect was identified for civilian killings. Data sources: GVA, MPV and News Exposure.

²⁹Our experiment focuses on the influence of narrative structures on perceptions of policing within the present media landscape. We acknowledge that respondents might have developed specific interpretations due to the prevailing narrative styles. There is a possibility that if narrative structures were to undergo more systematic changes, different inferences might arise.

5.3 Differential Obfuscation

Our experimental results showed that for any given case of police killings, obfuscation is most effective in changing attitudes of the audience when the victim is not reported to have a weapon. We explore the extent of differential obfuscation in the news coverage across this dimension, and test if there is higher obfuscation in the cases in which it matters most.

Table 13: Obfuscation in News of Police Killings by Whether Victim Was Allegedly Armed

Dimension of Obfuscation:	All Sentences		1st Sentence	
	Any (1)	No Explicit Agent (2)	Any (3)	No Explicit Agent (4)
Panel A: Unknown or no reported weapon for victim of police shooting				
Police Killing	0.130*** (0.011)	0.058*** (0.012)	0.176*** (0.013)	0.078*** (0.015)
Observations	201,781	201,781	79,293	79,293
Mean Civ.Shoot.	0.2935	0.2481	0.2815	0.2072
Panel B: Reported weapon for victim of police shooting				
Police Killing	0.058*** (0.008)	0.026*** (0.008)	0.104*** (0.010)	0.067*** (0.008)
Observations	411,453	411,453	158,591	158,591
Mean Civ.Shoot.	0.2935	0.2481	0.2815	0.2072
Controls	Story Controls+DMA FE			

Notes: This table examines differential obfuscation in stories about police killings relative to stories about civilian killings by whether the victim was allegedly armed, as defined by MPV. Panel A is for cases where the victim was allegedly unarmed, while Panel B is for cases where the victim was allegedly armed. The analysis is at the sentence level. Our sample includes incidents and news stories where a suspect was identified for civilian killings. We define “Any obfuscation” as a sentence with passive voice, with no agent, with intransitive verbs, or with nominalizations. We define “No explicit agent” as a sentence that has no agent or has intransitive verbs. See Section 2 for more details. Standard errors clustered by subject (*: $p < 0.1$, **: $p < 0.05$, ***: $p < 0.01$). Source: News Exposure.

The MPV data includes a variable that describes whether the victim was allegedly armed or not. Using this variable, we break out stories about police killings depending on whether the victim was allegedly armed or not. The results are presented in Table 13 and in Figure 3. We find that there is indeed more obfuscation for cases where the victim was not allegedly armed. Comparing the coefficients in Panel A to those in Panel B in Table 13 shows that, in fact, there is twice as much obfuscation when the victim is unarmed than allegedly armed. For all the pairwise comparison across the coefficients, the estimates for ‘no weapon’ are larger and statistically different from those with ‘weapon’.³⁰ These are likely to be the kinds of cases where obfuscation might be especially favorable for police officers, given results from our lab experiments.

³⁰The p-value is less than 0.01 for all cases, except *No explicit agent*, which has a p-value of 0.03

Beyond the question of whether the victim was armed, various factors can influence perceptions of situations, making them more or less favorable for police officers. There could be a substitution effect in determining officer responsibility based on either narrative structures or incident characteristics. In this case, other factors might similarly diminish perceived officer culpability, thereby reducing the necessity for obfuscation. The following features, measurable in the MPV data, can contribute to a diminished sense of officer responsibility: if the victim was fleeing, if a body-worn camera was absent, and if the event went viral. Appendix Figure C.4 presents estimates for these subgroups. Generally speaking, our findings suggest a discernible pattern of substitution between obfuscation and other elements that diminish relative officer responsibility.

The exact drivers of differential media coverage remain uncertain; it is possible that there are other observed or unobserved differences across incidents where the victim was allegedly armed or not, allegedly fleeing or not, etc. However, the results presented in this section suggest that obfuscation may be strategically employed to mitigate perceived responsibility in cases that could generate heightened scrutiny of police officers. In the paper's final section, we offer a conceptual framework to help weigh these different elements.

6 What Drives Media Obfuscation of Police Killings? A Conceptual Framework

The analysis presented in the previous two sections of the paper documents the heightened presence of obfuscation in media coverage of police killings relative to reporting on homicides with civilian perpetrators and indicates that these obfuscatory narrative structures significantly impact public perceptions about police killings. This evidence raises a natural follow-up question: What drives the greater use of obfuscation in media coverage of police killings?

In this last section of our paper, we propose a conceptual framework that outlines potential mechanisms driving the use of obfuscatory language and provide some suggestive empirical tests. Following the literature (Gentzkow *et al.*, 2015), we start by distinguishing demand- and supply-driven motivations for obfuscation in the coverage of police killings. Within the supply-driven channels, we further differentiate between mechanisms related to a desire by news stations to shape how consumers understand the news, potentially influenced, for example, by the political attitudes of station owners, and those reflecting the influence of the primary source of information about matters of crime and policing, namely, local police departments.

6.1 Demand for Obfuscation

Political affiliation is one of the strongest predictors of attitudes toward police, with Republicans, for instance, exhibiting more trust in the police force and more favorable perceptions of how well police conduct their jobs (Jones, 2013; Ekins, 2016; Brown, 2017). Thus, if media obfuscation is driven by a desire to support or avoid challenging their audience's views, we might expect to see greater obfuscation for police killings in Republican-leaning media markets. Appendix Figure C.5 reports the results of our main analysis of news stories stratified by quartile of the Republican vote share in the 2016 presidential election in the media market. We find no evidence of greater obfuscation of police killings in Republican-leaning markets. If anything, we find that obfuscation is more common in Democrat-leaning media markets. Although the

differences are not statistically significant, there is approximately 50 percent greater differential obfuscation in the bottom versus the top quartile of markets ranked by Republican vote share.

Appendix Figure C.5 reports results of a second demand-related analysis. In particular, we examine whether there is any correlation between obfuscation in police killings and television station advertising revenue. An association in this case might arise if obfuscating police killings helps stations keep viewers. We split our sample by whether a story was broadcast on a program with advertising revenue above or below the median value within our sample. Again, we see no difference in the effects of obfuscation across levels of advertising revenue.

The absence of systematic heterogeneity in obfuscation across media markets based on political leaning and the lack of correlation between obfuscation and advertising revenue suggest that demand-side considerations likely play a minor role in narrative structure choices for television newscasts.

6.2 Supply of Obfuscation: Station Ownership and Political Leaning

Next, we explore a first supply-side channel for obfuscation: television station ownership and political leaning. We consider three measures related to political leaning. First, we examine the influence of television station ownership groups on the use of obfuscatory narrative structures. Previous studies have demonstrated that station ownership has an impact on content, including both the topics covered and political slant of the news (DellaVigna & Kaplan, 2007; Martin & Yurukoglu, 2017; Ash & Poyker, 2023; Cagé *et al.*, 2022). Building on the research of Martin & McCrain (2019) and Mastrorocco & Ornaghi (2020), we focus, in particular, on the Sinclair Broadcast Group, recognized as a politically conservative media entity with an extensive portfolio of local news stations.³¹ To explore whether Sinclair ownership influences narrative choices, we examine whether there are differences in obfuscation on the basis of Sinclair ownership status across three categories: (i) stations never owned by Sinclair, (ii) those owned by Sinclair when the relevant story was aired, and (iii) those owned by Sinclair eventually but not at the time of airing. We also use an alternative indicator for conservative ownership: the station's ideological slant, as measured by Martin & McCrain (2019). Appendix Figure C.5 shows similar levels of obfuscation across these categories of Sinclair ownership status and levels of slant, suggesting little direct impact of conservative station ownership or political leaning on the use of obfuscatory narrative structures.

Finally, we directly compare obfuscation across TV stations. First, in Appendix Figure C.5, we break down our results by national versus local TV stations. Then, in Appendix Figure C.6, we present estimates for major national TV stations, ordered by viewership composition from liberal to conservative, following the classification by Mitchell *et al.* (2014). We see no discernible patterns in terms of differential obfuscation either between local and national channels or within national channels across news outlets.

The absence of patterns in heterogeneity across all of these specifications suggests that the political slant of television stations and their owners is likely not a primary driver of the media use of obfuscatory language for police killings: it is common across all outlet types.

³¹These papers show that Sinclair ownership not only affects the topics covered and political bias but also leads to a decline in violent crime clearance rates due to reduced coverage of local police activities.

6.3 Supply of Obfuscation: Primary Sources of Information

Beyond television ownership, another potential supply-side driver of obfuscation lies in the relationship between the media and local police departments, which could influence storytelling. Extensive research in sociology, communications, and criminology has highlighted how law enforcement agencies disseminate information and images to the media and the communication strategies that they employ to enhance their public image (Mawby, 2010, 2013; Chermak & Weiss, 2005; Colbran, 2018).³² Relatedly, Baron (2006) puts forth a theory on media bias stemming from private information accessed by journalists—police departments could be such a source. This body of work concentrates on the kinds of information that police provide to news outlets, such as details regarding crimes or ongoing investigations, rather than on language or grammatical choices.

In this subsection, we study whether the language used by police departments in their accounts of events surrounding police killings influences the narrative structures used in news media coverage of the same incidents. To explore this, we first collect data from the documents that many police departments issue after incidents in which police firearms are discharged. These so-called OIS reports (for “officer-involved shooting”) provide an official police department account of the circumstances surrounding such incidents. Matching these OIS reports on police killings to news coverage of the same incidents, we estimate a series of regressions to test whether the use of obfuscatory structures in police reports is reflected in local news media coverage.

To compile OIS records, we first rank police departments based on their size. Focusing on the top 25 police departments, we investigate whether any OIS documentation exists on the respective police department websites. We identify such records for six departments: Los Angeles, Houston, Philadelphia, Dallas, San Francisco, and Louisville.³³ Other departments either lacked an OIS-specific page or their OIS data do not include narratives (as exemplified by the NYPD). For each of the identified cities, we gather information on the date, city, and narrative content of each incident and matched these events to the Mapping Police Violence dataset. We successfully locate official police narratives for 105 individuals, representing approximately half of the police killings in these cities during our sample period. Our main News Exposure dataset has stories for 60 of these incidents, and they are associated with 1,366 television news stories and 2,840 distinct sentences. For all matched stories, we compute the obfuscation in the police narratives using the same methodology as we use for television coverage.

Table 14 presents estimates of regressions at the sentence level that relate the use of obfuscation in media coverage of police killings to obfuscation in the official police OIS reports. The estimates reported in the first column of each panel, which include both year and police department fixed effects, indicate that a higher degree of obfuscation in the relevant police reports corresponds to more obfuscation in media coverage of the same incident – i.e., how the police department describes an incident is highly correlated with its coverage on television. Importantly, because police department fixed effects are included in these specifications, the results imply that the narrative structures used *by the same department* across different police killings is directly mirrored in the associated media coverage of these events. These results suggest that the news media may rely fairly directly on official department narratives for their own reporting.

³²More recently, scholars have studied how citizen journalism can change these dynamics (Greer & McLaughlin, 2010; Deneff *et al.*, 2013) and how police departments use social media to shape perceptions (Heverin & Zach, 2010; Grunwald *et al.*, 2022).

³³Phoenix also has such documentation, but only after 2019, which falls outside our study period.

Table 14: Correlation Between Police Obfuscation and Media Obfuscation

	Sentence		Story	
	(1)	(2)	(3)	(4)
Panel A: Any Obfuscation				
	(Mean=0.3665)		(Mean=0.6091)	
Police Obf (Any)	0.133** (0.065)		0.102 (0.079)	
Police Obf (Any) $\times 1(\text{SameDMA} = 0)$		0.104* (0.059)		0.022 (0.075)
Police Obf (Any) $\times 1(\text{SameDMA} = 1)$		0.197** (0.079)		0.265** (0.123)
Panel B: No Explicit Agent				
	(Mean=0.2954)		(Mean=0.5249)	
Police Obf (No-Explicit Agent)	0.216*** (0.068)		0.253** (0.096)	
Police Obf (No-Explicit Agent) $\times 1(\text{SameDMA} = 0)$		0.184** (0.070)		0.169* (0.097)
Police Obf (No-Explicit Agent) $\times 1(\text{SameDMA} = 1)$		0.273*** (0.083)		0.409*** (0.130)
Police Department FE	X	X	X	X
Year FE	X	X	X	X
DMA FE	X	X	X	X
Observations	2,840	2,840	1,366	1,366

Notes: Sample matches all OIS Police statements across 6 police departments for which we match the obfuscation in the main sample from the MPV. We obtain 1,366 stories for 60 subjects. Police obfuscation is measured as the average obfuscation in the statement on sentences about the killing. When obfuscation by the police is interacted, the dummy for SameDMA is always included. Standard errors clustered by subject (*: $p < 0.1$, **: $p < 0.05$, ***: $p < 0.01$). Source: News Exposure

The second column of each panel interacts the measure of police obfuscation with an indicator for whether the television news station is in the same media market as the police department. The point estimates are greater in magnitude for television stations in the same market, consistent with the idea that local news stations may be especially likely to echo the language used by local police departments. The estimates are also generally positive and significant for television stations outside the local market, suggesting that police department accounts of events may ultimately influence how a specific incident is described in broader media circles.

Although not definitive, these results suggest that how local police departments represent their own actions likely constitutes an important mechanism influencing the prevalence of obfuscation in television stories.

Finally, it is worth noting that many of the results from the heterogeneity analysis presented earlier in the paper in Section 5.3 are also consistent with obfuscation originating from police departments as a way of defusing potential blame or criticism in cases that may be viewed more negatively by the public. In par-

ticular, we saw in Table 13 and in Appendix Figure C.4 that there is more differential obfuscation for police killings in which the victim did not have a weapon, when the victim was not fleeing, for more viral stories, and when there is body-worn camera footage. There are many reasons why this may be the case, but one potential explanation is that police departments may be especially likely to obfuscate in cases that might be viewed as more likely to lead to public outcry, given that high-profile incidents of police killings might change oversight, reform, or support for police.³⁴ Likewise, although not statistically significant, the higher obfuscation in Democrat-leaning media markets might indicate that audiences in these markets are more willing to hold the police accountable following a police killing, thus leading police departments in these areas to use greater obfuscation.

While not conclusive, our analyses suggest that neither demand factors nor TV channel-level supply explain the observed differential obfuscation in media coverage of police killings. Instead, our findings point to police departments themselves as a potential origin of obfuscation. Particularly in smaller local newsrooms constrained by limited resources (George & Waldfoegel, 2006), the necessity to form and uphold relationships with police departments or the tendency to draw directly from the language used by police in official accounts (Cagé *et al.*, 2020) could plausibly serve as a driving force behind obscuring responsibility following police killings.

7 Conclusion

The main aim of this paper is to provide new empirical and experimental evidence on two interrelated questions that have received a great deal of attention in recent years: whether there is systematic obfuscation of responsibility in media coverage of police killings and, if so, whether this matters for perceptions of both the incident in question and the potential harms from policing more generally (Cheng, 2021). To answer these questions, we collected comprehensive data on police killings, civilian homicides, and television news coverage in the United States covering the period 2013–19; we collected survey evidence on support for police funding; and we conducted an online experiment in which we varied the narrative structure used to report a story about a police killing.

The results of our analyses provide clear and robust evidence on both questions. First, we document that there is more obfuscation when an officer was responsible for the killing than in coverage of civilian homicides for which the media faced a similar choice of language. The use of obfuscatory language is especially common in the story’s lead sentence, which is most salient to viewers. Second, we show that obfuscation matters. Survey respondents are less prone to support police funding in the aftermath of a police killing if they were exposed to more obfuscatory coverage of this event. In our experiment, we find that respondents’ assessment of the situation varies with the degree of obfuscation. They are less likely to think that the officer is morally responsible and to ask for penalties when there is obfuscation—all the more so when we do not specify that the civilian was armed. Prompted by the experimental results, we ask a third question: whether there is differential obfuscation in cases where it might especially benefit the police—e.g., cases when the victim was not armed or when body camera footage is available. We find a doubling of the use

³⁴See Prendergast (2001); Shi (2009); Cunningham & Gillezeau (2019); Rivera & Ba (2018); Devi & Fryer (2020); Ang *et al.* (2021) and Premkumar (2019).

of obfuscation in these cases. Our findings align with research in social psychology indicating that positioning oneself as a moral patient rather than a moral agent restricts the perception of moral responsibility (Gray & Wegner, 2009; Waytz *et al.*, 2010; Gray *et al.*, 2012). Our paper indicates that even subtle alterations in narrative structures have the potential to influence perceptions of moral agency—a promising area for further investigation.

In the last part of our paper, we ask what drives obfuscation. We consider several potential motivations that television news channels and newspapers might have for using obfuscatory language. We first distinguish demand- and supply-driven motivations and, then, within the supply-driven mechanisms, we distinguish mechanisms related to a desire to shape how consumers understand the news (for example, due to the political attitudes of station owners), and to an echoing of the primary source of information (e.g., the local police department). Our evidence is most consistent with the last channel. This is important since our results indicate that the narrative structures employed by media outlets impact both how the public understands the harms of policing more generally, as well as its support for police accountability and reform. These broader effects of language are important given the growing discussion on policy changes and reforms that society might implement to improve public safety (Akbar, 2020; Bursztyn *et al.*, 2022b) in light of the significant negative externalities of police violence on cities and individuals (DiPasquale & Glaeser, 1998; Cunningham & Gillezeau, 2019; Ang, 2020; González & Prem, 2022). However, many questions on the origins of obfuscation remain. Documenting when terms such as “officer-involved” began to be widely used by the media and whether, for example, the use of obfuscatory language in police press conferences and press releases spills over directly to the language used by the local media are promising avenues for future research.

Finally, while our analysis focuses on the semantic structure of language in the context of the media’s coverage of police killings, our approach, studying the scale and consequences of obfuscation, offers a practical and widely applicable template for other topics covered in news outlets (both television outlets, as in our context, but also newspapers or radio) or social media. For example, how do media cover different forms of interpersonal violence, and how does this influence perceptions of responsibility and support for policy change? Beyond the crime and criminal justice space, how stories are told might also matter for many other economic and social issues, such as income inequality, immigration, climate change and health. This analysis could also easily be extended to study the language structures used by a wider set of actors—such as political speech, corporate messaging, and social media influencers—broadening our focus on traditional media.

References

- Abdul-Razzak, Nour, & Hallberg, Kelly. 2022. *Unpacking the Promise and Limitations of Behavioral Health Interventions on Interactions with Law Enforcement*. Working Paper.
- Akbar, Amna A. 2020. An Abolitionist Horizon for (Police) Reform. *California Law Review*, **108**, 1781.
- Alesina, Alberto, Ferroni, Matteo F, & Stantcheva, Stefanie. 2021. *Perceptions of Racial Gaps, their Causes, and Ways to Reduce Them*. Tech. rept.
- Andre, Peter, Haaland, Ingar, Roth, Chrisotpher, & Wohlfart, Johannes. 2021. Inflation narratives.
- Ang, Desmond. 2020. The Effects of Police Violence on Inner-City Students. *The Quarterly Journal of Economics*, 09.

- Ang, Desmond, Bencsik, Panka, Bruhn, Jesse, & Derenoncourt, Ellora. 2021. *Police Violence Reduces Civilian Cooperation and Engagement with Law Enforcement*. Working Paper.
- AP. 2020. *The Associated Press Stylebook 2020-2022*. The Associated Press.
- Ash, Elliott, & Poyker, Michael. 2023. Conservative news media and criminal justice: Evidence from exposure to fox news channel. *Columbia Business School Research Paper*.
- Ash, Elliott, Gauthier, Germain, & Widmer, Philine. 2023. Relatio: Text semantics capture political and economic narratives. *Political Analysis*, 1–18.
- Baker, Scott R., Bloom, Nicholas, & Davis, Steven J. 2016. Measuring Economic Policy Uncertainty. *The Quarterly Journal of Economics*, **131**(4), 1593–1636.
- Balko, Radley. 2014. The curious grammar of police shootings. *The Washington Post*.
- Baron, David P. 2006. Persistent media bias. *Journal of Public Economics*, **90**(1-2), 1–36.
- Bénabou, Falk, & Tirole. 2020. Narratives, imperatives, and moral persuasion. *University of Bonn, mimeo*.
- Blachor, Devorah. 2020. How to use the Past Exonerative Tense to Uphold White Supremacy. *McSweeney's*.
- Bohner, Gerd. 2001. Writing about rape: Use of the passive voice and other distancing text features as an expression of perceived responsibility of the victim. *British Journal of Social Psychology*, **40**(4), 515–529.
- Bonial, Claire, Babko-Malaya, Olga, Choi, Jinho D, Hwang, Jena, & Palmer, Martha. 2010. Propbank annotation guidelines. *Center for Computational Language and Education Research, CU-Boulder*, **9**.
- Brown, Anna. 2017. Republicans more likely than Democrats to have confidence in police.
- Bursztyn, Leonardo, Egorov, Georgy, Haaland, Ingar, Rao, Aakaash, & Roth, Christopher. 2022a. *Justifying Dissent*. Tech. rept. 141.
- Bursztyn, Leonardo, Rao, Aakaash, Roth, Christopher P, & Yanagizawa-Drott, David H. 2022b. *Misinformation During a Pandemic*. Working Paper. National Bureau of Economic Research.
- Cagé, Julia, Hervé, Nicolas, & Viaud, Marie-Luce. 2020. The production of information in an online world. *The Review of economic studies*, **87**(5), 2126–2164.
- Cagé, Julia, Hengel, Moritz, Hervé, Nicolas, & Urvoy, Camille. 2022. Hosting Media Bias: Evidence from the Universe of French Broadcasts, 2002-2020. *Available at SSRN 4036211*.
- Cagé, Julia. 2020. Media competition, information provision and political participation: Evidence from French local newspapers and elections, 1944–2014. *Journal of Public Economics*, **185**, 104077.
- Campbell, Romaine, & Redpath, Connor. 2022. *Officer Language and Subject Race: A Text Analysis of Police Reports*. Working Paper.
- Chan, Eugene Y, & Maglio, Sam J. 2020. The voice of cognition: Active and passive voice influence distance and construal. *Personality and Social Psychology Bulletin*, **46**(4), 547–558.
- Chen, M Keith. 2013. The effect of language on economic behavior: Evidence from savings rates, health behaviors, and retirement assets. *American Economic Review*, **103**(2), 690–731.
- Cheng, Tony. 2021. Social media, socialization, and pursuing legitimization of police violence*. *Criminology*, **59**(3), 391–418.
- Chermak, Steven, & Weiss, Alexander. 2005. Maintaining legitimacy using external communication strategies: An analysis of police-media relations. *Journal of criminal justice*, **33**(5), 501–512.
- Chestnut, Eleanor K, & Markman, Ellen M. 2018. “Girls are as good as boys at math” implies that boys are probably better: A study of expressions of gender equality. *Cognitive science*, **42**(7), 2229–2249.
- Chiang, Chun-Fang, & Knight, Brian. 2011. Media Bias and Influence: Evidence from Newspaper Endorse-

- ments. *The Review of Economic Studies*, **78**(3), 795–820.
- Colbran, Marianne P. 2018. Policing, social media and the new media landscape: can the police and the traditional media ever successfully bypass each other? *Policing and Society*.
- Conner, Andrew, Azrael, Deborah, Lyons, Vivian H, Barber, Catherine, & Miller, Matthew. 2019. Validating the National Violent Death Reporting System as a source of data on fatal shootings of civilians by law enforcement officers. *American journal of public health*, **109**(4), 578–584.
- Couttenier, Mathieu, Hatte, Sophie, Thoenig, Mathias, & Vlachos, Stephanos. 2021. Anti-Muslim Voting and Media Coverage of Immigrant Crimes. *The Review of Economics and Statistics*, 1–33.
- Cunningham, Jamein P., & Gillezeau, Rob. 2019. Don't Shoot! The Impact of Historical African American Protest on Police Killings of Civilians. *Journal of Quantitative Criminology*, December.
- Dahl, Gordon, & DellaVigna, Stefano. 2009. Does movie violence increase violent crime? *The Quarterly Journal of Economics*, **124**(2), 677–734.
- De Freitas, Julian, DeScioli, Peter, Nemirow, Jason, Massenkoff, Maxim, & Pinker, Steven. 2017. Kill or die: Moral judgment alters linguistic coding of causality. *J. Exp. Psychol. Learn. Mem. Cogn.*, **43**(8), 1173–1182.
- DellaVigna, Stefano, & Kaplan, Ethan. 2007. The Fox News Effect: Media Bias and Voting*. *The Quarterly Journal of Economics*, **122**(3), 1187–1234.
- DellaVigna, Stefano, & La Ferrara, Eliana. 2015. Economic and social impacts of the media. *Pages 723–768 of: Handbook of media economics*, vol. 1. Elsevier.
- DellaVigna, Stefano, Enikolopov, Ruben, Mironova, Vera, Petrova, Maria, & Zhuravskaya, Ekaterina. 2014. Cross-Border Media and Nationalism: Evidence from Serbian Radio in Croatia. *American Economic Journal: Applied Economics*, **6**(3), 103–32.
- Denef, Sebastian, Bayerl, Petra S, & Kaptein, Nico A. 2013. Social media and the police: Tweeting practices of British police forces during the August 2011 riots. *Pages 3471–3480 of: proceedings of the SIGCHI conference on human factors in computing systems*.
- Devi, Tanaya, & Fryer, Roland G, Jr. 2020 (June). *Policing the Police: The Impact of "Pattern-or-Practice" Investigations on Crime*. Working Paper 27324. National Bureau of Economic Research. Series: Working Paper Series.
- DiPasquale, Denise, & Glaeser, Edward L. 1998. The Los Angeles Riot and the Economics of Urban Unrest. *Journal of Urban Economics*, **43**(1), 52–78.
- Dixon, Travis L, & Williams, Charlotte L. 2015. The changing misrepresentation of race and crime on network and cable news. *Journal of Communication*, **65**(1), 24–39.
- Djourelouva, Milena. 2023. Persuasion through Slanted Language: Evidence from the Media Coverage of Immigration. *American Economic Review*, **113**(3), 800–835.
- Durante, Ruben, & Zhuravskaya, Ekaterina. 2018. Attack when the world is not watching? US news and the Israeli-Palestinian conflict. *Journal of Political Economy*, **126**(3), 1085–1133.
- Duxbury, Scott W, Frizzell, Laura C, & Lindsay, Sadé L. 2018. Mental illness, the media, and the moral politics of mass violence: The role of race in mass shootings coverage. *Journal of Research in Crime and Delinquency*, **55**(6), 766–797.
- Eisensee, Thomas, & Strömberg, David. 2007. News Droughts, News Floods, and U. S. Disaster Relief. *The Quarterly Journal of Economics*, **122**(2), 693–728.
- Ekins, Emily E. 2016. Policing in America: Understanding public attitudes toward the police. Results from a

- national survey. *Results from a National Survey (December 7, 2016)*.
- Engle, Robert F, Giglio, Stefano, Kelly, Bryan, Lee, Heebum, & Stroebel, Johannes. 2020. Hedging Climate Change News. *The Review of Financial Studies*, **33**(3), 1184–1216.
- Enikolopov, Ruben, Petrova, Maria, & Zhuravskaya, Ekaterina. 2011. Media and political persuasion: Evidence from Russia. *American Economic Review*, **101**(7), 3253–85.
- Frazer, Alexandra K, & Miller, Michelle D. 2009. Double standards in sentence structure: Passive voice in narratives describing domestic violence. *Journal of language and social psychology*, **28**(1), 62–71.
- Gardner, Matt, Grus, Joel, Neumann, Mark, Tafjord, Oyvind, Dasigi, Pradeep, Liu, Nelson F., Peters, Matthew, Schmitz, Michael, & Zettlemoyer, Luke S. 2017. AllenNLP: A Deep Semantic Natural Language Processing Platform.
- Gay, Victor, Hicks, Daniel L, Santacreu-Vasut, Estefania, & Shoham, Amir. 2018. Decomposing culture: an analysis of gender, language, and labor supply in the household. *Review of Economics of the Household*, **16**(4), 879–909.
- Genzckow, Matthew, & Shapiro, Jesse. 2004. Media, Education and Anti-Americanism in the Muslim World. *Journal of Economic Perspectives*, **18**(3), 117–133.
- Genzckow, Matthew, & Shapiro, Jesse. 2008. Preschool television viewing and adolescent test scores: Historical evidence from the Coleman study. *The Quarterly Journal of Economics*, **123**(1), 279–323.
- Genzckow, Matthew, & Shapiro, Jesse. 2010. What drives media slant? Evidence from US daily newspapers. *Econometrica*, **78**(1), 35–71.
- Genzckow, Matthew, Shapiro, Jesse, & Stone, Daniel. 2015. Media bias in the marketplace: Theory. *Pages 623–645 of: Handbook of media economics*, vol. 1. Elsevier.
- Genzckow, Matthew, Kelly, Bryan, & Taddy, Matt. 2019. Text as Data. *Journal of Economic Literature*, **57**(3), 535–74.
- George, Lisa M, & Waldfogel, Joel. 2006. The New York Times and the market for local newspapers. *American Economic Review*, **96**(1), 435–447.
- Giglio, Stefano, Kelly, Bryan, & Stroebel, Johannes. 2021. Climate Finance. *Annual Review of Financial Economics*, **13**(1), 15–36.
- Gilliam, Franklin D, Iyengar, Shanto, Simon, Adam, & Wright, Oliver. 1996. Crime in black and white: The violent, scary world of local news. *Harvard International Journal of press/politics*, **1**(3), 6–23.
- Gilliam, Franklin D Jr, & Iyengar, Shanto. 2000. Prime suspects: The influence of local television news on the viewing public. *American Journal of Political Science*, 560–573.
- González, Felipe, & Prem, Mounu. 2022. Police Violence, Student Protests, and Educational Performance. *The Review of Economics and Statistics*, 1–46.
- Gray, Kurt, & Wegner, Daniel M. 2009. Moral typecasting: divergent perceptions of moral agents and moral patients. *Journal of personality and social psychology*, **96**(3), 505.
- Gray, Kurt, Young, Liane, & Waytz, Adam. 2012. Mind perception is the essence of morality. *Psychological inquiry*, **23**(2), 101–124.
- Greer, Chris, & McLaughlin, Eugene. 2010. We predict a riot? Public order policing, new media environments and the rise of the citizen journalist. *The British Journal of Criminology*, **50**(6), 1041–1059.
- Grimmer, Justin. 2010. A Bayesian Hierarchical Topic Model for Political Texts: Measuring Expressed Agendas in Senate Press Releases. *Political Analysis*, **18**(1), 1–35.

- Grosjean, Pauline, Masera, Federico, & Yousaf, Hasin. 2022. Inflammatory Political Campaigns and Racial Bias in Policing.
- Grunwald, Ben, Nyarko, Julian, & Rappaport, John. 2022. *Police agencies on Facebook overreport on Black suspects*. Working Paper.
- Hassan, Tarek A, Hollander, Stephan, van Lent, Laurence, & Tahoun, Ahmed. 2019. Firm-Level Political Risk: Measurement and Effects*. *The Quarterly Journal of Economics*, **134**(4), 2135–2202.
- Henley, Nancy M, Miller, Michelle, & Beazley, Jo Anne. 1995. Syntax, semantics, and sexual violence: Agency and the passive voice. *Journal of Language and Social Psychology*, **14**(1-2), 60–84.
- Heverin, Thomas, & Zach, Lisl. 2010. Twitter for city police department information sharing. *Proceedings of the American Society for Information Science and Technology*, **47**(1), 1–7.
- Humphries, John Eric, Mader, Nicholas S, Tannenbaum, Daniel I, & van Dijk, Winnie L. 2019 (Aug.). *Does Eviction Cause Poverty? Quasi-Experimental Evidence from Cook County, IL*.
- Imai, Kosuke, & Khanna, Kabir. 2016. Improving ecological inference by predicting individual ethnicity from voter registration records. *Political Analysis*, 263–272.
- Jakiela, Pamela, & Ozier, Owen. 2018. Gendered language. *World Bank Policy Research Working Paper*.
- Jones, Jeffrey. 2013. *Honesty ratings of police, clergy differ most by party*.
- Kazenin, Konstantin I. 2001. *The passive voice*.
- Kearney, Melissa S., & Levine, Phillip B. 2015. Media Influences on Social Outcomes: The Impact of MTV’s 16 and Pregnant on Teen Childbearing. *American Economic Review*, **105**(12), 3597–3632.
- Khanna, K, Imai, K, & Hubert, J. 2017. *Who are You? Bayesian Prediction of Racial Category Using Surname and Geolocation*. Tech. rept. Technical report, The Comprehensive R Archive Network.
- Kuriwaki, Shiro. 2023. *Cumulative CES Common Content*.
- Lee, Kenton, He, Luheng, & Zettlemoyer, L. 2018. Higher-order Coreference Resolution with Coarse-to-fine Inference. *In: NAACL-HLT*.
- Logan, Trevon D., & Myers, Samuel L. Jr. 2022. Symposium: Race and Economic Literature—Introduction. *Journal of Economic Literature*, **60**(2), 375–376.
- Lussos, Rachael Graham, & Fernandez, Lourdes. 2018. Assault and accusation without agents: Verb voice in media narratives of campus sexual assault. *Journal of Mason Graduate Research*, **5**(2), 108–127.
- Martin, Gregory J, & McCrain, Joshua. 2019. Local news and national politics. *American Political Science Review*, **113**(2), 372–384.
- Martin, Gregory J, & Yurukoglu, Ali. 2017. Bias in cable news: Persuasion and polarization. *American Economic Review*, **107**(9), 2565–99.
- Mason, Patrick L., Myers, Samuel L. Jr., & Simms, Margaret. 2022. Racial Isolation and Marginalization of Economic Research on Race and Crime. *Journal of Economic Literature*, **60**(2), 494–526.
- Mastrorocco, Nicola, & Minale, Luigi. 2018. News media and crime perceptions: Evidence from a natural experiment. *Journal of Public Economics*, **165**, 230–255.
- Mastrorocco, Nicola, & Ornaghi, Arianna. 2020. *Who Watches the Watchmen? Local News and Police Behavior in the United States*. Tech. rept.
- Mawby, Rob. 2013. *Policing images*. Willan.
- Mawby, Rob C. 2010. Police corporate communications, crime reporting and the shaping of policing news. *Policing and society*, **20**(1), 124–139.

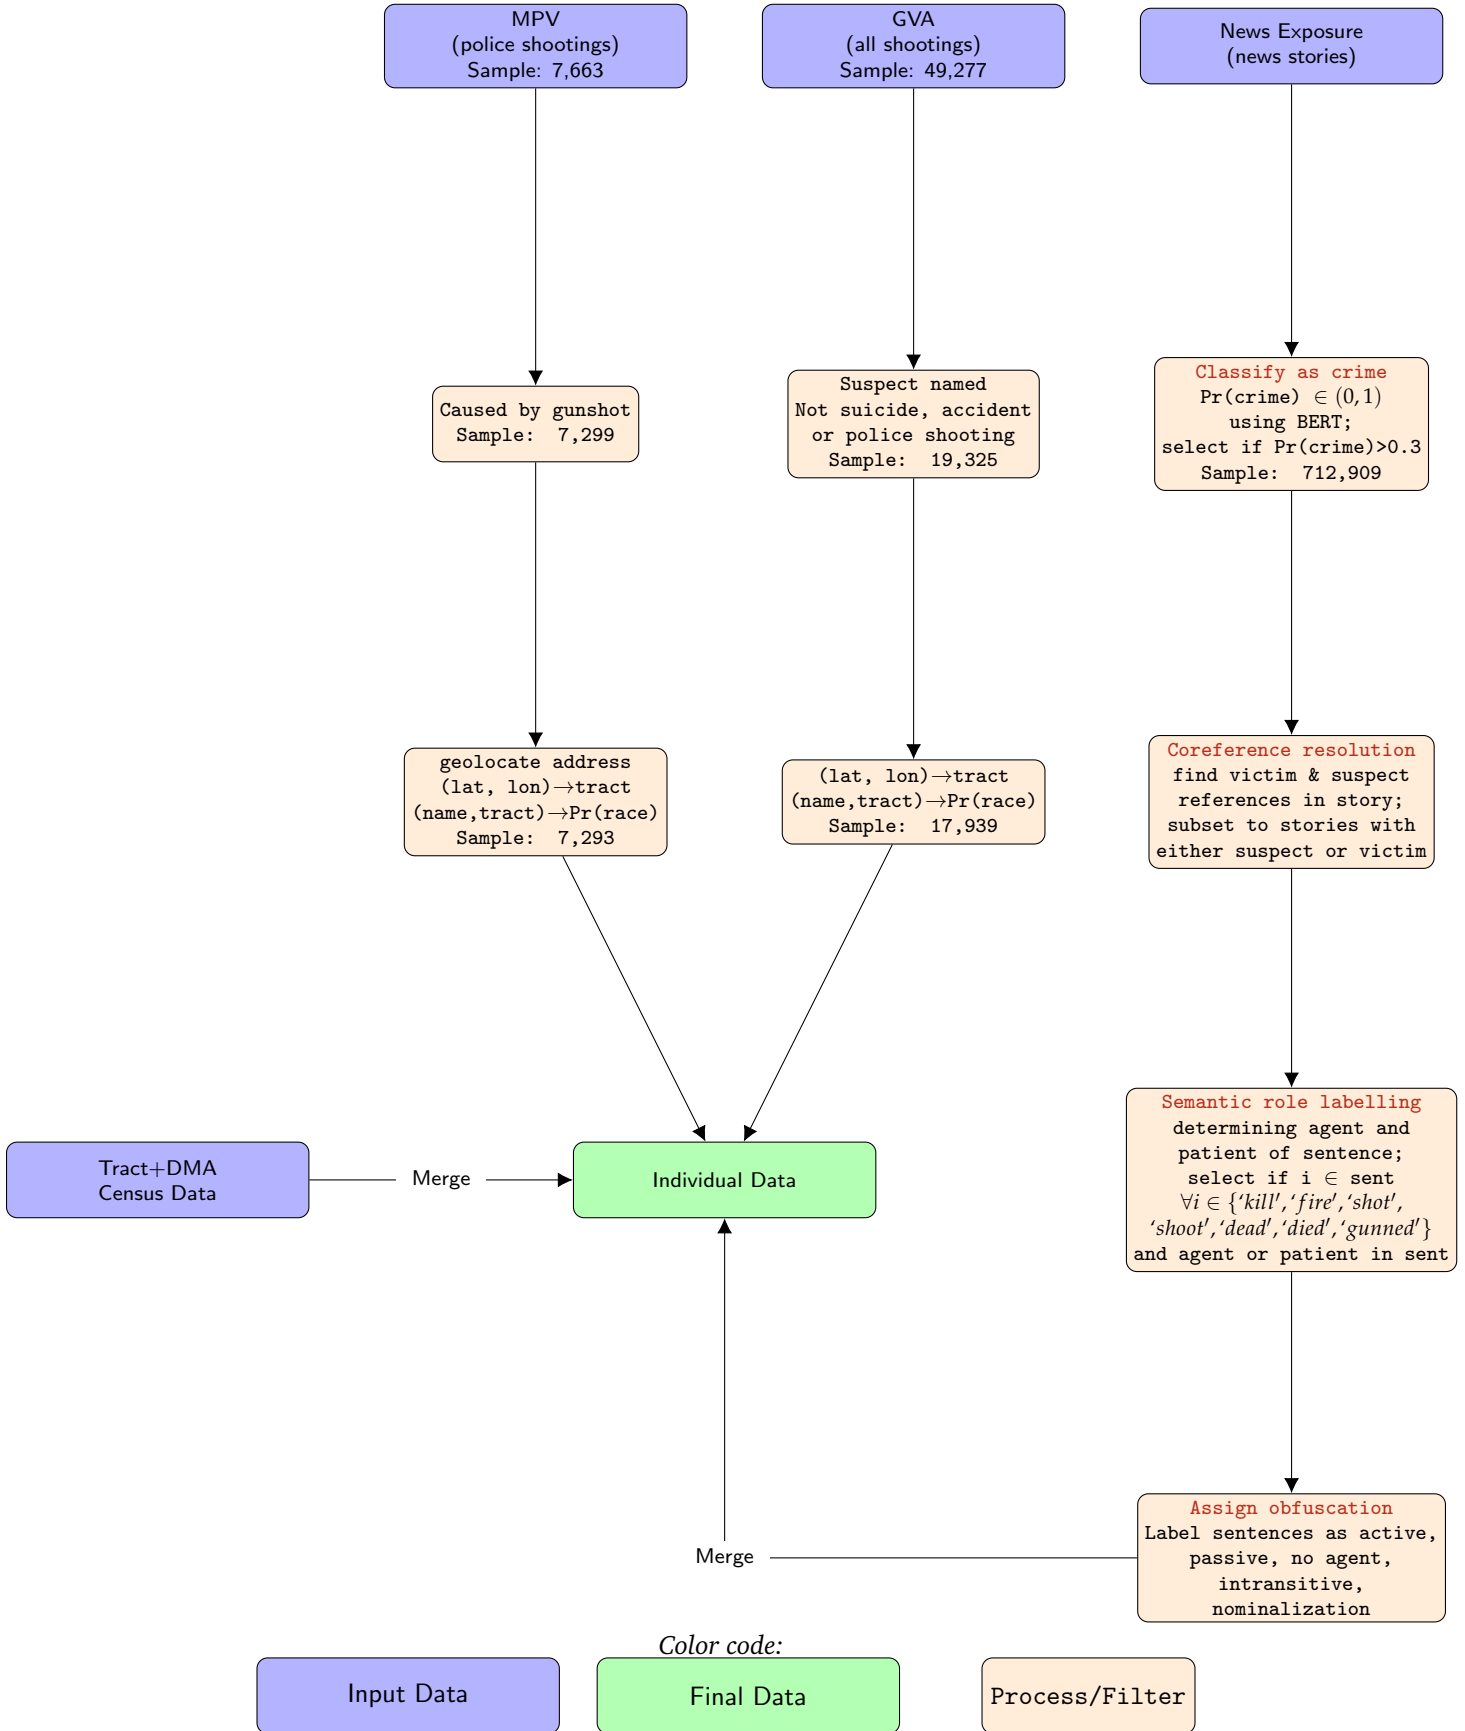
- Misra, Rishabh. 2018 (06). *News Category Dataset*.
- Mitchell, Amy, Gottfried, Jeffrey, Kiley, Jocelyn, & Matsa, Katerina Eva. 2014. Political polarization & media habits.
- Moreno-Medina, Jonathan. 2021. *Local Crime News: Extent, Causes and Consequences*. Tech. rept.
- Niemi, Laura, & Young, Liane. 2016. When and why we see victims as responsible: The impact of ideology on attitudes toward victims. *Personality and social psychology bulletin*, **42**(9), 1227–1242.
- Northcutt Bohmert, Miriam, Allison, Kayla, & Ducate, Caitlin. 2019. “A rape was reported”: construction of crime in a university newspaper. *Feminist Media Studies*, **19**(6), 873–889.
- Philippe, Arnaud, & Ouss, Aurélie. 2018. “No Hatred or Malice, Fear or Affection”: Media and Sentencing. *Journal of Political Economy*, **126**(5), 2134–2178.
- Pinker, Steven. 2007. *The stuff of thought: Language as a window into human nature*. Penguin.
- Premkumar, Deepak. 2019 (Sept.). *Public Scrutiny and Police Effort: Evidence from Arrests and Crime After High-Profile Police Killings*. Tech. rept.
- Prendergast, Canice. 2001. *Selection and oversight in the public sector, with the Los Angeles Police Department as an example*. Tech. rept. National Bureau of Economic Research.
- Rivera, Roman G., & Ba, Bocar A. 2018. The Effect of Police Oversight on Crime and Allegations of Misconduct: Evidence from Chicago. *Working Paper*.
- Sanh, Victor, Debut, Lysandre, Chaumond, Julien, & Wolf, Thomas. 2019. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.
- Schaffner, Brian, Ansolabehere, Stephen, & Shih, Marissa. 2023. *Cooperative Election Study Common Content, 2022*.
- Sharkey, Patrick. 2010. The acute effect of local homicides on children’s cognitive performance. *Proceedings of the National Academy of Sciences*, **107**(26), 11733–11738.
- Shi, Lan. 2009. The limit of oversight in policing: Evidence from the 2001 Cincinnati riot. *Journal of Public Economics*, **93**(1), 99–113.
- Shi, Peng, & Lin, Jimmy. 2019. Simple BERT Models for Relation Extraction and Semantic Role Labeling. *ArXiv*, **abs/1904.05255**.
- Toolan, Michael J. 2013. *Narrative: A critical linguistic introduction*. Routledge.
- Trew, Tony. 1979. Theory and ideology at work. *Pages 94–116 of: Fowler, Roger (ed), Language and control*. Routledge.
- Waytz, Adam, Gray, Kurt, Epley, Nicholas, & Wegner, Daniel M. 2010. Causes and consequences of mind perception. *Trends in cognitive sciences*, **14**(8), 383–388.
- Widmer, Philine, Ash, Elliott, & Galletta, Sergio. 2022. *Media Slant is Contagious*.
- Wolff, Phillip. 2003. Direct causation in the linguistic coding and individuation of causal events. *Cognition*, **88**(1), 1–48.
- Yasmin, Musarat. 2021. Asymmetrical gendered crime reporting and its influence on readers: A case study of Pakistani English newspapers. *Heliyon*, **7**(8), e07862.

Appendix A Data Processing

This appendix describes in detail the processing of the raw data from the MPV, GVA and News Exposure datasets into the final data used for text analysis. First, we present in figure A.1 a flowchart of our data cleaning steps. Second, in Sections A.1–A.4, we describe in turn each of our data processing steps:

1. Identifying news stories about shootings (Section A.1)
2. Coreference resolution (Section A.2)
3. Semantic role labeling (Section A.3)
4. Assigning degree of obfuscation to each sentence and story (Section A.4)

Figure A.1: Data Processing



A.1 Classification of a Story as Crime Related

This section briefly describes how we classify each matched text from the TV caption data as about crime or not. We employ state-of-the-art embedding based on *BERT* (bidirectional encoder representations from transformers).³⁵ *BERT* is very accurate in understanding contextual embeddings of words. This level of accuracy is important in our context because many noncrime stories could include words typically used in crime stories. One example would be the use of the word “shot” to describe a soccer or basketball action instead of the action of a gun.

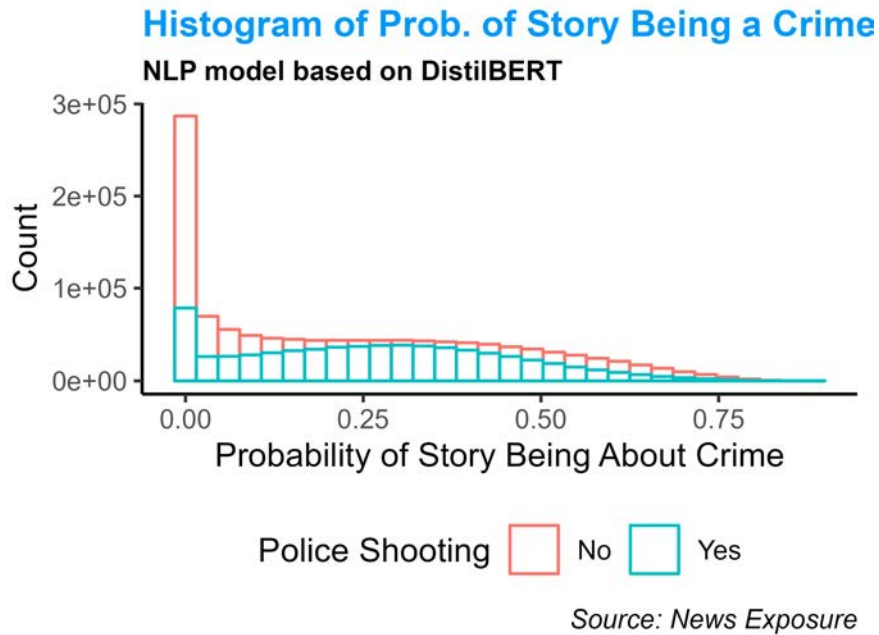


Figure A.2: Histogram of Probability of Story Being about Crime

We employ the pretrained smaller and faster version of *BERT*, called *DistilBERT*, which is another neural network with 6 layers (see [Sanh et al. \(2019\)](#) for more details). We retrain it to classify stories across “**crime**” and “**not crime**” categories. For this latter task, we need a database that accurately labels news stories on crime. We used an already labeled dataset of close to 200,000 short news stories from the *Huffington Post* run between 2012 and 2018 and collected and shared by [Misra \(2018\)](#). These data contain labels for **crime** news (3,405) and 20 other categories. We group together all other categories into the “**noncrime**” category. With this trained model, we estimate the probability that the matched caption is about crime. Manual inspection of a subsample of the predictions reveals that 99 percent of the stories with an estimated probability of being about crime of 0.3 or higher are indeed about crime. This is the threshold that we use for our classification.

³⁵*BERT* is a neural network model of language that has proven to be incredibly successful in a host of tasks in NLP. *BERT* has several technical features, but perhaps the most important is that it trains the model using not only the previous words in a text but also future ones. The standard model allows up to 512 words (or tokens) in a text. This means that for cases in which our stories have more than 512 tokens in a text, we truncate the length of the text at 512 and drop all the remaining tokens. In our sample, 16 percent of all stories have more than 512 tokens. The network has 7 layers, and it works with a type of word embedder model that captures the context in which the word is being used. In 2019, *Google Search* announced that it had started applying *BERT* models for English language search queries within the US.

With this classification, we go from 1,839,853 stories (1,170,573 on nonpolice killings, 669,280 on police killings) to 712,909 stories (426,728 on nonpolice killings, 286,181 on police killings).

Figure A.2 presents the histogram of the probabilities across police and nonpolice killings.

A.2 Coreference Resolution

A.2.1 Basic Coreference Resolution

Coreference resolution is the task of finding all expressions that refer to the same entity in a text. This task is important as in sentences where “him” or “man” is used, we want to know whether the terms refer to the same individual (victim or perpetrator).

We use the model proposed by Lee *et al.* (2018), with the library provided by Gardner *et al.* (2017). The model is also a neural network trained with BERT embedding and has a structure similar to the original BERT model's. Here is an example of how the library works. Consider the following text in our database, which contains a story on a police killing³⁶:

Police say [victim's name] fled before officers shot him, and say officers found the rifle in a nearby apartment. We've learned [victim's name] had several run-ins with the law. In 2007, he pleaded guilty to possessing a firearm. [victim's name] was released from prison more than a month after violating parole. He is married with six children. Their youngest is 5 months old. [victim's brother name] says his brother [victim's first name] was shot in the head, and the bullet went through his cheek and hit his spinal cord. Police are still investigating. No officers were injured in the shooting. Heather Hope 17 news.

The output of the algorithm is a set of “clusters” describing all the tokens recognized to be describing the same entity. Figure A.3 presents the output for this example.

This example is define by the following “spans” (or sets of tokens):³⁷ ['[Victim's Name]', 'HIM', '[Victim's Name]', 'HE', '[Victim's Name]', 'HE', '[Victim's Brother Name]', 'HIS', 'HIS BROTHER [Victim's First Name]', 'HIS', 'HIS'].

Notice that the algorithm does a very good job identifying all the instances in which the *victim* ([Victim's Name]) is referenced in the text. There are only 2 small mistakes with the output: “[Victim's Brother's Name]” and the immediately subsequent “HIS” refer not to the victim but to his brother. In our case, this is not particularly worrisome because if we focus only on sentences describing a victim being killed (as we discuss in Section A.3), we will correctly identify all instances in which the victim was described as being shot.

³⁶We have removed the victim's name from the text.

³⁷To be even more specific: the algorithm provides the spans that capture the same entity with coordinates in terms of characters. So, for the first mention of '[Victim's Name]', the algorithm would show (11,22), which are the position of the first 'A' in '[Victim's Name]' and 22, which is the last 'A' in the same span.

POLICE SAY **0** ABEL GURROLA FLED BEFORE OFFICERS SHOT **0** HIM , AND SAY OFFICERS FOUND **10** THE RIFLE IN A NEARBY APARTMENT . WE'VE LEARNED **0** ABEL GURROLA HAD SEVERAL RUN -
 INS WITH THE LAW . IN 2007 , **0** HE PLEADED GUILTY TO POSSESSING A FIREARM . **0** ABEL GURROLA WAS RELEASED FROM PRISON MORE THAN A MONTH AFTER VIOLATING PAROLE . **0** HE IS
 MARRIED WITH **1** SIX CHILDREN . **1** THEIR YOUNGEST IS 5 MONTHS OLD . **0** GABRIEL GURROLA SAYS **0** HIS BROTHER ABEL WAS **2** SHOT IN THE HEAD , AND THE BULLET WENT
 THROUGH **0** HIS CHEEK AND HIT **0** HIS SPINAL CORD . POLICE ARE STILL INVESTIGATING . NO OFFICERS WERE INJURED IN **2** THE SHOOTING . HEATHER HOPE 17 NEWS . [A10]TERRANCE WAY
 OIS - TV INTRO **3** POLICE IN SACRAMENTO SAY **3** THEY HAVE **4** THE MAN **3** THEY BELIEVE KILLED TWO PEOPLE AT A BAR ON NEW YEAR 'S EVE . AUTHORITIES SAY **4** THE 22 - YEAR - OLD WAS
 AMONG THREE OTHERS WOUNDED LAST NIGHT . **4** HIS INJURIES AND THOSE TO A SECURITY GUARD AND **7** A WOMAN ARE NOT LIFE- THREATENING . POLICE SAY **5** IT STARTED AS AN ARGUMENT
 INSIDE THE BAR THAT ESCALATED AND WHEN **6** AN EMPLOYEE TRIED TO BREAK **5** IT UP , **6** HE WAS SHOT , AS WERE **7** THE WOMAN AND **8** ANOTHER MAN . **8** THAT MAN AND
6 THE EMPLOYEE DIED AT THE SCENE . [A16]SA SHOOTING - VO **9** A MAN CAUGHT NAKED AND WIELDING **12** A SWORD HAS BEEN TAKEN INTO POLICE CUSTODY IN SAN JOSE . OFFICERS TRACKED
9 THE SUSPECT DOWN AFTER SOMEONE CALLED 9 - 1 - 1 SAYING **9** THE MAN HAD **10** AN ASSAULT RIFLE . POLICE SPOTTED **11** **9** THE MAN 'S VEHICLE NEAR AN EXPRESSWAY .
13 OFFICERS SAY **9** THE MAN STOPPED **11** THE CAR AND GOT OUT WEARING NOTHING AND HOLDING **12** A LARGE SAMU SWORD . **13** THEY SAY **9** HE TRIED TO RUN AWAY , BUT OFFICERS
 CAUGHT **9** HIM .

Figure A.3: Example of Coreference Resolution

A.2.2 Using Coreference Resolution to Identify Relevant Stories

As the previous example elucidates, locating the relevant part of the story is key. We use coreference resolution to identify the story that we care about in the following way:

1. Find all coreferences in raw text.
2. Replace all references by “Victim” if any span in the cluster includes the name of the victim.
3. Replace all references by “Perpetrator” (what we call the suspect, not including police) if any span in the cluster includes the name of the suspect.
4. Divide the text into sentences.
5. Define the story as all sentences between the first and last mention of Victim or Perpetrator.

We illustrate the above process with another example.³⁸ Consider this raw text, which includes the story that we are looking for (references to the victim are in **BOLD**) and some other text not about the victim/shooting (in *italics*):

What the proud parents are saying about having their bundle of joy at the start of 20 - 13.
[a6]open 5pm 365 - deko 2. A **BAKERSFIELD MAN** is in critical condition as a result of an officer-involved shooting that happened just after the new year arrived. The **MAN’S** family claims Bakersfield police did not have to shoot 26-year-old **[VICTIM’S NAME]** who police say had a rifle and refused to put it down after being told to do so. The shooting happened at an apartment complex on the 700 block of Terrace Way. Police say they were responding to calls of shots fired shortly after midnight. When they got there, police claim they saw **[VICTIM’S NAME]** holding a rifle and told **HIM** to drop it.

This becomes the following text after steps 1–3:

What the proud parents are saying about having their bundle of joy at the start of 20 - 13.
[a6]open 5pm 365 - deko 2 **VICTIM** is in critical condition as a result of an officer-involved shooting that happened just after the new year arrived. **VICTIM’S** family claims Bakersfield police did not have to shoot 26-year-old **VICTIM** who police say had a rifle and refused to put it down after being told to do so. The shooting happened at an apartment complex on the 700 block of Terrace Way. Police say they were responding to calls of shots fired shortly after midnight. When they got there, police claim they saw **VICTIM** holding a rifle and told **VICTIM** to drop it.

Finally, after steps 4–5, we define the story to be:

VICTIM is in critical condition as a result of an officer-involved shooting that happened just after the new year arrived. **VICTIM’S** family claims Bakersfield police did not have to shoot 26-year-old **VICTIM** who police say had a rifle and refused to put it down after being told to do so.

³⁸Steps 2 and 3 above are not quite exact. We replace references to the victim with the name “Pete” (same starting letter as “Patient”) and to the perpetrator with “Adam” (same starting letter as “Agent”). We do this because the semantic role labeling model has been trained with human names, which makes the prediction explained in Section A.3 more accurate.

The shooting happened at an apartment complex on the 700 block of Terrace Way. Police say they were responding to calls of shots fired shortly after midnight. When they got there, police claim they saw **VICTIM** holding a rifle and told **VICTIM** to drop it.

Once we have defined the story and replaced all the instances in which the victim and suspect are referred to in the text with the explicit “Victim” and “Perpetrator”, respectively, we can proceed to uncover who is the agent or the patient in each sentence reporting on the killing of the victim.

A.3 Semantic Role Labeling

In this section, we present how we identify who (agent) does what (verb) to whom (patient)—the objective of semantic role labeling. We select for the analysis sentences within each story (as defined in Section A.2.2) that mentions the victim or perpetrator and involves one of the relevant verbs or adverb modifiers that we are interested in: “kill,” “fire,” “shoot,” “dead,” “die” or “gunned.”

For this task, we use another BERT-type model proposed by Shi & Lin (2019). The output uses the PropBank annotation dataset (Bonial *et al.*, 2010).³⁹ The model takes each argument of each predicate in the sentence and annotates it with the semantic roles in relation to the predicate. Each verb is defined to possibly have several type of predicates, such as:

- ARG-0 is usually PROTO-AGENT (who executes the verb)
- ARG-1 is usually PROTO-PATIENT (who is affected by the action)
- ARG-2 is usually benefactive, instrument, attribute
- ARG-3 is usually start point, benefactive, instrument, attribute
- ARG-4 is usually end point (e.g., for “move”- or “push”-style verbs)

The following illustrates how the algorithm works in our case. Consider the following sentence:

According to authorities, [Perpetrator] shot [Victim] in the left side of [Victim’s] head, killing [Victim] instantly.⁴⁰

The output from the model is presented in Figure A.4. It shows that the model identifies 3 actions or verbs (“according,” “shot,” and “killing”) and identifies whether there is an agent or patient for that verb, plus other information related to that action.

Once we have this output, we can assign the category for each sentence in our sample. We describe this process next in Section A.4.

³⁹This model is composed of 4 models for different subtasks: predicate detection, predicate sense disambiguation, argument identification, and argument classification. Nonetheless, all tasks use a neural network with a structure similar to BERT’s, with differential training for the output—7 layers, with the final output being fed by the concatenation of the final hidden state in each direction from the bidirectional long short-term memory (BiLSTM) neural network, fed through a multilayer perceptron (MLP).

⁴⁰We remind the reader that, at this point, we have replaced all references to either the victim with “Pete” and to the suspect with “Adam.” This is the coreference resolution task described in Section A.2.

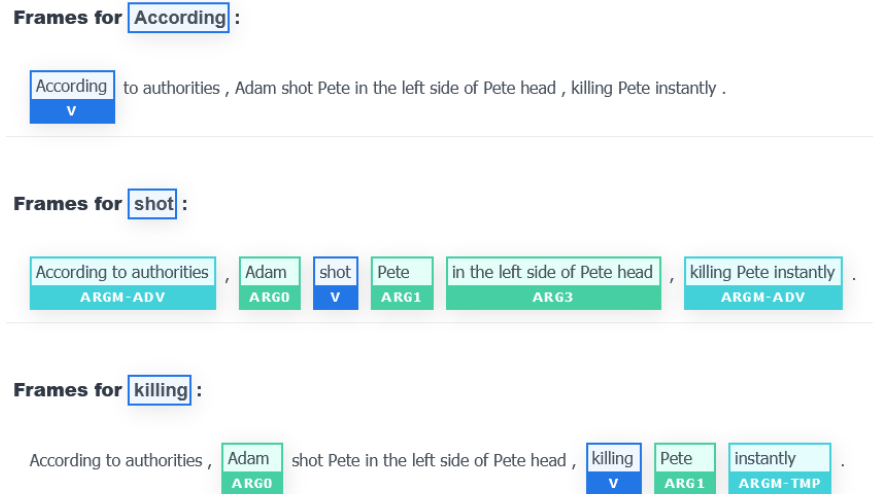


Figure A.4: Example of Semantic Role Labeling

A.4 Assignment of Obfuscation Dimension for Each Sentence and Story

Based on the linguistic framework laid out in Section 2, we define the following classification for each sentence.

Passive: We categorize the sentence as *Passive* if any of the following transitive verbs appear in the sentence with the patient being “Victim” (based on Section A.3) and the passive voice is used⁴¹: “kill,” “gun,” “murder,” “shoot,” “hit,” “fire,” “open (fire),” “strike.”

No agent: We subset those sentences classified as passive and classify them as *No agent* if the agent identified is either empty or different from the known one (in civilian killings it is “Perpetrator,” while in police killings it is “officer,” “deputy,” “sheriff,” “sergeant,” “detective,” “they,” “SWAT” or a slight modification of these words).

Intransitive: We classify the sentence as *Intransitive* if the following intransitive verbs appear in the sentence with the patient being “Victim” (based on Section A.3): “die,” “(is) dead,” “(declared/found/pronounced) dead.”⁴²

Nominalization: We classify the sentence as containing a nominalization if it includes a description of a shooting in the form of “[X]-involved shooting,” “X-related shooting,” “shooting (death) of” or “shooting (killing) of.”⁴³

⁴¹The passive voice is relatively easy to identify based on the use of a form of the auxiliary verb “to be” (such as “be,” “was,” “were,” and “are”) followed by the past participle of the main verb.

⁴²Although “declare,” “find” and “pronounce” are transitive verbs, the agent of killing is not explicitly acknowledged, and thus they effectively function like the other verbs in this category.

⁴³The “death” and “killing” in the parentheses may or may not appear in the text.

Last, the classification at the story level is equal to 1 if it includes sentences categorized in the respective obfuscation category above.

Appendix B Online Experiment: Additional Materials

B.1 Treatment Arms

Our experiment has 8 treatment arms: *Intransitive/Active/Passive/No agent*, interacted with whether the story mentions that the civilian killed was armed. We present each sentence structure in turn below.

Active

A police officer killed a 52-year-old man on Friday night.

According to the Police Department, an officer responded to a home near 21st Street and Avenue C for a report of domestic violence just before 9:30 p.m. As the officer arrived, he came into contact with a 52-year-old man [who was reportedly armed]. The police officer shot the man. The man was taken to the hospital, where he later died. No officer was hurt in the incident.

Passive

A 52-year-old man was killed by a police officer on Friday night.

According to the Police Department, an officer responded to a home near 21st Street and Avenue C for a report of domestic violence just before 9:30 p.m. As the officer arrived, he came into contact with a 52-year old man [who was reportedly armed]. The man was shot by police officers. The man was taken to the hospital, where he later died. No officer was hurt in the incident.

No agent

A 52-year-old man was killed in an officer-involved shooting on Friday night.

According to the Police Department, an officer responded to a home near 21st Street and Avenue C for a report of domestic violence just before 9:30 p.m. As the officer arrived, he came into contact with a 52-year old man [who was reportedly armed]. The man was shot in an officer-involved shooting. The man was taken to the hospital, where he later died. No officer was hurt in the incident.

Intransitive

A 52-year-old man died in an officer-involved shooting on Friday night. According to the Police Department, an officer responded to a home near 21st Street and Avenue C for a report of domestic violence just before 9:30 p.m. As the officer arrived, he came into contact with a 52-year old man [who was reportedly armed]. The man got wounded in an officer-involved shooting. The man was taken to the hospital where he later died. No officer was hurt in the incident.

B.2 Experiment Questions

Below are all the questions included in our online experiment. We set in italic red font the questions that we use to construct our primary outcomes, as specified in our AER preanalysis plan (AEARCTR-0009052).

Attention checks.

1. Was the person who died a man or a woman? (Man/Woman)
2. How old was the person who died? (20/52/89)

Participant retelling of the story.

Please summarize what happened, based on this news article.

Responsibility and consequences.

1. *Is the officer morally responsible for the man's death? (Yes/No)*
2. On a scale from 1 to 7, how much do you agree with these statements? (7 being the strongest and 1 being the weakest):
 - The police officer was justified in shooting the person
 - *The police officer should face some penalties in their department for their actions*
 - *The police officer should face some legal penalties*

Sentiment.

How do you think the police officer is depicted in the story? (Positively/Negatively/Neutrally)

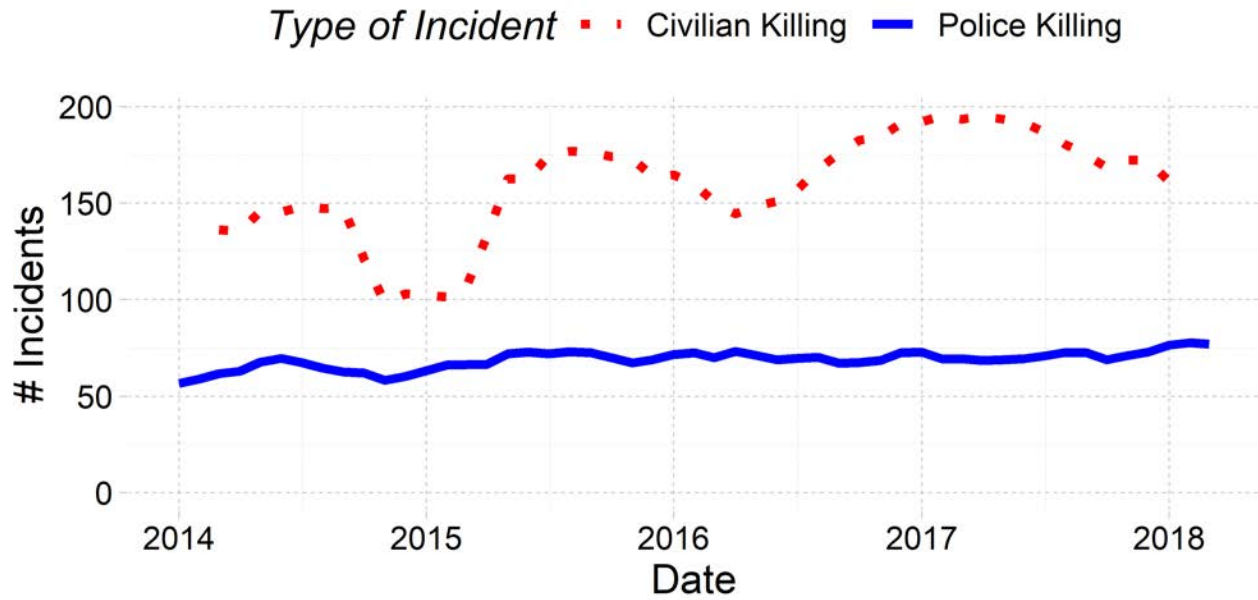
Donation.

One in 100 participants in this study will have the opportunity to donate \$100 to an organization. This is in addition to their payment for participating in the study. *Please choose how you want to split your donation among the two organizations below.* If you are randomly selected, we will make an anonymous donation to each organization as you have decided below. (Answers need to add up to 100 percent.)

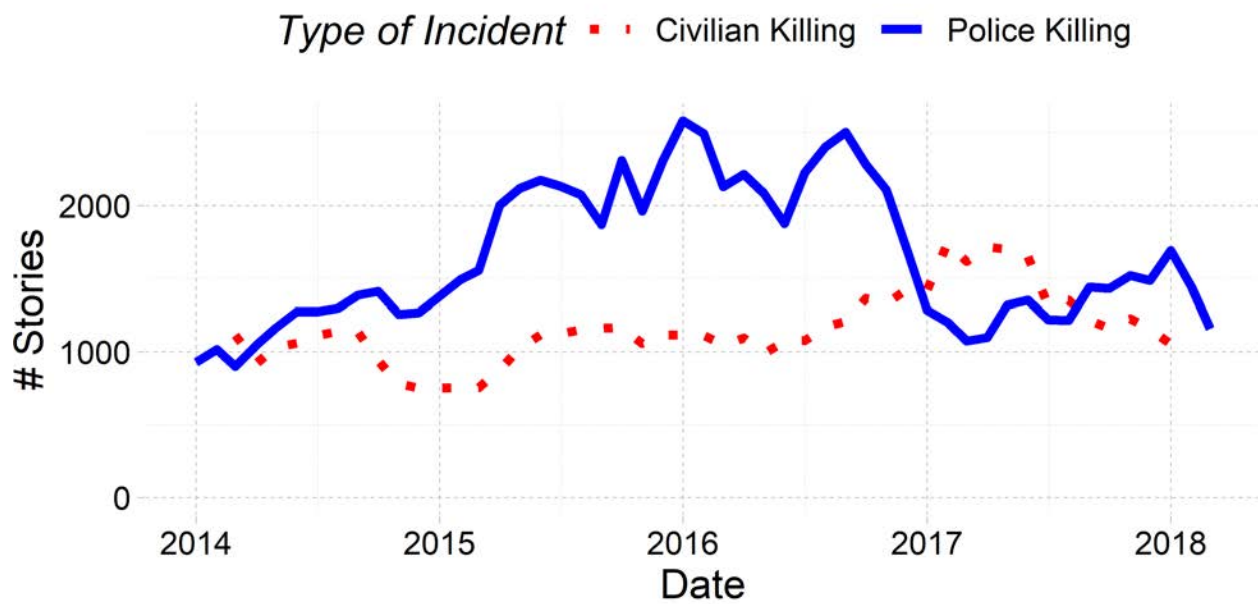
- An organization that aims to improve officer safety as well as health and wellness in police
- An organization that advocates to reform the police by increasing accountability, for example through officer training

Appendix C Additional Figures

Figure C.1: Number of Incidents and Stories in Sample



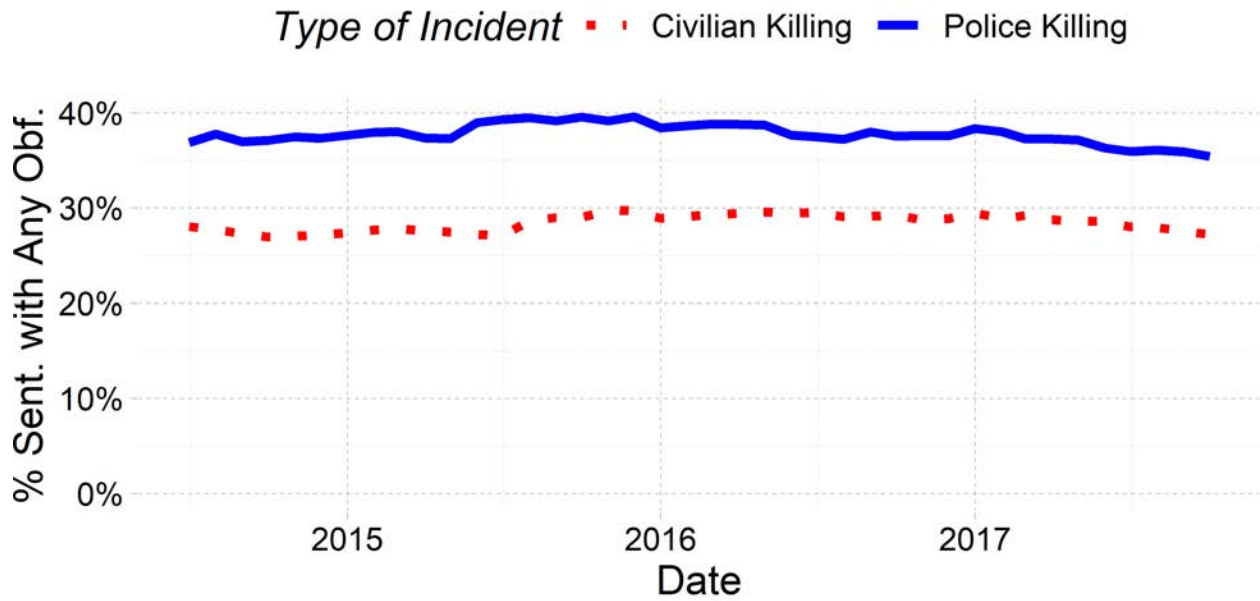
(a) Number of incidents



(b) Number of stories

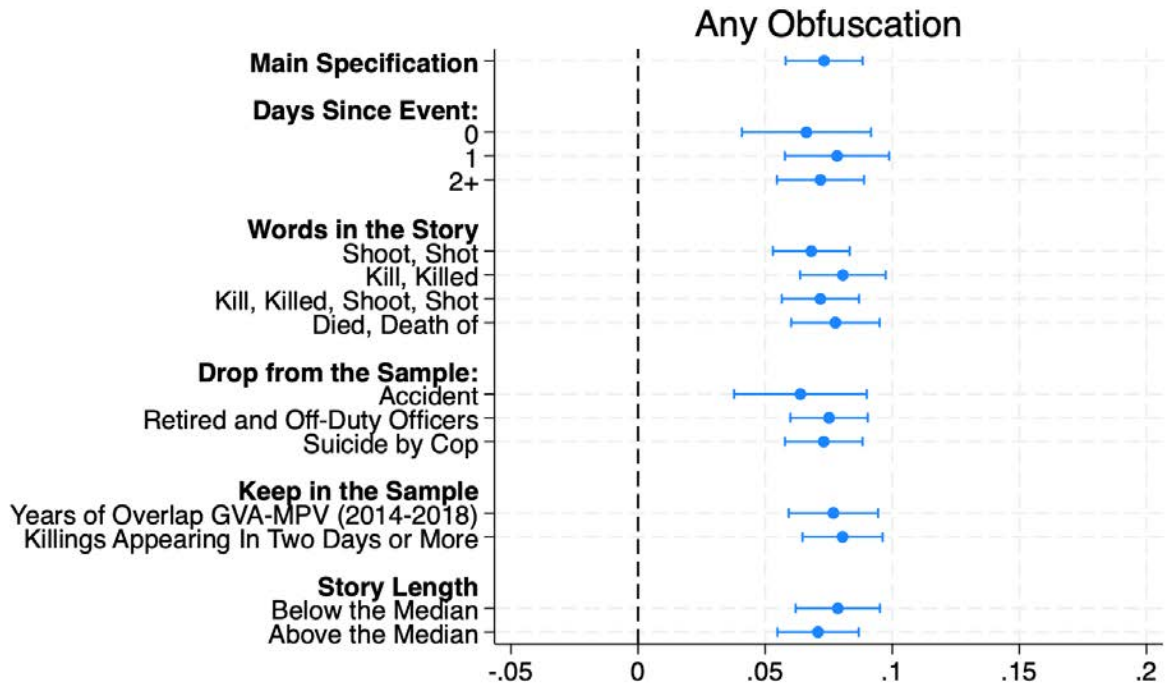
Notes: This figure plots the number of incidents (Panel a) and new stories (Panel b) on police and civilian killings in our main analysis sample. Our sample includes incidents and news stories where a suspect was identified for civilian killings. Data sources: GVA, MPV and News Exposure.

Figure C.2: Percentage of Stories with Any Obfuscation



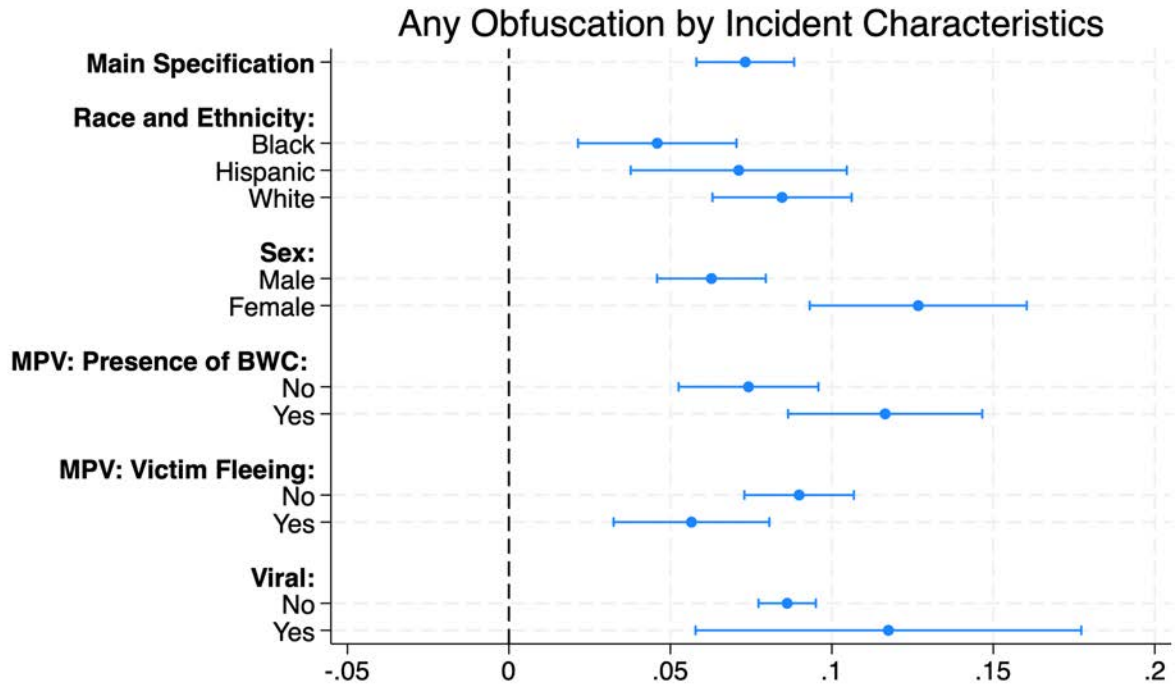
Notes: This figure presents the percent of stories over time that have some obfuscation for police and civilian killings. Our sample includes incidents and news stories where a suspect was identified for civilian killings. We define obfuscation as a sentence that has passive voice, has no agent, has intransitive verbs, or has a nominalization. See Section 2 for more details. Source: News Exposure.

Figure C.3: Robustness Tests



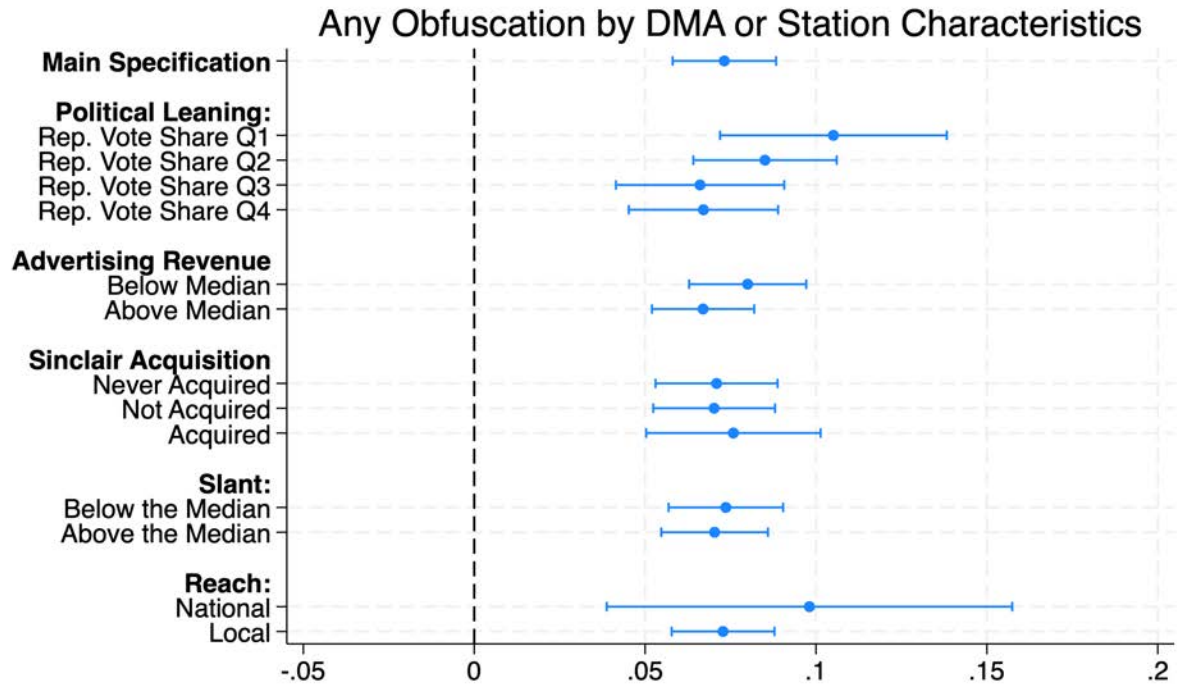
Notes: This figure presents differences in point estimates and 95 percent confidence intervals across different subgroups of news stories to test the robustness of our specification. We estimate Equation 1 separately for each subgroup. The checks include limiting our sample to the years for which we have both GVA and MPV data (2014–2018); limiting our sample to killings covered two or more days; running estimates separately for news stories aired on the day of the shooting, on the next day, or on subsequent days; limiting our sample to stories that include in the text the words “shoot,” “kill,” both “shoot” and “kill,” or “die”; dropping accidents from the GVA sample, where accidents are defined as entries with “accidental shooting” in the incident characteristics field of GVA; dropping from MPV cases flagged as involving retired or off-duty officers; dropping from MPV cases flagged by GVA as “suicide by cop” in the incident characteristics field; and splitting our sample by whether a news story is above or below the median length. We define “Any obfuscation” as sentences with passive voice, with no agent, with intransitive verbs, or with nominalizations. See Section 2 for more details. Source: News Exposure, MPV, GVA.

Figure C.4: Heterogeneity in Obfuscation, by Incident Characteristics



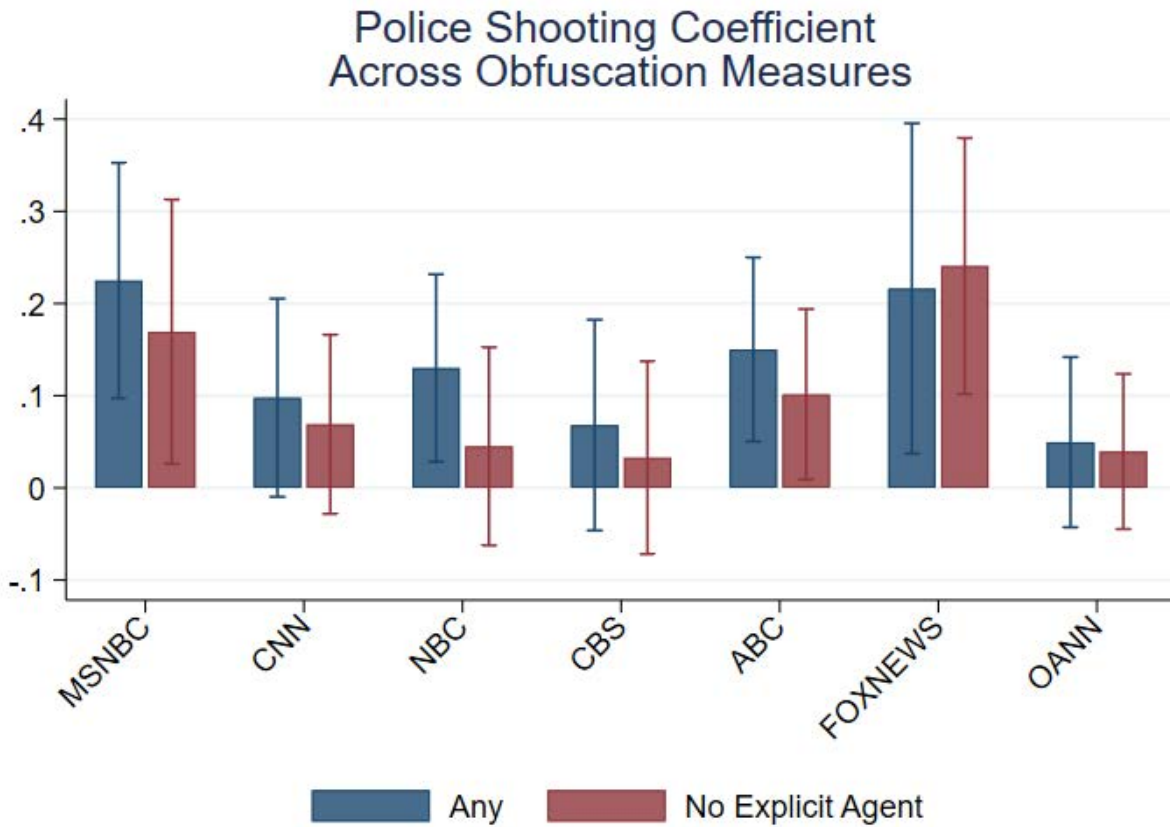
Notes: This figure presents differences in point estimates and 95 percent confidence intervals across different subgroups of news stories, by incident and victim characteristics. We estimate Equation 1 separately for each subgroup. The political leaning coefficients refer to media markets by quartile of Republican vote share in the DMA. Both the presence of BWC and whether the victim is fleeing come from MPV classifications. Virality is defined as an incident's being covered in more than 100 stories in our sample. Last, the sex and race of defendants are based on information on victims from MPV and GVA. We define "Any obfuscation" as a sentence with passive voice, with no agent, with intransitive verbs, or with nominalizations. See Section 2 for more details. Source: News Exposure, MPV, GVA.

Figure C.5: Heterogeneity in Obfuscation, by DMA or Station Characteristics



Notes: This figure presents differences in point estimates and 95 percent confidence intervals across different subgroups of news stories, by station or DMA characteristics. We estimate Equation 1 separately for each subgroup. The political leaning coefficients refer to media markets by quartile of Republican vote share in the DMA. We also consider advertising revenue (We split our sample by whether a story was broadcast on a program with advertising revenue above or below the median value within our sample); by TV station Sinclair ownership (never owned, eventually owned by Sinclair but not when the story was aired, owned by Sinclair when the story was aired); by slant (whether the story is aired on a station whose average slant, as measured by [Martin & McCrain \(2019\)](#), is above of below the median slant); and by reach (local vs. national TV stations). Data sources: GVA, MPV and News Exposure. We define “Any obfuscation” as a sentence with passive voice, with no agent, with intransitive verbs, or with nominalizations. See Section 2 for more details. Source: News Exposure, MPV, GVA.

Figure C.6: Obfuscation Estimates for Main National TV Stations



Notes: This figure presents the estimated coefficient for the level of obfuscation for the main national TV stations. Our sample includes incidents and news stories where a suspect was identified for civilian killings. Stations are ranked from least to most conservative, following the classification established by [Mitchell et al. \(2014\)](#). We define “Any obfuscation” as a sentence with passive voice, with no agent, with intransitive verbs, or with nominalizations. See Section 2 for more details. Source: News Exposure.

Appendix D Additional Tables

Table D.1: Descriptive Statistics Across Different Samples

	Police Killings			Civilian Killings		
	No Filter (1)	Subject Filter (2)	Story+Sent.Filter (3)	No Filter (4)	Subject Filter (5)	Story+Sent.Filter (6)
	Mean					
General Chars:						
Has Name of Victim	0.97	0.99	1.00	0.84	1.00	1.00
Age	36.80	36.77	36.79	32.15	33.31	33.90
Male	0.95	0.95	0.95	0.82	0.75	0.72
Black	0.17	0.18	0.18	0.37	0.31	0.28
Hispanic	0.18	0.17	0.16	0.11	0.10	0.09
White	0.55	0.58	0.58	0.29	0.42	0.47
Other/Unknown	0.10	0.08	0.07	0.24	0.16	0.15
Share Vote Rep. DMA	0.48	0.48	0.49	0.47	0.51	0.52
Caused by Gunshot	0.95	1.00	1.00	1.00	1.00	1.00
MPV Chars:						
Body Camera	0.07	0.07	0.07	.	.	.
Victim Not Fleeing	0.44	0.45	0.46	.	.	.
Unarmed/Unknown	0.18	0.15	0.15	.	.	.
GVA Chars:						
Has Name of Suspect	.	.	.	0.72	1.00	1.00
Suicide	.	.	.	0.02	0.00	0.00
Domestic Violence	.	.	.	0.09	0.18	0.21
Murder and Suicide	.	.	.	0.07	0.11	0.14
Gang-Related	.	.	.	0.04	0.04	0.02
Near School	.	.	.	0.00	0.00	0.01
Home Invasion	.	.	.	0.02	0.04	0.03
Number Victims in Incident	.	.	.	1.24	1.37	1.40
Observations	7663	6070	5759	49277	14011	7943

Notes: This table presents descriptive statistics across different filters. “No Filter” represents the overall data across the MPV and GVA datasets, except that the GVA data are filtered of police killings. “Subject Filter” is the sample once we filter by characteristics of the incident. “Story+Sent.Filter” is our main final sample, which includes a filter for sentences where both a victim and an alleged perpetrator were identified, for the story’s being predicted to be about crime with probability over 30%, and so on. *Data sources:* GVA, MPV and News Exposure.

Table D.2: Newspaper Data: Descriptive Statistics by Individual

	All (1)	Police Killings (2)	Civilian Killings (3)
<i>Panel A: Subject Level</i>			
Victim Chars.:			
Age	35.72	36.48	34.71
Male	0.82	0.95	0.66
Black	0.22	0.20	0.25
Hispanic	0.14	0.17	0.11
White	0.53	0.56	0.50
Other/Unknown	0.11	0.07	0.15
Incident Chars.:			
Body Camera	0.12	0.12	.
Victim Not Fleeing	0.67	0.67	.
Share Vote Rep. DMA	0.48	0.47	0.49
Observations	4915	2792	2123
<i>Panel B: Sentence Level</i>			
Obfuscation Dims.:			
Passive	0.16	0.18	0.13
Nominalization	0.05	0.04	0.07
No Agent	0.13	0.14	0.12
Intransitive	0.12	0.13	0.11
Any Obfuscation	0.33	0.34	0.30
No Explicit Agent	0.25	0.26	0.23
Observations	49024	32987	16037

Notes: This table presents the mean of different variables for newspaper stories. Our sample includes sentences where both a victim and an alleged perpetrator were identified. *Data sources:* GVA, MPV and NexisUni.

Table D.3: Obfuscation in Newspaper Coverage of Police Killings: Sentence-Level Analyses

Outcomes: Dimension of Obfuscation	Mean Civ. Shoot.	Police Killing				
		(1)	(2)	(3)	(4)	(5)
Panel A: Aggregate Dimensions						
Any Obfuscation	0.3026	0.036*** (0.009)	0.029*** (0.010)	0.030*** (0.010)	0.032*** (0.009)	0.030*** (0.010)
No Explicit Agent	0.2253	0.038*** (0.008)	0.031*** (0.009)	0.038*** (0.008)	0.039*** (0.008)	0.041*** (0.009)
Panel B: Individual Dimensions						
Intransitive	0.1123	0.018*** (0.006)	0.015** (0.006)	0.024*** (0.006)	0.028*** (0.006)	0.030*** (0.007)
No Agent	0.1207	0.021*** (0.006)	0.017** (0.007)	0.015** (0.007)	0.012* (0.007)	0.012 (0.007)
Nominalization	0.0675	-0.023*** (0.004)	-0.025*** (0.005)	-0.030*** (0.005)	-0.028*** (0.005)	-0.029*** (0.005)
Passive	0.1345	0.045*** (0.007)	0.042*** (0.007)	0.040*** (0.007)	0.036*** (0.007)	0.033*** (0.008)
Story Controls			X	X	X	X
DMA FE				X	X	X
Newspaper FE					X	
Month-Year FE						X
Observations		49,021	49,021	49,021	49,021	49,021

Notes: This table presents differential obfuscation for stories about police killings relative to stories about civilian killings from our estimation of Equation 1. Our analyses are at sentence level. We vary which controls are included across columns, including a time trend. Each row presents a separate regression coefficient on a dummy equal to 1 if the story is about a police killing rather than a civilian killing, for the different measures of obfuscation described in the first column. Our sample includes incidents and news stories where a suspect was identified for civilian killings. All sentences include some mention of either the victim or suspect. We define “Any obfuscation” as a sentence with passive voice, with no agent, with intransitive verbs, or with nominalizations. We define “No explicit agent” as sentences with no agent, with intransitive verbs, or with nominalizations. See Section 2 for more details. Standard errors clustered by subject (*: $p < 0.1$, **: $p < 0.05$, ***: $p < 0.01$). Source: News Exposure.

Table D.4: Obfuscation in Newspaper Coverage of Police Killings: First Sentence in News Story

Outcomes: Dimension of Obfuscation	Mean Civ. Shoot.	Police Killing				
		(1)	(2)	(3)	(4)	(5)
Panel A: Aggregate Dimensions						
Any Obfuscation	0.3144	0.015 (0.014)	0.011 (0.015)	0.017 (0.015)	0.021 (0.015)	0.013 (0.017)
No Explicit Agent	0.1599	0.064*** (0.011)	0.063*** (0.011)	0.069*** (0.012)	0.070*** (0.012)	0.068*** (0.013)
Panel B: Individual Dimensions						
Intransitive	0.0619	0.030*** (0.007)	0.031*** (0.007)	0.036*** (0.007)	0.038*** (0.007)	0.037*** (0.009)
No Agent	0.1036	0.034*** (0.009)	0.031*** (0.010)	0.033*** (0.010)	0.032*** (0.010)	0.029*** (0.011)
Nominalization	0.1477	-0.069*** (0.009)	-0.075*** (0.010)	-0.077*** (0.011)	-0.073*** (0.011)	-0.078*** (0.012)
Passive	0.1184	0.058*** (0.010)	0.059*** (0.011)	0.064*** (0.011)	0.063*** (0.011)	0.056*** (0.012)
Story Controls			X	X	X	X
DMA FE				X	X	X
Station FE					X	
Month-Year FE						X
Observations		12,409	12,409	12,409	12,409	12,409

Notes: This table presents differential obfuscation for stories about police killings relative to stories about civilian killings from our estimation of Equation 1. Our analyses examine only the first sentence of each news story. We vary which controls are included across columns, including a time trend. Each row presents a separate regression coefficient on a dummy equal to 1 if the story is about a police killing rather than a civilian killing, for the different measures of obfuscation described in the first column. Our sample includes incidents and news stories where a suspect was identified for civilian killings. All sentences include some mention of either the victim or suspect. We define “Any obfuscation” as a sentence with passive voice, with no agent, with an intransitive verb, or with nominalizations. We define “No explicit agent” as a sentence with no agent, with intransitive verbs, or with nominalizations. See Section 2 for more details. Standard errors clustered by subject (*: $p < 0.1$, **: $p < 0.05$, ***: $p < 0.01$). Source: News Exposure.

Table D.5: Summary Statistics for CES

	(1)	(2)	(3)
	All	No Obfuscation	Any Obfuscation
Male	0.47	0.46	0.48
White	0.73	0.73	0.73
Black	0.11	0.11	0.11
Hispanic	0.08	0.08	0.08
Other Race	0.08	0.08	0.08
Age	50.52	50.74	50.33
Ideology: Liberal	0.29	0.28	0.30
Ideology: Moderate	0.32	0.32	0.31
Ideology: Conservative	0.33	0.33	0.32
Ideology: Not Sure	0.07	0.07	0.06
Observations	29465	13894	15571

Notes: The table provides the mean values of the covariates for respondents participating in the Cooperative Election Survey across the years 2014, 2016, and 2018. Column 1 displays the aggregate results for the entire sample. Columns 2 and 3 break down these mean values by level of obfuscation.

Table D.6: Difference in Video Content of News Stories on Police Killings, by Level of Obfuscation

	Anchor (1)	Field reporter (2)	Police car (3)	Yellow tape (4)	Interview (5)	Victim photo (6)
Any obfuscation	0.05* (0.03)	0.00 (0.03)	0.04 (0.03)	0.04 (0.03)	0.01 (0.04)	0.02 (0.03)
Mean Dep. Var.	0.67	0.31	0.74	0.60	0.39	0.20
SD Dep. Var.	0.47	0.47	0.44	0.49	0.49	0.40
N	251	251	251	251	251	251

Notes: “Any obfuscation” is equal to 1 if the sentence structure has passive voice, has no agent, or has intransitive verbs. All outcomes are indicator variables equal to 1 if the following appears in the news story: an anchor, a field journalist, a police car, yellow tape, an interview (with a witness, a neighbor, law enforcement personnel, etc.), or a photo of the victim. We report the mean and standard deviation of the dependent variable for the “Active” narrative structure. The sample is all videos about police killings aired from April to June 2023 for which we found video footage. Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Source: News Exposure.

Table D.7: Online Experiment: Balance Test

	Age (1)	Male (2)	Prolific Experience (3)	Non-US Born (4)	Black (5)	White (6)	Hispanic (7)	Asian (8)	College Degree (9)	Passed Attention (10)
Passive	1.11 (0.82)	0.00 (0.03)	47.94 (37.73)	-0.01 (0.01)	-0.00 (0.02)	-0.01 (0.02)	-0.00 (0.02)	0.02 (0.02)	0.01 (0.03)	0.01 (0.01)
No agent + nominalization	0.65 (0.78)	0.01 (0.03)	64.40* (37.78)	-0.00 (0.01)	0.02 (0.02)	0.01 (0.02)	0.00 (0.02)	-0.00 (0.02)	-0.01 (0.03)	0.01** (0.01)
Intransitive + nominalization	0.92 (0.81)	-0.02 (0.03)	46.97 (36.82)	-0.01 (0.01)	-0.02 (0.02)	0.00 (0.02)	0.02 (0.02)	-0.00 (0.02)	-0.03 (0.03)	0.01 (0.01)
Weapon	-0.86 (0.57)	0.00 (0.02)	17.67 (27.12)	0.00 (0.01)	-0.01 (0.01)	-0.03 (0.02)	-0.01 (0.01)	0.03** (0.01)	-0.02 (0.02)	0.00 (0.00)
Mean Dep. Var.	36.51	0.44	881.85	0.06	0.09	0.77	0.08	0.11	0.61	0.98
SD Dep. Var.	13.88	0.50	621.81	0.24	0.29	0.42	0.27	0.31	0.49	0.13
N	2397	2402	2397	2402	2402	2402	2402	2402	2402	2402

Notes: The outcomes in each column capture respondent characteristics. We report the mean and standard deviation of the dependent variable for the “Active” narrative structure. Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table D.8: Online Experiment: Additional Results on Effect of Narrative Structure on Judgment of Event

	Officer Justified (1)	Negative Sentiment (2)
Panel A: No mention of victim weapon		
Passive	0.22* (0.14)	-0.07* (0.04)
No agent + nominalization	0.28** (0.14)	-0.17*** (0.04)
Intransitive + nominalization	0.36*** (0.14)	-0.17*** (0.04)
Mean Dep. Var.	2.60	0.35
SD Dep. Var.	1.71	0.48
N	1201	1201
Panel B: Victim has a weapon		
Passive	-0.14 (0.14)	-0.04 (0.03)
No agent + nominalization	0.21 (0.14)	-0.11*** (0.03)
Intransitive + nominalization	0.29** (0.14)	-0.08** (0.03)
Mean Dep. Var.	3.49	0.21
SD Dep. Var.	1.79	0.41
N	1201	1201

Notes: Panel A presents results when we do not specify and Panel B when we do specify that the victim was armed. The outcome in Column 1 is agreement on a scale from 1 to 7 with the statement that the police officer was justified in shooting the person. The outcome in Column 2 is a dummy equal to 1 if the person thought that the police officer was depicted negatively in the story. See Appendix B for the full questions. We report the mean and standard deviation of the dependent variable for the “Active” narrative structure. Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.