

1. Introduction

Public demands for police reform have accelerated in the years following highly publicized police-involved deaths. One common proposed reform is the creation of external civilian review boards (CRBs) with varied powers to review, recommend, and impose decisions on agency personnel and policy (Alpert et al., 2015; Walker & Bumphus, 1992a). Despite calls for this form of democratic policing for decades, relatively few such external review boards exist (Hope, 2021). Scholars have pointed to strong opposition from police executives as a primary reason for slow adoption of CRBs (Rosenthal, 2018).

The purpose of the current study is to better understand the influence of potentially competing sovereign demands on the policing institutional environment (Matusiak, 2016; Matusiak et al., 2017). How police executives—namely chiefs and sheriffs—develop policy preferences and learn in the face of external influences is an underdeveloped area of the literature (Nix et al., 2020). This inquiry is vital as political science literature suggests that political elites, including those in criminal justice, are responsive to public opinion (Aspin & Hall, 1993; Baumgartner et al., 2008; Canes-Wrone et al., 2014), yet criminology has explored little about the responsiveness of police executives to such opinions. To address this gap, we investigate whether police executives update their strongly held policy views when they are exposed to information on public opinion and peer practices. Our research questions are as follows: How do police executives' beliefs about Civilian Review Boards change in response to exposure to information on up-to-date, localized public opinion data and elite peer practices?

We demonstrate a methodological advance, using Structural Topic Modeling (STM) (Roberts et al., 2014) within an experimental framework. We analyze open-ended responses related to CRBs from the largest single experimental survey among police executives to date ($n=1,331$). Our novel approach enables a more authentic capture of police executives' views, offering valuable insights for policymakers and researchers alike. By analyzing what leaders say in their own words, our study establishes mechanisms through which public opinion and peer influence can shape the policy preferences of policing executives and characterizes how these elite political actors react to their various sovereigns (Matusiak, 2016).

Our study makes two primary academic contributions. First, it advances experimental methodology in criminal justice research by capturing shifts in respondents' attitudes through their own words, rather than through Likert scales or forced-choice questions. This approach alleviates concerns around expressive responding or researcher satisficing, which is increasingly salient given the partisan polarization of attitudes around police (Mourtgos & Adams, 2019; Novoa et al., 2023; Pickett, 2019). Second, it focuses on the concept of policy learning among police executives, offering insights into how leaders gather information and form judgments regarding controversial policy reforms in policing, a critical area of study given the call for effective policing reform and external oversight.

2. Literature Review

2.1 Civilian Review Boards

Civilian oversight of law enforcement agencies has existed in various forms since the inception of formal policing in America, with the first modern incarnation appearing in Washington, D.C. in 1948 (Fairley, 2020). Civilian oversight of the police is exercised through elected mayors, city councils, prosecutors' offices, the judiciary, and legislative bodies (Finn, 2001; Ofer, 2016; Stephens et al., 2018b; Walker, 2001). However, these forms of civilian oversight are not generally acknowledged by the public

(Stephens et al., 2018b) or are argued to be ineffective (Clarke, 2009). Indeed, there appears to be a growing sentiment that the above-listed oversight mechanisms, along with control mechanisms internal to law enforcement agencies (i.e., internal affairs units), do not work (Ofer, 2016). A 2014 poll conducted by USA Today and Pew Research Center found that Americans by a 2-to-1 margin believe that police departments nationwide do not do a good job in holding officers accountable for police misconduct (Page, 2014).

Calls for additional forms of civilian oversight often occur whenever controversial police actions garner national attention (Angelis et al., 2016; Fairley, 2020; Ofer, 2016; Stephens et al., 2018b). And while recent calls for additional civilian oversight have been voluminous (President's Task Force on 21st Century Policing, 2015), they are not new. Angelis and colleagues (2016) provide a historical perspective on the evolution of Civilian Review Boards (CRBs), tracing their roots back to police commissions established during the Progressive Era reforms, which were intended to diminish the influence of political machines over municipal policing. Despite these intentions, effective oversight was not achieved, primarily because the commissioners, often chosen through political connections, lacked policing expertise and typically deferred to the judgment of police executives.

Post-Progressive Era, cities like Washington, D.C., Philadelphia, and New York City pioneered further civilian oversight mechanisms (Fairley, 2020; Walker, 2001). These entailed the establishment of review boards tasked with examining the conclusions of internal police misconduct investigations. Nonetheless, due to inadequate resources and opposition from law enforcement and politicians, these boards were short-lived. The Civil Rights Movement catalyzed new civilian oversight efforts from the late 1960s to the 1980s, transitioning from mere review to the establishment of civilian investigative oversight agencies with the authority to conduct independent investigations, a notable advancement from their predecessors. During this period, agencies such as Berkeley's Public Review Commission and San Francisco's Office of Citizen Complaints were founded.

The 1991 Rodney King incident marked the beginning of a new wave of civilian oversight, characterized by the emergence of hybrid bodies that amalgamated review and investigative functions. Additionally, the institution of police auditor roles targeted systemic issues within police departments. The evolution of police oversight and the advent of the hybrid model have led to a complex landscape where CRB models are diverse and no two CRBs in the U.S. are exactly alike.

There have been many attempts to create a taxonomy of CRBs. Hope (2021) identifies a five-part taxonomy: Investigation, Review, Appeals, Auditor/Monitor, and Hybrid/Mixed. Fairley (2020) identifies a seven-part, non-exclusive taxonomy: Investigative, Review, Audit, Adjudicative, Appeals, Supervisory, and Advisory. Vitoroulis and colleagues (2021) identify a broad three-part taxonomy: Review-focused, Investigation-focused, and Audit/Monitor-focused. Other researchers provide several other taxonomies (Finn, 2001; Prenzler & Ronken, 2001a; Savage, 2013; Stephens et al., 2018b; Walker & Bumphus, 1992b).

We adopt the Vitoroulis and colleagues (2021) definition: "[I]ndependent, external, and ongoing review of a law enforcement agency and its operations by individuals outside of the law enforcement agency being overseen" (p. 2). This definition allows for the inclusion of hybrid CRBs combining functions of several models, which has become the norm. While most CRBs remain review-focused, many have combined investigative and auditor/monitor powers. This allows many CRBs to review internal police investigations, receive complaints and conduct independent investigations, hear appeals, recommend dispositions, recommend discipline, revise police policies, and even subpoena records or witnesses. While the authority to make recommendations is held by most CRBs and subpoena powers for records is not uncommon, subpoena power for witnesses is rare, and the ability to impose discipline (up to and

including termination) independent of an agency's chief executive is even rarer (Hope Sr., 2021; Vitoroulis et al., 2021), and more contentious than any other accountability power (Ali, 2023).

Despite the long history of calls for CRBs and the increasing contemporary prevalence of CRBs, research on the efficacy of their stated goals is rare (Angelis et al., 2016; Bartels & Silverman, 2005; Prenzler, 2000; Prenzler & Ronken, 2001a; Stephens et al., 2018b; Walker & Bumphus, 1992b; Worden et al., 2018), perhaps as a result of their rarity or diversity. Research that does exist suggests that CRBs rarely disagree with internal affairs units' findings, and in some cases, internal dispositions are more likely to recommend discipline for officers compared to CRB recommendations (Prenzler & Ronken, 2001a; Walker & Bumphus, 1992b). For example, examination of the Los Angeles Police Department's new civilian discipline panel found that the panel comprised entirely of civilians was more lenient on accused officers than the agency's more traditional panel of two officers and one civilian (Rector, 2021). These observations have led some researchers to suggest that CRBs may improve public trust not of those who make allegations against the police but instead of the broader public by signaling an agency's commitment to accountability (Worden et al., 2018).

Our work examines how empirical data can inform police executives' learning processes regarding policy reforms. This understanding is fundamental for those engaged in the discourse on policing reforms, as it provides a foundation for strategies aimed at enhancing the democratic governance of police forces, regardless of the varied outcomes of CRB effectiveness. Despite popular and repeated calls for more CRBs over multiple decades, they are far from ubiquitous in American policing. One proposed explanation for the low use of CRBs is resistance from policing executives (Fairley, 2020; Finn, 2001; Hope Sr., 2021; Walker & Bumphus, 1992b). We turn our attention to police executives and CRBs next.

2.2 Policing Executives and CRBs

Police resistance to CRBs has been noted since their inception. There is a well-documented perception among police of all ranks that the public does not understand the realities of police work and is therefore unqualified to judge the appropriateness of police actions (Mourtgos et al., 2020). Other arguments against CRBs by police include misconduct not being as widespread as the media has led the public to believe and that oversight would undermine police job performance (Fairley, 2020; Flamm, 2022). More specific to police executives, however, is that police executives may suspect that their ability to effectively manage the organization is imperiled by external review. This concern may sharpen under specific conditions, such as when CRBs hold independent authority to discipline officers, rather than merely to review their actions and propose recommendations (Finn, 2001).

One potential source of resistance to CRBs is pushback from police executives. Police executives are powerful policy actors with the power to slow, hasten, impose, and block major police reforms in the country's 18,000 policing agencies. While police unions may influence reform efforts, it is police executives who arguably play the most critical role in determining whether a CRB will be established or discouraged in a given community: their decisions set the priorities, culture, and behavior of their agency and the employees within (Ingram et al., 2022). Indeed, twentieth-century police reforms have been most successful when spearheaded by cohorts of police executives rather than by Progressive-era reformers (Uchida, 2020). Historically and notably, unions have been more skeptical of CRBs than police executives. Consequently, efforts to invite civilian review of departmental policy and outcomes have largely been enabled (ultimately) by chiefs and sheriffs (Juris & Feuille, 1974). Baseline preferences amongst police executives thus present significant obstacles to CRB adoption; however, these barriers may not be insurmountable and overcoming them likely necessary.

The above concern(s) may be outweighed by the influence of extra-institutional demands on police executives. As explained by Crank and Langworthy (1992), police department executives do not operate

in a vacuum. They constantly interact with other powerful actors – sovereigns – in an institutional environment that have a great deal of power over receipt of their legitimacy and the continued provisions of resources necessary for organizational survival. That is, police executives are not the sole authority in determining whether a CRB (or any other institutional structure) is created but are often influenced by these sovereigns (Crank, 2003). Police executives thus serve at the intersection of the police agency and its broader institutional environment, such that sovereigns' influence is exerted through executives to the policing agency (Matusiak et al., 2017).

Some of the institutional sovereigns that can exert significant influence over police organizations' operations and structures include large, influential police agencies and professional associations such as the International Association of Chiefs of Police (IACP), the Major Cities Chiefs Association (MCCA), and Police Executive Research Forum (PERF). These institutional sovereigns can determine widely accepted technical procedures and orthodoxies of the policing profession, with police agencies tending to emulate their peers (Matusiak et al., 2017).

Institutional sovereigns are not restricted to these organizations, however. Other sovereigns that shape organizational priorities include the media, police employee associations, elected officials, state and federal institutions, and community organizations (Krimmel & Lindenmuth, 2001; Li & Brown, 2019; Matusiak, 2016; Matusiak et al., 2017). Indeed, police chiefs perceive their institutional environment as multidimensional (Matusiak, 2016), and chiefs place more importance on maintaining relationships with some sovereigns over others (Matusiak et al., 2017). However, the literature has not yet developed to the degree that we can adequately theorize how police executives balance these competing demands from various institutional sovereigns. Each can exert overlapping or competing influence on policy adoption. The limited available research indicates that police executives are responsive to public opinion (Mughan et al., 2020) and their peer agencies (Matusiak et al., 2017). Yet, how much weight police executives assign to these potentially competing sources is not well understood (Nix et al., 2020), including how police executives may balance demands from competing sovereigns within their institutional environment to meet the minimum expectations of that environment.

Complicating matters is the dual and interactive nature of relationships with sovereigns. That is, while police executives must navigate their institutional environment, meeting the minimum requirements of their sovereigns to maintain legitimacy, police executives also assign assessments of legitimacy to these sovereigns, which impacts interactions and decision making (Bottoms & Tankebe, 2012; Matusiak, 2016; Matusiak et al., 2017). For example, Nix and colleagues (2020) found that hostile local media coverage reduces police executives' perceptions of audience legitimacy of the public. They theorize that perhaps local media serves as a symbolic representation of the public to police executives, possibly influencing executives' willingness to share oversight responsibilities with an external civilian board. If true, this may influence resistance to CRBs following a controversial incident where a police executive believes their agency is being mistreated.

The influence of public opinion on police executives is also undertheorized (Bottoms & Tankebe, 2012; Nix et al., 2020). The lack of understanding of public opinion on police executives is even more problematic when considering that some policing executives are elected officials rather than appointed bureaucrats. That is, the two primary police executive positions within American policing are police chiefs (typically an appointed position within a municipality) and sheriffs (typically an elected county official). Sheriffs in the U.S. are elected in all states except two, meaning that they may need to secure public support and respond to public preferences, as well as potentially obtain partisan support to be elected and retain office (Hayes, 2001). One may expect this electoral connection (Mayhew, 2004) to possibly influence how sheriffs respond to public opinion as compared to an appointed police chief who

may be more concerned about pleasing the mayor, city council, or city manager who appointed (and can terminate) them (Krimmel & Lindenmuth, 2001; Li & Brown, 2019; Mughan et al., 2020). That is, the elected sheriff is directly accountable to their citizenry in a way appointed police chiefs may not be. It is reasonable to think that sheriffs are attentive to this political environment in which they operate (Helms, 2008).

As an example in the case of elected sheriffs, Mughan and colleagues (2020) analyzed more than 3,500 sheriff's offices over a fourteen-year period. They found that sheriffs take advantage of asset forfeiture processes (a politically unpopular activity) less often than municipal police agencies out of fear of electoral reprisal. In contrast, some evidence suggests that elected sheriffs are likely to give less weight to public opinion. When examining a similarly sized sheriff dataset, and in relation instead to immigration policy, Thompson (2020) found that the average Democratic sheriff cooperates with federal immigration authorities at nearly the same rate as the average Republican sheriff. This suggests that elected sheriffs largely have consistent views and are less likely to explicitly respond to partisan electoral considerations. Our understanding of how police executives engage in policy learning, then, remains incomplete.

With such limited research surrounding how policing executives balance the influence of competing sovereigns, we focus on discerning the average effects of two such sovereign influences on police executives' policy preferences: public opinion and the diffusion of peer practices. Further, given the wide variation in powers given to CRBs across the U.S., we also assess if these sovereign influences affect the powers police executives are willing to grant to CRBs. We now turn our attention to further explaining the current study.

3. Methodology and Findings

3.1 Sampling and data collection

Our data come from an experiment embedded in a survey distributed to 13,287 U.S. police chiefs and sheriffs via email between February 1 and March 7, 2022.¹ Our sampling frame was drawn from a database containing the individual contact details of law enforcement leaders in the U.S. across all levels of government. We reduce the larger dataset to just municipal police chiefs and sheriffs in agencies larger than one officer and with a listed email address. In total, 1,331 respondents completed the survey, providing a 9.98% response rate, higher than most web-based, police-centered survey efforts (Nix et al., 2019)² and in line with surveys of other elites in political science and economics disciplines. The respondents are representative of the broader population of chiefs and sheriffs in the United States, and are also representative of chiefs and sheriffs who did not participate in our survey. Overall, we show that respondents who are situated in areas with a slightly larger population and greater educational attainment were marginally more likely to respond, although the predicted probability discrepancy was minor (see Appendix Tables 2, 3, and 4). The appendix also includes additional information on power analysis.

The nation's municipal police departments and sheriff's offices are led by a fairly homogeneous group. Among the surveyed sample, the average police executive is a white, male chief with 30 years of total law enforcement experience and has held his current position for just over seven years. This average

¹ The survey was fielded using Qualtrics for approximately five weeks, with three reminder emails sent after the original recruitment email.

² See survey size comparisons provided in Adams et al. (2022).

respondent is likely to have at least a bachelor's (33%) or master's degree (30%) and oversees an agency with fewer than 25 sworn officers (56%). Descriptive statistics for the sample are reported in Table 1.

Table 1: Descriptive Statistics

Characteristic	(N = 1,284)
	n / N (%); Mean (SD)
Has Current CRB	187 / 1,276 (15%)
Current Position	
Chief	1,040 / 1,283 (81%)
Sheriff	243 / 1,283 (19%)
Formal Education Level	
High school graduate	150 / 1,108 (14%)
Associate degree	225 / 1,108 (20%)
Bachelor's degree	361 / 1,108 (33%)
Master's degree	331 / 1,108 (30%)
JD, MD, PhD degree	41 / 1,108 (3.7%)
Age (years)	
< 34	8 / 1,113 (0.7%)
35 - 44	112 / 1,113 (10%)
45 - 54	519 / 1,113 (47%)
55 - 64	379 / 1,113 (34%)
65 - 74	91 / 1,113 (8.2%)
75 +	4 / 1,113 (0.4%)
Sex	
Male	1,074 / 1,122 (96%)
Female	48 / 1,122 (4.3%)
Race	
Asian/Pacific Islander	8 / 1,110 (0.7%)
Black	46 / 1,110 (4.1%)
Hispanic	36 / 1,110 (3.2%)
Other	37 / 1,110 (3.3%)
White	983 / 1,110 (89%)
Partisan Identification	

	Democrat	115 / 1,075 (11%)
	Independent	309 / 1,075 (29%)
	Republican	651 / 1,075 (61%)
Years in Current Position		7.20 (6.04)
Full Time Officers in Agency		
	Fewer than 25	637 / 1,133 (56%)
	Between 25 and 49	230 / 1,133 (20%)
	Between 50 and 99	153 / 1,133 (14%)
	Between 100 and 499	98 / 1,133 (8.6%)
	Between 500 and 999	8 / 1,133 (0.7%)
	More than 1000	7 / 1,133 (0.6%)

3.2 Description of the experimental survey design

In our experimental design, we provided participants with current evidence regarding state-level public attitudes or the adoption of CRBs in major policing agencies. We randomly assigned police executives to a control group or one of three informational treatment groups: 1) state-specific public support for CRBs without disciplinary power, 2) state-specific public support for CRBs with disciplinary power (e.g., the ability to terminate officers), and 3) MCCA information about peer adoption of CRBs.³

To obtain state-level public opinion data, we collaborated with a national polling firm to survey 16,840 individuals about their support for CRBs, both with and without independent disciplinary power—a key distinction in the range of CRB powers (De Angelis et al., 2016; Prenzler & Ronken, 2001b). We find CRBs benefit from clear majority support: 65% of respondents support a CRB with investigative power and 60% support a stronger CRB model with disciplinary power. To acquire information on peer adoption, we relied on data curated by the Major Cities Chiefs Association (Stephens et al., 2018a).

3.2.1 Structural topic modeling approach

Our first task is to analyze the words of police executives, as reported in their unstructured responses to a broadly-scoped substantive question asked at the end of the survey: *“In your own words: What are your thoughts about civilian review boards?”* Of the 1,131 respondents who completed some portion of the survey, a total of 1,029 provided open-text responses (91%). One modern method to analyze this type of data is structural topic modeling (Roberts et al., 2014). We use an R package developed for this type of analysis (Roberts et al., 2019). STM is an unsupervised machine learning approach to text data, but includes a “human in the loop” that allows for substantive human expertise to better understand the topic structures that appear in the underlying matrices. For example, scholars have previously used this method to understand the views of street-level officers towards proactive policing (Mourtgos &

³ Considering that our peer treatment only shares whether peers have adopted CRBs and not the success or outcomes of such adoption, the mechanism arguably resembles policy imitation rather than policy learning.

Adams, 2019). This method can provide the benefits of in-depth qualitative methods with the replicability of quantitative approaches (Chung et al., 2021; Roberts et al., 2014).

More to the point, we are interested in police executives' thoughts regarding CRBs *in their own words*. Identifying police executives' perceptions through their own words is critical to understanding baseline policy preferences over CRB adoption, and whether various information provisions—whether public opinion or peer practices—lead to policy learning. Analyzing open-ended responses from police executives increases our confidence that what is being measured is the fuller attitudes of the respondents beyond their ability to articulate a relatively narrow response to a Likert scale or forced-choice question (Chung et al., 2021; Mourtgos & Adams, 2019; Roberts et al., 2019). While responses to closed-ended questions have utility, they present three specific issues when attempting to understand police executives' opinion regarding CRBs. First, when faced with closed-ended questions, respondents are forced to choose from available options, even if the best answer is not included. Relatedly, the use of open-ended questions mitigates the problem of cueing respondents to think of particular explanations, thus biasing responses. Third, collecting open-ended responses after randomized informational treatments allows for the measurement of attitudes that are potentially shaped by treatments.

Topic modeling is an ideal method to use when trying to identify latent themes within text data from open-ended questions in a statistically sound and reproducible manner (Mourtgos & Adams, 2019). In this approach, topics are represented as distributions over a set of words, reflecting semantically interpretable themes. These distributions are determined by analyzing documents as vectors of word counts, where each word is assigned to a specific topic within a document. Consequently, each document can be represented as a vector of proportions, indicating the fraction of words belonging to each topic (Roberts et al., 2014). This implies that topics are composed of mixtures of words, with each word having a probability of belonging to a particular topic.

The generation of topics in topic modeling relies on probabilistic Bayesian inference, which allows for an inductive approach using the available data. Unlike clustering algorithms that assign each document to a single topic, topic modeling permits a document to encompass multiple topics, thereby forming a mixed-membership model that better captures the complexity and multidimensionality of real-world documents. (Roberts et al., 2014). This flexibility enhances the representation of documents, enabling them to capture diverse themes and topics within a single coherent framework.

STM is an extension of the more general topic modeling approach. In the context of STM, the process includes indexing the relationship between each document's metadata and control variables, as well as independent variables. This indexing allows researchers to incorporate covariates that may contribute to variance (Roberts et al., 2014). Unlike methods that assume topic consistency across all participants, STM provides researchers with the ability to: 1) control for specific factors of interest by maintaining them at their sample median, and 2) examine variations in the independent variable(s) under investigation. This capability grants a significant advantage over other topic modeling approaches in which topic constancy across all units of analysis is assumed.

Grimmer and colleagues (2022) define the process of STM. As they describe, for each document i , there is a set of P covariates X_i^P which are used to explain topic prevalence. There is also a separate set of P' covariates X_i^C that explain topical content. Using these sets of covariates allows one to specify the data-generating process:

$$\gamma_k \sim \text{Normal}_p(0, \sigma^2 I_p)$$

$$\begin{aligned}\pi_i &\sim \text{LogisticNormal}_{K-1}(\Gamma'(X_i^P)', \Sigma) \\ Z_{im} &\sim \text{Multinomial}_k(\pi_i) \\ \tilde{W}_{im} &\sim \text{Multinomial}_J(\mu_c Z_{im}) \\ \mu_{ckj} &= \frac{\exp(k_j^{(0)} + k_{kj}^{(t)} + k_{X^c j}^{(c)} + k_{X^c kj}^{(int)})}{\sum_j \exp(k_j^{(0)} + k_{kj}^{(t)} + k_{X^c j}^{(c)} + k_{X^c kj}^{(int)})}\end{aligned}$$

where i = an index for each document; π_i = a vector of topic weights; k = an index of topics; Γ is a $P \times (K - 1)$ matrix of coefficients for the topic prevalence of the model; Z_{im} = a multinomial distribution from which each word draws a corresponding topic indicator; m = the m^{th} word from a particular document; W_{im} = a draw from a multinomial distribution where μ_c describes the probability of a particular word's association with covariate c ; and $\{k_{..}^{(t)}, k_{..}^{(c)}, k_{..}^{(int)}\}$ describe a set of rate deviations for topics, covariates, and topic-covariate interactions from the overall vocabulary distribution within the corpus denoted by $k^{(0)}$. This process generates an $M_i \times 1$ matrix \tilde{W}_i . By summing over the columns, one obtains the corresponding row in the document term matrix, $W_i = \sum_{m=1}^{M_i} \tilde{W}_{i,m}$. The data-generating process is completed with priors, assuming both π_i and μ_c are drawn from a Dirichlet distribution.⁴

We apply this approach to the text data obtained from police executives in our survey. A unique contribution of our study is that we also assess whether different experimental conditions influence how police executives express their thoughts on CRBs. This approach provides an additional (and novel) lens through which to assess the effect of information provision on police executives' policy learning and attitudes regarding CRBs.

3.3 STM Analysis and Results

We measure impacts on how police executives think about CRBs by calculating the impact of each treatment on elements of the topic model described above. To accomplish this, we extract latent themes from the executives' open-ended responses, and then perform Bayesian regression on the resulting topic distributions to understand experimental effects on their responses.

To undertake STM, we follow the methods outlined in Mourtgos and Adams (2019), beginning with preprocessing of the raw text data, which is necessary to convert it into a format that computers can comprehend and learn from. One commonly used method for analyzing text data is the "bag of words" approach, which treats each word as a distinct feature in a document while preserving word co-

⁴ As stated by Grimmer and colleagues (2022, p. 158): "The key idea is that each k [topic] incorporates how the use of a word varies across a set of covariates, how the use of a feature varies with the prevalence of a topic, and how the two interact for a particular topic...STM, then, enables researchers to encode the ways topic prevalence and content carries across documents based on the document's characteristics, or metadata...The topic prevalence covariates inform our estimate of the prominence of topics in documents, while the topic content covariates alter how the topics are discussed."

occurrences. The corpus was pre-processed,⁵ including stemming⁶ and removal of stopwords.⁷ We applied these preprocessing steps to our corpus of 1,029 open-ended responses. Upon listwise deletion, 934 open-ended responses remained (90.8%; mean response length = 49 words; range = 1 – 886 words).⁸

Once the preprocessing steps are completed, the words in the corpus are organized into a document term matrix, where each document is a row and each word is a column. Next, we conducted an estimation of nine structural topic models for the corpus, ranging from two to ten topics. The posterior of structural topic models, like other mixed-membership topic models, is challenging to compute due to its intractability and non-convex nature (Roberts et al., 2014), making structural topic models sensitive to the initial seed used and resulting in instability across different model solutions.

To address this issue, we employed a spectral learning algorithm for initializing the topic models during estimation. Spectral learning algorithms offer consistent estimators by identifying an "anchor" word for each topic. By determining the anchor word, the model parameters can be estimated without iteration, thereby mitigating sensitivity to the starting values of the algorithm. This approach circumvents concerns related to multi-modality and guarantees consistent results irrespective of the seed used (Arora et al., 2013; Roberts et al., 2015).

After examining data-driven diagnostics (see Mourtgos and Adams (2019) for details regarding this process), we determined that the range of models to explore was between three and ten topics. The final decision on number of topics was made by closely reading the documents associated with each topic in the retained models and considering semantic coherence and word exclusivity. This exploration helps determine the number of topics that align with the content's meaning and minimize inter-topic similarities. If similar subjects are spread across multiple topics, it suggests inadequate exclusivity of words to individual topics. Through this manual examination, the five-topic model was found to have the best goodness of fit (Roberts et al., 2014). The resulting topic distributions are displayed in Figure 1.

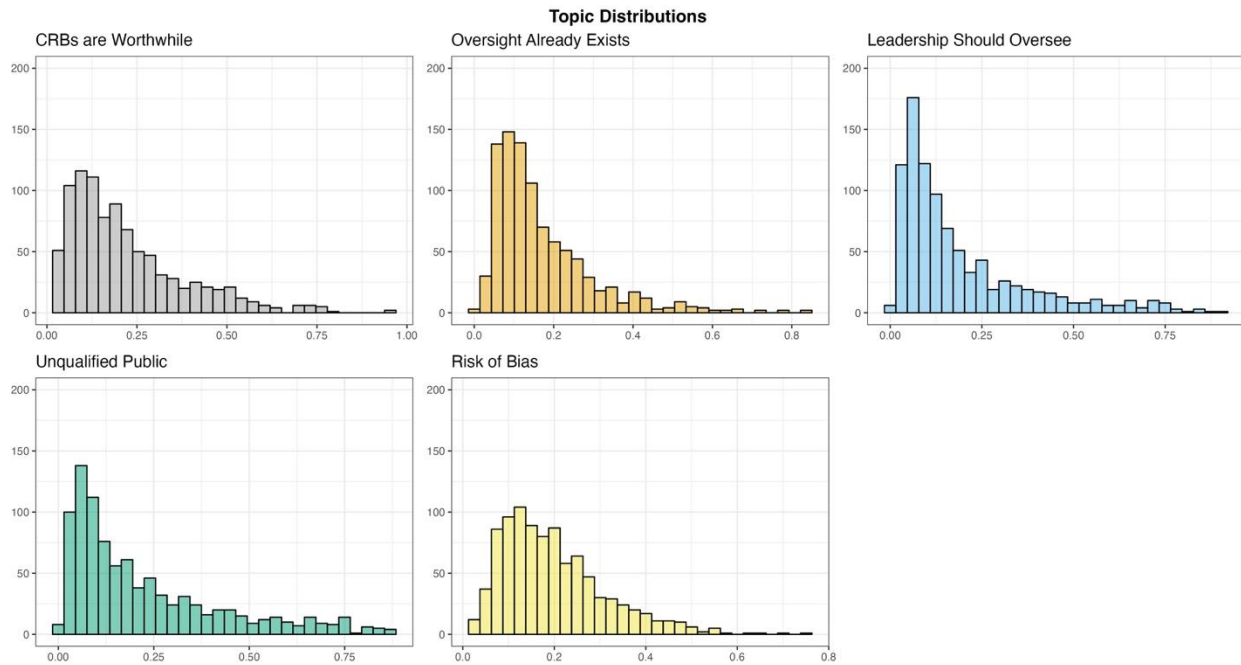
⁵ The bag of words method involves several preprocessing steps. First, the text needs to be tokenized, which entails identifying and separating individual words by treating whitespace and punctuation as word boundaries (Ignatow & Mihalcea, 2017). Second, stopwords, which are commonly occurring but non-substantive words such as pronouns, prepositions, and determiners, are removed since they provide little analytical value (Albert, 2020; Ignatow & Mihalcea, 2017). Next, various transformations are applied. All letters are converted to lowercase, punctuation is eliminated, non-alphanumeric characters are removed, and numerical values are discarded. Additionally, any HTML present in the text is stripped away. Sparse terms, typically words appearing in only one document, are also removed as they contribute minimal information and increase computational complexity (Roberts et al., 2018). The sparsity level was set to .999, meaning that words in at least 0.001% of the responses were included.

⁶ Stemming is performed to extract the root form of words, allowing related words to be grouped together in the resulting document-term matrix. For example, "police" and "policing" would both stem to "polic" (Ignatow & Mihalcea, 2017; Mourtgos & Adams, 2019).

⁷ It is common to remove custom stopwords in sparse text data, as virtually every respondent uses at least some of these words due to the specific question and topic. In a sparse data frame, not removing these overly common words can generate skewed word distributions, overwhelming other meaning-generating words, thus obscuring the underlying latent themes. The stopwords removed here included: review, police, law, board, enforcement, civilian, crb, boards, officers, agency, department, officer, citizen, civilians, citizens, public, crbs, crb's, le, departments, CRB, and CRBs.

⁸ We removed responses that did not have the accompanying metadata: 1) currently having/not having a CRB, 2) chief or sheriff, 3) education, 4) gender, 5) race, 6) self-identified partisanship, 7) tenure in position, 8) size of agency. These are required for developing the regression weights in the topic models.

Figure 1, Topic Distributions



Next, we examined ten highly associated documents for each topic. Through several close readings of the documents, an underlying theme for each topic was identified.

The task is to identify the common theme across responses highly associated with each topic. Topic 1, the **CRBs are Worthwhile** topic, contains responses that evaluate CRBs as a positive tool in furthering accountability, transparency, and the police-public relationship. Topic 2, the **Oversight Already Exists** topic contains responses that convey that policing already has many forms of oversight in place such that CRBs are an unneeded addition. Topic 3, the **Leadership Should Oversee** topic, expresses the idea that police executives are the individuals with proper responsibility who should have oversight of their personnel. Topic 4, the **Unqualified Public** topic, expresses the idea that members of the public, including appointees to a CRB, are not sufficiently qualified to pass judgment on police actions. Finally, topic 5, the **Risk of Bias** topic, expresses concerns that CRBs will suffer from bias amongst their members, not allowing for fair judgments of officers and their actions.

Topic names and examples of corresponding and highly associated responses are provided in Tables 2 through 6. Note that each response may have more than one topic represented within it.

Table 2

'CRBs are Worthwhile' Topic Documents

1. Independent oversight has clear and well-established benefits to law enforcement and the communities we serve. While some review board members may have political agendas, public opinion, and/or activism driving their opinions far more than objective fact gathering efforts, the benefits of listening to those opinions, just like any from the community, have great value to any law enforcement organization. The dialog is key to improvement and must remain open; this is most successful when neither side digs in and/or makes demands of the other,

but seeks first to understand, and shares a goal of providing better service. The best review boards and oversight bodies work with law enforcement, seeking improvements that are meaningful and realistic. As we seek to constantly improve policing in America, we must listen and be responsive to the needs of our customers, but not at the expense of the law or the rights of officers, witnesses or victims. These rights must be protected, respected and precede any agendas, trends, or public opinion. *-Male, White, Republican Sheriff at a Large Agency with a Bachelor's Degree*

2. Using civilian review boards increases transparency, helps foster trust between the police and the community, and facilitates police/community communication. *-Female, White, Republican Chief at a Small Agency with a Bachelor's Degree*
 3. Civilian review boards help increase the transparency of an agency. They CRB's can be useful providing there is a balance of civilian and law enforcement personnel. *-Male, White, Republican Chief at a Small Agency with a Bachelor's Degree*
-

Table 3

'Oversight Already Exists' Topic Documents

1. All incidents involving the use of deadly force, traffic crashes resulting in death or serious injury and allegations of misconduct where criminal charges may arise are investigated by outside law enforcement agencies at my request. *-Male, White, Republican Chief at a Medium-Sized Agency with an Associate's Degree*
 2. Our agency is under civil service and has a review board for the officers under the civil service system. The Tennessee Bureau of Investigation conducts all officer involved shooting incidents and other incidents that involve the use of excessive force. They would be no benefit to our agency since we already have oversight. *-Male, White, Republican Sheriff at a Medium-Sized Agency with an Associate's Degree*
 3. In my agency, situations that are serious in nature involving officers are investigated by an outside agency and then the case is reviewed by the grand jury for potential criminal charges. We do not need an independent board to be involved as the situations are reviewed by the independent grand jury. *-Male, White, Republican Chief at a Medium-Sized Agency with a High School Diploma*
-

Table 4

'Leadership Should Oversee' Topic Documents

1. ...I understand the cry for transparency and I agree but it only takes one rouge member on the board to get a leadership role and havoc begins. A simple fix to this issue is for the Leadership of agencies to be stand up and do the right things and do it loudly. So if an officers is wrong, his due process has been afforded and the investigation is completed, that Chief needs to be telling the community there was an issue and it has been dealt with. If the officer is vindicated then that Chief needs to be even louder telling the community they were wrong. As this type of vocal leadership progresses and it will be hard and many Chiefs will be fired

but it will change the perception the citizens have and begin the migration back to entrusting those who protect us with out doubt. I am willing to be that Chief to start this out and am doing that now as reform passes through Utah. Its my belief this new trend of police distrust was directly caused by Chiefs who refuse to do the right thing and that has led us to where we are today... *-Male, White, Independent Chief at a Small Agency with a Master's Degree*

2. Sheriffs in CA are constitutionally elected positions with absolute authority over all aspects of their Office. Civilian oversight of Sheriffs is equivalent to civilian oversight of District Attorneys, and other state and local elected officials. Since there are no civilian review boards for these elected officials, it is unnecessary for Sheriffs to be burdened with this type of oversight. *-Male, White, Independent Sheriff at a Large Agency with a Master's Degree*
 3. Officer conduct should be regulated by written policy from the Chief of Police. The Chief is hired by a citizen panel (City Council, etc) and serves as the person responsible for ensuring proper and accountable policing. If the City Council is unhappy with the conduct of the police department or its members, they should hold the Chief accountable. If citizens are not receiving the service they require from their police department, they should hold the City Council accountable. A police Chief should be hired to provide the level of policing and accountability that the community wants. Injecting a CRB into the mix causes confusion, distrust, and a lack of accountability (if you cannot hold a single person accountable, e.g. the Police Chief, then the responsibility and thus accountability gets diluted). *-Male, White Independent Chief at a Medium-Sized Agency with a Master's Degree*
-

Table 5

'Unqualified Public' Topic Documents

1. After twenty-seven years in law enforcement my knowledge, opinions, perceptions and expectations from the public are continuing changing. Law enforcement training and practices are also changing. I am a uniformed Sheriff and still work the road and respond to calls. We face danger today unlike anytime in my twenty-seven year career. Challenges, dangers, low staffing levels, and higher call volumes result in stresses unlike our profession has ever faced. We are also facing more violent citizens. When reviewing use of force issues I have to draw upon all my experience, training and knowledge to determine if the actions of the suspect, officers and witnesses warranted the force. I do not and cannot understand when someone who has no experience, knowledge or training in my field is now given the responsibility to determine if I or another officer acted reasonable and with due regard in their situation. I compare it to me determining if a doctor's actions were reasonable during a surgery where complications occurred that were affected the patient's outcome. I am not an expert or have any experience in that field. Watching surgeries on tv or taking a forty hour civilian academy on surgeries still does not provide me adequate experience or the expertise to be able to determine if their actions were reasonable. This needs to be done by leaders within their own organization and if they are not capable or trustworthy to carry this out then they need to be replaced. *-Male, White, Independent Sheriff at a Small Agency with an Associate's Degree*
2. Thank you for the opportunity to address this issue. When a doctor loses a patient, that doctor answers to a board of experienced doctors. These experts understand the science, the methods, and the best practices of their profession; therefore, these professional boards are

in the best position to address the mistakes or intent of that doctor. The same holds true for judges, teachers, and other specialized professions. Use of force issues should be addressed by use of force experts. CALEA and national standards are the best solution towards assisting agencies that are struggling with the necessary expertise. Civilian review boards are very poor products to understand the dynamics of use of force issues. They are inherently political. CRB's will have a tendency to either support or undercut what police present based on their community's ideology instead of a case by case professional evaluation. They rarely have the training and hands on experienced to understand case law, dynamic decision making under duress, or how to recognize failure in training. This is why other professions have experts and why law enforcement agencies need similar commissions, not political appointed CRB's. - *Male, White, Democratic Chief at a Medium-Sized Agency with a Master's Degree*

3. Civilian's have no idea how the law enforcement world works, nor do they have the knowledge to understand the situations we deal with on a day-to-day basis, and how we train to deal with them. -*Male, White, Republican Chief at a Small Agency with an Associate's Degree*

Table 6

'Risk of Bias' Topic Documents

-
1. I do not think that civilian review boards are necessary or, in most cases, appropriate, especially in cases involving discipline. Most review boards that recommend discipline do so by recommending lesser degrees of discipline than is desired by the agency. I also believe that civilian review boards become political pawns for whatever administration created them and that they are often staffed by people who have little to no understanding of case law, department policies or procedures, or officer safety issues. -*Male, White, Republican Chief at a Small Agency with a Master's Degree*
 2. I think it would be hard to find non-bias citizens that have predetermined thoughts about law enforcement. -*Male, White, Republican Chief at a Small Agency with an Associate's Degree*
 3. We do have a civilian panel as a part of the third step of disciplinary hearings that officers can request for grievance procedures. I just feel that civilians often do not understand the policies and procedures of a para-military organization unless they too have been in the military or a part of law enforcement. Often times it appears they think with their hearts and not with heads. Policies and procedures seem to go over their heads. -*Male, White, Democratic Chief at a Small Agency with an Associate's Degree*
-

3.4 Experimental Analysis

Can provision of different information influence police leaders' beliefs on controversial policy topics, specifically CRBs? We turn next to our analysis of this core question of policy learning. We use general linear modeling in a Bayesian framework to assess the effect of our experimental conditions on how police executives discuss CRBs. Recall that STM is a mixed-membership model, meaning that each document can be represented as a mixture of topics, rather than each document being restricted to only one topic. As such, topic proportions (theta values) can be extracted for each document, where a response's theta value refers to the proportion of each topic within it. We regress the experimental

conditions, along with controls, onto each latent theme to assess whether differential information provision influences how a police executive expresses their thoughts about CRBs.

We use Bayesian modeling because these methods assume a probability distribution for each parameter coefficient, rather than assuming one 'true' population parameter. The ability to acknowledge and leverage uncertainty through Bayesian methods allows researchers to better guard against false-negative results, inflated false-discovery rates, and inflated effect sizes (Barnes et al., 2020).

Due to the distributions of the identified topics, we estimate a gamma model for each topic regression. For each model, weakly informative priors were used and ten thousand iterations of four Markov-chain Monte Carlo (MCMC) chains were sampled to estimate the posterior distribution. All chains converged with Rubin-Gelman statistics of one. Adequate sample sizes were obtained for each parameter in the models. Finally, posterior predictive checks were examined. The resulting distributions in each model indicate a very good fit to the observed data providing confidence in the modeling strategy.

3.5 Experimental Results

STM enables us to examine the experimental effects on each topic in turn, focusing on changes in the prevalence of topics within the short responses of police executives. It is important to note that these changes in topic prevalence do not necessarily indicate a change in beliefs per se. Instead, they reflect a shift in the prominence or prioritization of certain topics in the executives' expressions when asked for their broad opinion. This might manifest as belief change or a reordering or a change in emphasis of the priorities among the topics measured.

Figure 2 plots the standardized effects of the experimental conditions on police executive attitudes. Table 7 provides model statistics for all parameters in each model. The standardized coefficients in our results can be interpreted as the estimated change in the prevalence of a topic (related to attitudes towards CRBs) for a one standard deviation change in the independent variable, with all other variables held constant. This standardization allows for comparison of the magnitude of effects across different variables in terms of standard deviations.

For example, a positive coefficient in the 'CRBs are Worthwhile' column in Table 7 indicates an increase in the prevalence of positive attitudes towards CRBs corresponding to a one standard deviation increase in the independent variable. Conversely, a negative coefficient suggests a decrease in the topic's prevalence. The magnitude and direction of these coefficients, combined with their 95% credible intervals, offer insights into the strength and relative importance of these topic shifts in the narratives of police executives.

Figure 2, Experimental Condition Effects on Police Executive Attitudes

Experimental Condition Effects

95% CIs




Table 7, All Parameter Effects for Topic Models

Variable	CRBs are Worthwhile				Oversight Already Exists				Leadership Should Oversee				Unqualified Public				Risk of Bias			
	Mean	Error	95% CI		Mean	Error	95% CI		Mean	Error	95% CI		Mean	Error	95% CI		Mean	Error	95% CI	
Tx Public Support	0.135*	0.062	0.013	0.258	0.195*	0.060	0.077	0.311	0.022	0.074	-0.123	0.169	-0.104	0.080	-0.260	0.054	-0.240*	0.048	-0.334	-0.148
Tx Public Support w/ Firing Power	0.003	0.062	-0.119	0.125	0.119*	0.059	0.003	0.236	-0.319*	0.075	-0.466	-0.171	-0.019	0.079	-0.172	0.138	0.176*	0.048	0.082	0.269
Tx Peer Adoption	0.254*	0.061	0.133	0.373	-0.169*	0.059	-0.283	-0.053	-0.006	0.073	-0.150	0.136	-0.068	0.079	-0.223	0.087	-0.124*	0.047	-0.217	-0.031
Current CRB (yes)	0.531*	0.062	0.412	0.656	-0.088	0.058	-0.201	0.028	0.069	0.076	-0.076	0.219	-0.445*	0.079	-0.600	-0.288	-0.174*	0.046	-0.264	-0.085
Sheriff	-0.129*	0.062	-0.251	-0.007	-0.186*	0.057	-0.299	-0.074	0.460*	0.077	0.312	0.612	-0.078	0.076	-0.225	0.070	-0.132*	0.046	-0.221	-0.040
Tenure (years)	-0.003	0.004	-0.010	0.004	0.001	0.004	-0.006	0.008	0.012*	0.005	0.003	0.021	-0.003	0.005	-0.013	0.006	-0.007*	0.003	-0.013	-0.002
Male	-0.132	0.113	-0.360	0.087	0.002	0.108	-0.214	0.209	0.467*	0.138	0.187	0.730	0.048	0.146	-0.246	0.326	-0.382*	0.086	-0.553	-0.218
Education																				
Associate	-0.084	0.078	-0.237	0.068	0.072	0.074	-0.074	0.216	-0.050	0.094	-0.236	0.134	0.184	0.098	-0.010	0.374	-0.104	0.059	-0.221	0.013
Bachelor	0.111	0.073	-0.033	0.255	-0.158*	0.069	-0.296	-0.023	0.199*	0.089	0.023	0.373	0.061	0.092	-0.120	0.241	-0.199*	0.055	-0.308	-0.092
Master	0.073	0.077	-0.077	0.223	0.069	0.072	-0.073	0.210	0.332*	0.092	0.150	0.513	-0.056	0.097	-0.248	0.132	-0.292*	0.059	-0.408	-0.178
PhD/JD	0.277*	0.130	0.023	0.533	0.274*	0.122	0.037	0.520	0.197	0.154	-0.100	0.504	-0.292	0.162	-0.606	0.032	-0.347*	0.098	-0.537	-0.156
Race																				
Asian/Pacific	0.912*	0.256	0.434	1.441	-0.065	0.245	-0.517	0.439	-0.904*	0.308	-1.461	-0.260	-0.718*	0.332	-1.320	-0.021	-0.153	0.191	-0.516	0.235
Black	0.079	0.117	-0.145	0.313	-0.187	0.113	-0.405	0.037	-0.254	0.141	-0.522	0.029	-0.228	0.150	-0.516	0.075	0.407*	0.088	0.238	0.581
Hispanic	0.210	0.122	-0.021	0.453	0.077	0.113	-0.143	0.305	-0.039	0.144	-0.312	0.250	-0.003	0.157	-0.301	0.310	0.031	0.091	-0.144	0.213
Other	-0.156	0.135	-0.412	0.118	-0.574*	0.127	-0.814	-0.316	0.115	0.161	-0.192	0.440	-0.082	0.171	-0.408	0.260	0.392*	0.101	0.199	0.592
Partisanship																				
Independent	-0.164*	0.083	-0.327	-0.003	0.045	0.077	-0.106	0.196	0.126	0.094	-0.060	0.307	0.227*	0.105	0.018	0.431	-0.280*	0.061	-0.400	-0.162
Republican	-0.032	0.077	-0.184	0.118	0.104	0.072	-0.038	0.244	-0.151	0.089	-0.328	0.022	0.083	0.098	-0.112	0.272	-0.011	0.057	-0.126	0.101
Agency Size																				
25-49	0.004	0.058	-0.107	0.118	0.210*	0.054	0.104	0.317	0.052	0.069	-0.082	0.189	-0.088	0.074	-0.231	0.056	-0.095*	0.044	-0.181	-0.009
50-99	0.035	0.070	-0.100	0.171	0.219*	0.064	0.094	0.346	0.138	0.082	-0.023	0.301	-0.164	0.087	-0.332	0.008	-0.193*	0.052	-0.294	-0.091
100-499	0.022	0.085	-0.139	0.191	0.035	0.080	-0.120	0.192	0.453*	0.101	0.259	0.653	-0.395*	0.110	-0.610	-0.178	-0.216*	0.063	-0.339	-0.091
500-999	0.111	0.250	-0.354	0.626	-0.294	0.235	-0.734	0.189	0.533	0.299	-0.019	1.160	-1.026*	0.329	-1.633	-0.337	0.086	0.189	-0.269	0.467
>1000	0.340	0.287	-0.194	0.937	-0.002	0.268	-0.491	0.546	-0.037	0.342	-0.657	0.684	-1.290*	0.353	-1.928	-0.544	-0.518*	0.206	-0.905	-0.098
Intercept	-1.577*	0.141	-1.852	-1.298	-1.918*	0.134	-2.177	-1.654	-2.442*	0.173	-2.779	-2.098	-1.454*	0.187	-1.815	-1.081	-0.824*	0.108	-1.034	-0.611

Standardized coefficients reported. * = 95% Credible Intervals that do not cross zero

The "CRBs are Worthwhile" topic model, which assessed attitudes that regard CRBs as advantageous for accountability and police-public relations, showed diverse reactions to the treatments. When presented with public support for CRBs without disciplinary power, a significant positive effect was noted (mean= 0.135, 95% CI [0.012, 0.256]). This indicates that public support for CRBs without disciplinary power positively affects police executives' views towards CRBs. Conversely, public support for CRBs with disciplinary power yielded a non-significant effect, suggesting that public preference for more robust (or aggressive) oversight does not induce more positive views of CRBs amongst police executives. Meanwhile, the treatment involving peer adoption demonstrated a significant positive effect (mean= 0.254, 95% CI [0.134, 0.373]), indicating that peer adoption information also positively sways attitudes towards CRBs.

In the "Oversight Already Exists" topic model, which reflects the viewpoint that existing policing oversight mechanisms are adequate, the treatments elicited varied responses. Public support for CRBs without disciplinary power led to a significant positive effect (mean= 0.195, 95% CI [0.079, 0.312]), reinforcing the belief in sufficient existing oversight. Public support for CRBs with disciplinary power similarly produced a significant positive effect (mean= 0.119, 95% CI [0.004, 0.235]). However, the treatment for peer diffusion showed a significant negative effect (mean= -0.168, 95% CI [-0.282, -0.053]), suggesting that information on peer adoption decreases the perception of CRBs as an unnecessary addition. Peer agencies may be considered as having weighed competing sovereign influences and trade-offs with respect to CRBs, such that they represent a more reliable source of policy learning.

The "Leadership Should Oversee" topic model, centered on the belief that police executives should oversee their personnel, responded significantly to public support for CRBs with disciplinary power. This treatment resulted in a significant negative effect (mean= -0.319, 95% CI [-0.466, -0.172]), indicating that public support for authoritative CRBs reduces the mention for executive-led oversight. Neither the public support for weak CRBs, nor peer adoption information, had a significant effect on these concerns. While we cannot say with certainty, we do not believe that this result is a function of executive de-emphasis of leadership when faced with support for authoritative CRBs. Rather, we believe that when faced with public support for CRBs with the authority to terminate employees, executives are more likely to focus on the topics of Oversight Already Exists and Risk of Bias than on Leadership, as suggested by Figure 2.

In the "Unqualified Public" topic model, none of the experimental treatments generated a significant change in police executive beliefs about CRBs, indicating that this is a core belief unmoved by external public opinion or peer effects. This finding aligns with other research suggesting that police officers generally believe the public is naïve about policing and its realities (Mourtgos et al., 2020; Mourtgos & Adams, 2019). It may be that the consistent belief in the public's unqualified status to competently judge police actions does not vary enough to affect police executive policy learning.

Finally, the "Risk of Bias" topic model, addressing concerns about potential bias against police within CRBs, presented contrasting effects based on the treatments. Public support for "weak" CRBs had a significant negative effect (mean= -0.240, 95% CI [-0.331, -0.148]) on bias concerns, perhaps suggesting less concern about bias in this lower-risk context where the public only has advisory powers. In contrast, public support for stronger CRBs with independent powers to impose discipline intensified bias concerns, as indicated by a significant positive effect (mean= 0.176, 95% CI [0.084, 0.268]). The treatment with peer adoption information produced a significant negative effect (mean= -0.123, 95% CI [-0.214, -0.032]), indicating that knowledge about major city agencies' experience with CRBs can also attenuate concerns about anti-police bias within CRBs.

3.6 Descriptive Associations

The non-experimental results of our study, which include individual respondent and agency-level correlates, offer insights into the attitudes of police executives towards Civilian Review Boards (CRBs). These non-experimental results offer a preliminary but insightful glimpse into the complex and varied factors that correlate with police executives' attitudes towards CRBs. As depicted in Table 7, the factors examined include personal experience, organizational role, tenure, gender, education, race, political

affiliation, and agency size, each potentially contributing uniquely to the overall perspective on civilian oversight in policing.

Compared to those without CRB experience, executives in agencies with an existing CRB were more likely to view CRBs as worthwhile (mean= 0.532, 95% CI [0.411, 0.656]) and less likely to perceive bias against police within CRBs (mean= -0.174, 95% CI [-0.265, -0.083]). This suggests that firsthand experience with CRBs may positively impact perceptions of their effectiveness and fairness, an important lesson for policy learning.

Sheriffs, compared to chiefs, showed a significantly different attitude in the "Leadership Should Oversee" and "Risk of Bias" models. They were more likely to believe in leadership-centric oversight (mean= 0.459, 95% CI [0.311, 0.612]) and less likely to perceive bias within CRBs (mean= -0.132, 95% CI [-0.222, -0.041]). This difference could reflect variations in organizational culture, professional identity, and operational environment between sheriff and police departments.

Longer tenure among police executives was significantly associated with a preference for executive-led oversight in the "Leadership Should Oversee" model (mean= 0.012, 95% CI [0.003, 0.022]) and a reduced perception of bias in CRBs (mean= -0.007, 95% CI [-0.013, -0.002]). This might indicate that more experienced executives favor traditional oversight mechanisms and are less concerned about bias in CRBs that may threaten their leadership. Male executives showed a significantly higher likelihood of supporting leadership-centric oversight (mean= 0.468, 95% CI [0.193, 0.729]) and a lower perception of bias against police within CRBs (mean= -0.382, 95% CI [-0.547, -0.219]) compared to their female counterparts. This could reflect gender differences in perceptions of organizational control and skepticism towards external oversight.

Education level was also a significant correlate in attitudes towards CRBs. Compared to those with only a high school degree, those with a PhD or JD degree exhibited more positive views towards CRBs (mean= 0.277, 95% CI [0.023, 0.533]) and less belief in public lack of qualification for oversight roles (mean= -0.292, 95% CI [-0.606, 0.032]). Bachelor's degree holders showed a slightly positive attitude towards CRBs (mean= 0.111, 95% CI [-0.033, 0.255]) but maintained a significant negative perception of public qualification (mean= -0.199, 95% CI [-0.308, -0.092]).

The racial background of respondents, with white executives as the reference group, also influenced views. Asian or Pacific Islander executives were significantly more positive towards CRBs (mean= 0.912, 95% CI [0.434, 1.441]). Notably, black executives showed no significant difference in their positivity towards CRBs (mean= 0.079, 95% CI [-0.145, 0.313]) but were more likely to perceive bias against police within CRBs (mean= 0.407, 95% CI [0.238, 0.581]). Hispanic executives did not significantly differ across any of the models.

Partisanship was not strongly associated with distinct attitudes, other than in the case of Independents. With Democrats as the reference group, Republican executives showed no significant difference in their overall positivity towards CRBs (mean= -0.032, 95% CI [-0.184, 0.118]) and were neutral about the public's qualifications for oversight (mean= -0.011, 95% CI [-0.126, 0.101]). Independent executives were less positive about CRBs (mean= -0.164, 95% CI [-0.327, -0.003]) and less concerned about anti-police bias (mean= -0.280, 95% CI [-0.400, -0.162]).

Agency size, with small agencies with fewer than 25 officers as the reference group, also correlated with attitudes. Executives in mid-sized agencies with 50-99 officers showed a significant belief that existing oversight is sufficient (mean= 0.219, 95% CI [0.094, 0.346]). Executives in mid-sized agencies with 100-499 officers and agencies with 500-999 officers showed less skepticism of public qualifications for oversight (mean= -0.395, 95% CI [-0.610, -0.178]) and (mean= -1.026, 95% CI [-1.633, -0.337]), respectively. Executives in larger agencies (>1000 officers) are much less likely to be skeptical of the public's

qualifications (mean= -1.290, 95% CI [-1.928, -0.544]) and are also less skeptical of potential bias within CRBs (mean= -0.518, 95% CI [-0.905, -0.098]).

4. Discussion

The findings of this study offer a new perspective on the role of policy learning and the attitudes of police executives towards CRBs. The use of STM (or other text-based statistical analysis) in analyzing police leadership narratives provides an advancement in criminal justice research methodology. This approach not only enhances the understanding of policy learning dynamics among law enforcement leaders but also provides insight into the factors influencing their stance on crucial governance issues and underlying policy rationales and assumptions. The use of open-text responses, complemented by a randomized experimental approach, provides a richer, more nuanced picture of police executives' attitudes, unencumbered by the constraints of closed-ended survey questions, offering valuable insights for policymakers and researchers alike.

A wealth of evidence in political science argues that political elites are responsive to public opinion generally (Mayhew, 2004), but also specifically in the criminal justice context through judicial and prosecutorial outcomes (Aspin & Hall, 1993; Baumgartner et al., 2008; Canes-Wrone et al., 2014). However criminology knows little about whether, and how, police executives respond to public opinion (Matusiak et al., 2017). We align with previous scholarship arguing that the police executive is little known, little understood, and yet likely has substantial impacts on policing outcomes (Adams et al., 2022; Hunt & Magenau, 1993; Matusiak, 2016).

Analyzing what leaders say in their own words can help establish the mechanisms through which public opinion can shape the internal policies of law enforcement agencies and better characterize how leaders view their various sovereigns. As a practical research matter, we argue, researchers should consider putting open text responses into an *experimental framework*. Through such a design, researchers can begin to understand the precise causal mechanisms that shape the thinking of these powerful executives and how their information environment determines their policy attitudes. Survey designs such as this one can further alleviate concerns around expressive responding or researcher satisficing in a forced choice framework—an especially salient problem given the increasing partisan polarization of attitudes around police (Mourtgos & Adams, 2019; Novoa et al., 2023; Pickett, 2019).

Regarding core policy implications, one of the most significant findings of this research is that public opinion and peer information can exert significant influence on the mindset of police executives. In particular, our study demonstrates that exposure to evidence of public sentiment and adoption rate of their peers can lead to sizable shifts in attitudes towards CRBs, including shifts supportive of and critical of these reforms. This suggests that police leadership's perspectives are not rigid but are instead open to influence and learning. For those interested in police reform, this finding highlights the potential for informed strategies to reshape police practices and civilian oversight, along with other accountability-related reforms. Understanding the determinants of police leadership's attitudes towards oversight mechanisms is essential for crafting interventions that could lead to a more responsive policing culture.

Taken together, the results suggest that executives are more inclined to express positive views about CRBs and less inclined to voice concerns about bias and existing oversight mechanisms when presented with information from peer agencies, especially regarding the diffusion of CRBs in (elite) peer institutions, such as members of the MCCA. Conversely, exposure to public opinion data showing support for CRBs with firing power increases executives' likelihood of expressing concerns about public bias, and public opinion data in general makes them more likely to assert that adequate oversight already exists. We believe these findings have substantial implications for understanding how police executives process information. Trust in the sources of information may play a major role in either driving or undermining credibility of policy reform proposals.

5. Conclusion

This study makes significant contributions to the understanding of policy learning among police executives in the context of Civilian Review Boards (CRBs). By employing Structural Topic Modeling (STM) within an experimental framework, we have provided a novel methodological approach in criminal justice research (Mourtgos & Adams, 2019; Roberts et al., 2014). This approach has allowed us to capture the nuances of shifting police executives' core policy beliefs in response to new information, offering insights that extend beyond traditional survey methods.

Our experimental design varied the information provided to the executives, incorporating state-level public opinion data from a representative sample of 16,840 U.S. residents and information on peer practices in major city police agencies. The findings reveal that police executives, while generally consistent in their views, are open to modifying their beliefs when presented with cohesive local public opinion and information about peer practices. This readiness underscores the value of our methodological innovation, as it captures a richer, more authentic understanding of executive beliefs. There is a vast potential for applying advanced analytical tools like STM in other areas of law enforcement and public administration, especially in understanding the intricate dynamics of police-community relations and the efficacy of civilian oversight. The application of STM underscores the multifaceted nature of police executives' attitudes towards CRBs, providing a nuanced understanding that is critical for developing policy interventions that are well-designed and implemented.

The study not only demonstrates the flexibility of police executives in adapting their core policy beliefs but also emphasizes the importance of informed decision-making in the dynamic field of police governance. These insights are crucial for policy adaptation and the enhancement of effective civilian oversight in policing. Our study thus marks a significant step forward in research methodology, offering a more authentic understanding of police executive opinions, and setting a new standard for future inquiries in this domain.

Our findings reveal a notable capacity for belief modification among police executives when presented with cohesive, current, and local public opinion, as well as peer practice information. This underscores the dynamic nature of policy preferences within police leadership and highlights the potential impact of informed, evidence-based communication strategies in shaping policing policies (Matusiak, 2021; Matusiak et al., 2017). The responsiveness of police executives to peer practices and public sentiment provides a promising avenue for driving meaningful reform in police governance.

This study also suggests new pathways for future research in criminal justice. Building on these findings, subsequent research could explore the effectiveness of different types of information and communication strategies in further influencing the attitudes and behaviors of law enforcement personnel, not limited to executives. Additionally, further research is warranted to explore the broader implications of these findings on community-police relations and public trust (Nix et al., 2020). For example, future studies could build on our findings by examining different communication strategies and their impact on policy adoption and implementation in policing.

The insights garnered from this research thus hold significant implications for policymakers, law enforcement agencies, and community advocates, as well as scholars. They highlight the critical role of informed dialogue in the process of transforming law enforcement practices and enhancing the legitimacy and effectiveness of police governance. By bridging the gap between empirical evidence and police leadership perspectives, the study offers a blueprint for better understanding and transforming law enforcement practices and enhancing the legitimacy and effectiveness of police governance.

References:

- Adams, I. T., McCrain, J., Schiff, D. S., Schiff, K. J., & Mourtgos, S. M. (2022). *Public Pressure or Peer Influence: What Shapes Police Executives' Views on Civilian Oversight?* SocArXiv. <https://doi.org/10.31235/osf.io/mdu96>
- Albert, Z. (2020). Partisan policymaking in the extended party network: The case of cap-and-trade regulations. *Political Research Quarterly*, 73(2), 476–491.
- Ali, M. U. (2023). The adoption of culturally contentious innovations: The case of citizen oversight of police. *Policy Studies Journal*, n/a(n/a). <https://doi.org/10.1111/psj.12499>
- Alpert, G. P., Cawthray, T., Rojek, J., & Ferdik, F. V. (2015). Citizen Oversight in the United States and Canada: Applying Outcome Measures and Evidence-Based Concepts. In *Civilian Oversight of Police*. Routledge.
- Angelis, J. D., Rosenthal, R., & Buchner, B. (2016). *Civilian oversight of law enforcement: Assessing the evidence*.
- Arora, S., Ge, R., Halpern, Y., Mimno, D., Moitra, A., Sontag, D., Wu, Y., & Zhu, M. (2013). *A practical algorithm for topic modeling with provable guarantees*. Proceedings of the 30th International Conference on Machine Learning, Atlanta, Georgia.
- Aspin, L. T., & Hall, W. K. (1993). Retention Election and Judicial Behavior. *Judicature*, 77(6), 306–315.
- Barnes, J. C., TenEyck, M. F., Pratt, T. C., & Cullen, F. T. (2020). How powerful is the evidence in criminology? On whether we should fear a coming crisis of confidence. *Justice Quarterly*, 37(3), 383–409.
- Bartels, E. C., & Silverman, E. B. (2005). An exploratory study of the New York City Civilian Complaint Review Board mediation program. *Policing: An International Journal of Police Strategies & Management*, 28(4), 619–630.
- Baumgartner, F. R., Boef, S. L. D., & Boydston, A. E. (2008). *The Decline of the Death Penalty and the Discovery of Innocence*. Cambridge University Press.
- Blair, G., Cooper, J., Coppock, A., Humphreys, M., & Sonnet, L. (2022). *estimatr: Fast estimators for design-based inference* [Computer software]. <https://CRAN.R-project.org/package=estimatr>
- Bottoms, A., & Tankebe, J. (2012). Beyond procedural justice: A dialogical approach to legitimacy in criminal justice. *The Journal of Criminal Law and Criminology*, 102(1), 119–170.
- Canes-Wrone, B., Clark, T. S., & Kelly, J. P. (2014). Judicial Selection and Death Penalty Decisions. *American Political Science Review*, 108(1), 23–39. <https://doi.org/10.1017/S0003055413000622>
- Chung, S. K. G., Rodriguez, M., Lanier, P., & Gibbs, D. (2021). *Text-mining open-ended survey responses using structural topic modeling: A practical demonstration to understand parents' coping methods during COVID-19 pandemic in Singapore*. OSF Preprints. <https://doi.org/10.31219/osf.io/enzst>
- Clarke, S. (2009). Arrested oversight: A comparative analysis and case study of how civilian oversight of the police should function and how it fails. *Columbia Journal of Law and Social Problems*, 43(1), 1–50.
- Crank, J. P. (2003). Institutional theory of police: A review of the state of the art. *Policing: An International Journal of Police Strategies & Management*, 26(2), 186–207. <https://doi.org/10.1108/13639510310475723>
- Crank, J. P., & Langworthy, R. (1992). Institutional perspective on policing. *The Journal of Criminal Law and Criminology*, 83(2), 338–363.
- De Angelis, J., Rosenthal, R., & Buchner, B. (2016). *Civilian Oversight of Law Enforcement: Assessing the Evidence* (250287; pp. 1–76). Office of Justice Programs. <https://www.ojp.gov/ncjrs/virtual-library/abstracts/civilian-oversight-law-enforcement-assessing-evidence>
- Fairley, S. R. (2020). Survey says? U.S. cities double down on civilian oversight of police despite challenges and controversy. *Cardozo Law Review De-Novo*, 2020, 1–54.
- Finn, P. (2001). *Citizen review of police: Approaches and implementation*.

- Flamm, M. W. (2022). "Law and Order" at Large: The New York Civilian Review Board Referendum of 1966 and the Crisis of Liberalism. 24.
- Grimmer, J., Roberts, M. E., & Stewart, B. M. (2022). *Text as data: A new framework for machine learning and the social sciences*. Princeton University Press.
- Hayes, C. (2001). The office of sheriff in the United States. *The Police Journal*, 74(1), 50–54. <https://doi.org/10.1177/0032258X0107400106>
- Helms, R. (2008). Locally elected sheriffs and money compensation: A quantitative analysis of organizational and environmental contingency explanations. *Criminal Justice Review*, 33(1), 5–28. <https://doi.org/10.1177/0734016808315588>
- Hope Sr., K. R. (2021). Civilian oversight for democratic policing and its challenges: Overcoming obstacles for improved police accountability. *Journal of Applied Security Research*, 16(4), 423–455.
- Hunt, R. G., & Magenau, J. M. (1993). *Power and the Police Chief: An Institutional and Organizational Analysis* (p. 172). Office of Justice Programs. <https://www.ojp.gov/ncjrs/virtual-library/abstracts/power-and-police-chief-institutional-and-organizational-analysis>
- Ignatow, G., & Mihalcea, R. (2017). *Text mining: A guidebook for the social sciences*. Sage.
- Ingram, J. R., Paoline, E. A., Rockwell, A. R., & Azimi, A. M. (2022). Police Chief Culture: A View From the Top. *Police Quarterly*, 10986111221114332. <https://doi.org/10.1177/10986111221114332>
- Juris, H. A., & Feuille, P. (1974). *The Impact of Police Unions: Summary Report*. U.S. Government Printing Office.
- Krimmel, J. T., & Lindenmuth, P. (2001). Police chief performance and leadership styles. *Police Quarterly*, 4(4), 469–483. <https://doi.org/10.1177/109861101129197950>
- Li, Y., & Brown, B. (2019). Police chief turnover in Texas: An exploratory analysis of peer-evaluation survey data pertinent to police performance and turnover. *Police Quarterly*, 22(4), 391–415. <https://doi.org/10.1177/1098611119845664>
- Matusiak, M. C. (2016). Dimensionality of local police chiefs' institutional sovereigns. *Policing and Society*, 26(7), 753–770.
- Matusiak, M. C. (2021). Do New Chiefs Mean Fresh Perspectives? The Institutionalization of the Police Executive Environment. *Criminal Justice and Behavior*, 48(10), 1390–1410. <https://doi.org/10.1177/00938548211006753>
- Matusiak, M. C., King, W. R., & Maguire, E. R. (2017). How perceptions of the institutional environment shape organizational priorities: Findings from a survey of police chiefs. *Journal of Crime and Justice*, 40(1), 5–19. <https://doi.org/10.1080/0735648X.2016.1155302>
- Mayhew, D. R. (2004). *Congress: The electoral connection* (2nd ed.). Yale University Press. <https://doi.org/10.2307/2129195>
- Mourtgos, S. M., & Adams, I. T. (2019). The rhetoric of de-policing: Evaluating open-ended survey responses from police officers with machine learning-based structural topic modeling. *Journal of Criminal Justice*, 64, 61–73. <https://doi.org/10.1016/j.jcrimjus.2019.101627>
- Mourtgos, S. M., Mayer, R. C., Wise, R. A., & O'rourke, H. (2020). The overlooked perspective of police trust in the public: Measurement and effects on police job behaviors. *Criminal Justice Policy Review*, 31(5), 639–672. <https://doi.org/10.1177/0887403419851850>
- Mughan, S., Li, D., & Nicholson-Crotty, S. (2020). When Law Enforcement Pays: Costs and Benefits for Elected Versus Appointed Administrators Engaged in Asset Forfeiture. *The American Review of Public Administration*, 50(3), 297–314. <https://doi.org/10.1177/0275074019891993>
- Nix, J., Pickett, J. T., Baek, H., & Alpert, G. P. (2019). Police research, officer surveys, and response rates. *Policing and Society*, 29(5), 530–550. <https://doi.org/10.1080/10439463.2017.1394300>
- Nix, J., Pickett, J. T., & Wolfe, S. E. (2020). Testing a theoretical model of perceived audience legitimacy: The neglected linkage in the dialogic model of police–community relations. *Journal of Research in Crime and Delinquency*, 57(2), 217–259. <https://doi.org/10.1177/0022427819873957>
- Novoa, G., Echelbarger, M., Gelman, A., & Gelman, S. A. (2023). Generically partisan: Polarization in political communication. *Proceedings of the National Academy of Sciences*, 120(47), e2309361120. <https://doi.org/10.1073/pnas.2309361120>

- Ofer, U. (2016). Getting It right: Building effective civilian review boards to oversee police. *Seton Hall Law Review*, 46(4), 1033–1062.
- Page, S. (2014, August 25). *Poll: Whites and blacks question police accountability*. <https://www.usatoday.com/story/news/nation/2014/08/25/usa-today-pew-poll-police-tactics-military-equipment/14561633/>
- Pickett, J. T. (2019). Public Opinion and Criminal Justice Policy: Theory and Research. *Annual Review of Criminology*, 2(1), 405–428. <https://doi.org/10.1146/annurev-criminol-011518-024826>
- Prenzler, T. (2000). Civilian oversight of police: Test of capture theory. *British Journal of Criminology*, 40(4), 659–674.
- Prenzler, T., & Ronken, C. (2001a). Models of police oversight: A critique. *Policing and Society*, 11(2), 151–180. <https://doi.org/10.1080/10439463.2001.9964860>
- Prenzler, T., & Ronken, C. (2001b). Models of police oversight: A critique. *Policing and Society*, 11(2), 151–180. <https://doi.org/10.1080/10439463.2001.9964860>
- President’s Task Force on 21st Century Policing. (2015). *Final Report of the President’s Task Force on 21st Century Policing*.
- Rector, K. (2021, May 11). Civilian disciplinary panels more lenient on LAPD officers. *Los Angeles Times*. <https://www.latimes.com/california/story/2021-05-11/new-civilian-discipline-panels-more-lenient-on-accused-lapd-officers-review-finds>
- Roberts, M. E., Stewart, B. M., & Tingley, D. (2014). Structural Topic Models for Open-Ended Survey Responses. *American Journal of Political Science*, 58(4), 1064–1082.
- Roberts, M. E., Stewart, B. M., & Tingley, D. (2015). *Navigating the local modes of big data: The case of topic models*.
- Roberts, M. E., Stewart, B. M., & Tingley, D. (2019). stm: An R Package for Structural Topic Models. *Journal of Statistical Software*, 91(1), Article 1. <https://doi.org/10.18637/jss.v091.i02>
- Rosenthal, R. A. (2018). Perspectives of directors of civilian oversight of law enforcement agencies. *Policing: An International Journal*, 41(4), 435–447. <https://doi.org/10.1108/PIJPSM-01-2018-0018>
- Savage, S. P. (2013). Seeking “Civilianness”: Police Complaints and the Civilian Control Model of Oversight. *British Journal of Criminology*, 53(5), 886–904. <https://doi.org/10.1093/bjc/azt033>
- Sjoberg, D. D., Whiting, K., Curry, M., Lavery, J. A., & Larmarange, J. (2021). Reproducible summary tables with the gtsummary package. *The R Journal*, 13(1), 570–580. <https://doi.org/10.32614/RJ-2021-053>
- Stephens, D. W., Scrivner, E., & Cambareri, J. F. (2018a). *Civilian Oversight of Police in Major Cities*. Office of Community Oriented Policing Services. <https://cops.usdoj.gov/ric/Publications/cops-w0861-pub.pdf>
- Stephens, D. W., Scrivner, E., & Cambareri, J. F. (2018b). *Civilian oversight of the police in major cities*.
- Thompson, D. M. (2020). How Partisan Is Local Law Enforcement? Evidence from Sheriff Cooperation with Immigration Authorities. *American Political Science Review*, 114(1), 222–236. <https://doi.org/10.1017/S0003055419000613>
- Uchida, C. D. (2020). The Development of the American Police. In R. G. Dunham, Alpert, Geoffrey P., & K. D. McLean (Eds.), *Critical Issues in Policing: Contemporary Readings* (Eighth Edition, p. 11). Waveland Press.
- Vitoroulis, M., McElhiney, C., & Perez, L. (2021). *Civilian oversight of law enforcement: Report on the state of the field and effective oversight practices*.
- Walker, S. (2001). Police accountability. *Belmont, CA: Wadsworth*, 199–214.
- Walker, S., & Bumphus, V. W. (1992a). The Effectiveness of Civilian Review: Observations on Recent Trends and New Issues regarding the Civilian Review of the Police. *American Journal of Police*, 11(4), 1–26.
- Walker, S., & Bumphus, V. W. (1992b). The effectiveness of civilian review: Observations on recent trends and new issues regarding the civilian review of the police. *American Journal of Police*, 11(4), 1–26.
- Worden, R. E., Bonner, H. S., & McLean, S. J. (2018). Procedural Justice and Citizen Review of Complaints Against the Police: Structure, Outcomes, and Complainants’ Subjective Experiences. *Police Quarterly*, 21(1), 77–108. <https://doi.org/10.1177/1098611117739812>

Appendix Materials

Treatment Balance

Balance is assessed with the `gtsummary` package (Sjoberg et al., 2021). We find that treatments were effectively balanced across the available covariates, lending further credence to our experimental results. Balance test results are reported in Appendix Table 1. Note that the race and agency size covariates have been collapsed in order to account for very small cell sizes in some levels of those variables.

Appendix Table 1: Covariate Balance Across Treatment Conditions

Characteristic	Control (No information), (N = 307)	Public Support CRB, (N = 324)	Public Support CRB (with firing), (N = 325)	Peer Diffusion, (N = 323)	p-value ¹
Position					0.7
Chief	246 (80%)	267 (82%)	267 (82%)	256 (79%)	
Sheriff	61 (20%)	57 (18%)	58 (18%)	67 (21%)	
Age					0.8
< 34	2 (0.7%)	1 (0.4%)	3 (1.1%)	2 (0.7%)	
35 - 44	21 (7.8%)	25 (9.0%)	27 (9.6%)	39 (14%)	
45 - 54	126 (47%)	128 (46%)	131 (47%)	134 (47%)	
55 - 64	96 (36%)	102 (37%)	93 (33%)	88 (31%)	
65 - 74	24 (8.9%)	22 (7.9%)	25 (8.9%)	20 (7.0%)	
75 +	1 (0.4%)	0 (0%)	2 (0.7%)	1 (0.4%)	
Unknown	37	46	44	39	
Gender					0.2
Male	262 (96%)	269 (96%)	275 (97%)	268 (94%)	
Female	11 (4.0%)	11 (3.9%)	8 (2.8%)	18 (6.3%)	
Unknown	34	44	42	37	
Race					0.3
Nonwhite	38 (14%)	27 (9.7%)	34 (12%)	28 (10.0%)	
White	233 (86%)	251 (90%)	246 (88%)	253 (90%)	
Unknown	36	46	45	42	
Partisan					0.5

Democrat	32 (12%)	28 (10%)	25 (9.4%)	30 (11%)	
Independent	65 (25%)	85 (32%)	74 (28%)	85 (31%)	
Republican	167 (63%)	156 (58%)	168 (63%)	160 (58%)	
Unknown	43	55	58	48	
Agency Size					0.6
< 25	147 (54%)	159 (56%)	174 (61%)	157 (54%)	
25 to 49	54 (20%)	54 (19%)	59 (21%)	63 (22%)	
50 to 99	43 (16%)	38 (13%)	28 (9.8%)	44 (15%)	
> 100	30 (11%)	31 (11%)	26 (9.1%)	26 (9.0%)	
Unknown	33	42	38	33	

n (%); ¹Pearson's Chi-squared test

Response/Non-response Balance

The present study also examines the comparative distribution of respondents and non-respondents, within the groups of chiefs and sheriffs. To this end, Appendix Table 2 and 3 present a balanced distribution of response across the county sheriff and police chief cohorts, respectively. Additionally, Appendix Table 4 features the outcomes of a linear model that forecasted agency-level participation in the survey, based on population traits and departmental officer count. As evidenced by the table, municipal police departments demonstrated a higher propensity to respond than their county sheriff counterparts. Departments situated in areas with a slightly larger population and greater educational attainment were marginally more likely to respond, although the predicted probability discrepancy was minor. Overall, these tables enhance the ability to generalize from the study's findings.

Appendix Table 2: Sheriff Response Comparison

Average	Non-Response Sample	Response Sample
Population	152463	156241
Percent White	0.82	0.84
Percent Black	0.09	0.07
Percent College	0.32	0.36
Income Below 50k	0.47	0.43
Income Above 100k	0.22	0.25
Employed Percent	0.58	0.61
2020 GOP Vote	0.65	0.62
2020 Dem Vote	0.33	0.36

Appendix Table 3: Chief Response Comparison

Average	Non-Response Sample	Response Sample
Percent College	0.34	0.37
Employed Percent	0.57	0.59
Median Income	59215	62298
Total Population	18996	21269
Percent White	0.74	0.75
Percent Black	0.10	0.08

Appendix Table 4: Linear Prediction of Survey Response

	Pooled	Municipal	County
<i>log Officers</i>	-0.001 (0.004)	-0.010+ (0.006)	0.018* (0.007)
<i>log Population</i>	0.015*** (0.004)	0.024*** (0.004)	-0.009 (0.006)
Percent White	0.035 (0.022)	0.031 (0.024)	0.146* (0.062)
Percent Black	-0.031 (0.025)	-0.029 (0.028)	0.012 (0.065)
Percent College Degree	0.073*** (0.020)	0.056** (0.022)	0.188** (0.067)
Municipal Dept.	0.045*** (0.008)		
Trump Vote 2020			-0.083 (0.054)
N	12067	9092	2975
R²	0.01	0.01	0.02
R² Adj.	0.009	0.009	0.015
F	19.737	18.307	8.418

Note: Outcome is a binary variable indicating the police department executive completed our survey.

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Power Considerations

In the current study, we build power simulations to determine the appropriate sample size required to examine the central hypothesis surrounding the differences between the control and treatment groups, using the package `estimatr` (Blair et al., 2022). With an 80% power threshold, the simulations indicated that standardized effect sizes of (small) 0.1 and (moderate) 0.2 necessitated approximately 4,200 respondents (30% response rate) and 1,600 respondents (12% response rate), respectively. The actual obtained sample size was 1,331 (10% response rate), signaling potentially underpowered analyses and an increased risk of false negatives in cases where very small “true” effects went undetected. Appendix Figure 1 displays the power analysis outcomes, accounting for two distinct effect sizes.

Appendix Figure 1: Power Simulation

