

# Institutional Factors Driving Citizen Perceptions of AI in Government: Evidence from a Survey Experiment on Policing

Kaylyn Jackson Schiff<sup>1</sup>, Daniel S. Schiff<sup>1</sup>, Ian T. Adams<sup>2</sup>, Josh McCrain<sup>3</sup>, and Scott M. Mourtgos<sup>3</sup>

<sup>1</sup>*Department of Political Science, Purdue University, West Lafayette, IN*

<sup>2</sup>*Department of Criminology and Criminal Justice, University of South Carolina, Columbia, SC*

<sup>3</sup>*Department of Political Science, University of Utah, Salt Lake City, UT*

June 28, 2024

## Abstract

Law enforcement agencies are increasingly adopting AI-powered tools. While prior work emphasizes the technological features driving public opinion, we investigate how public trust and support for AI in government vary with the *institutional* context. We administer a pre-registered survey experiment to 4,200 respondents about AI use cases in policing to measure responsiveness to three key institutional factors: *bureaucratic proximity* (i.e., local sheriff versus national FBI), *algorithmic targets* (i.e., public targets via predictive policing versus detecting officer misconduct through automated case review), and *agency capacity* (i.e., necessary resources and expertise). We find that the public clearly prefers local over national law enforcement use of AI, while reactions to different algorithmic targets are more limited and politicized. However, we find no responsiveness to agency capacity or lack thereof. The findings suggest the need for greater scholarly, practitioner, and public attention to organizational, not only technical, prerequisites for successful government implementation of AI.

**Keywords:** policing, artificial intelligence, survey experiment, public opinion, trust in government

**Word count:** 9,804

## Evidence for Practice

- The institutional and organizational structures into which algorithms are introduced matter, as individuals have baseline trust, expectations, and values associated with government institutions and bureaucrats, not only with AI systems themselves.
- We find a strong preference for *bureaucratic proximity*—members of the public indicate they are more supportive of, trusting of, and willing to pay for local, compared to national, law enforcement use of new AI tools.
- We find some evidence that responsiveness to different *algorithmic targets* falls along political and racial lines. However, we find no evidence of responsiveness to *agency capacity* or lack thereof to implement such tools effectively and responsibly.
- Efforts should be directed toward improving public understanding of the organizational, not only the technical, prerequisites necessary for government implementation of AI.

## 1 Introduction

Practitioners of public administration are increasingly expected to use artificial intelligence (AI) tools in their work. These tools, sometimes described as automated decision systems or decision support systems, can help to achieve administrative goals such as improved efficiency and service quality, and may even advance publicly held values such as fairness and transparency. At the same time, scholars have noted the mixed implications of AI adoption in public administration, including concerns like decreased bureaucratic discretion, potential algorithmic bias, and overreliance on automated systems (Wirtz, Langer and Fenner 2021, Zuiderwijk, Chen and Salem 2021). Indeed, the details of policy design and implementation choices, along with underlying technical features, shape the extent to which AI tools in a given policy sector further administrative goals or induce public value failures (Alon-Barkat and Busuioc 2022, Johansson, Thomsen and Åkesson 2022, Schiff, Schiff and Pierson 2021).

As the adoption of AI expands rapidly, few sectors are deemed more controversial than criminal justice, with heightened concern towards the use of AI in policing. Applications of AI such as predictive policing and facial recognition are at the center of national and

international policy debates. For example, one of the most controversial debates surrounding the European Union’s landmark AI Act is the extent to which the use of facial recognition for security and criminal justice purposes should be allowed or banned (Access Now 2021). While large coalitions of civil society organizations have argued for significant restrictions on predictive and biometric techniques in policing and immigration, prominent government actors have called for carve-outs, noting the benefits of these technologies for safety and security. At the same time, there is a legitimacy crisis in policing today surrounding issues like racial justice, community oversight, use of force, and accountability. Overall then, the stakes for carefully shaping the usage of AI in policing are incredibly high given the potential impact, including on vulnerable groups.

While predictive techniques represent the most widespread and heavily-studied form of AI usage in policing, there are numerous other applications of AI under development (Berk 2021, Ferguson 2017*a*). These include the use of audio and visual evidence in retrospective forensic analysis such as measuring gun shot trajectories and identifying actors involved in crimes. AI can also be used for internal, operational purposes, such as identifying behavioral patterns or concerns in police reporting that may be indicative of misconduct. For example, Truleo is an AI application that audits footage from body-worn cameras and promises to detect unreported use-of-force while measuring professionalism among police officers.

Public attitudes toward these new developments in policing practices matter for democratic accountability in policing, for responding to elevated calls for increased public oversight, and for advancing scholarly work on shared governance expressed through the New Public Service and related paradigms in public administration (Denhardt and Denhardt 2000). To date, scholars have expanded knowledge on how street-level bureaucrats and the public evaluate AI tools, including by experimentally assessing how public support and trust depend on an AI tool’s level of performance, transparency, or on the fairness of its automated recommendations (de Boer and Raaphorst 2021, Grimmelikhuijsen 2023). The plurality of this

work, however, has focused on the technical features of the tools or algorithms themselves. Meanwhile, scholars have urged that more work is needed to understand the contextual factors related to the institutional and organizational settings surrounding AI use (König et al. 2022, Meijer, Lorenz and Wessels 2021, Miller and Keiser 2021, Wenzelburger et al. 2022). Yet, there is little evidence to date beyond simple comparisons of public evaluations of AI in different sectors, such as in healthcare versus in child welfare applications.

This study uses a pre-registered<sup>1</sup> survey experiment to examine three such important and under-explored contextual factors that could shape public reactions to AI use in policing. We employ a 2x2x2 factorial design to provide vignettes with varying information about the actors, institutions, and targets involved in police use of algorithms. First, we examine *bureaucratic proximity*, or the institutional and geographic closeness of the implementers of the algorithms to the public in a federalist system (i.e., local sheriff versus national FBI). Second, to provide richer insight on within-sector variation in AI usage, we examine how the public responds to different *algorithmic targets*, or whether the AI tools are aimed externally on the public (i.e., for predictive policing) or internally on the bureaucrats themselves (i.e., for automated review of cases to detect officer misconduct). Third, because the skill and readiness of agencies to implement tools effectively and ethically matters, we explore how the public responds to concerns about *agency capacity* to fully use new AI-based tools.

We administer the experiment to 4,200 American adults in partnership with the polling organization Data for Progress. We assess public attitudes in response to the vignettes along a range of outcomes, including overall support for use of the AI tool, trust in the agency’s responsible use of the tool, whether individuals would be willing to pay more in taxes to fund the use of the tool, and willingness to share personal data to help improve the accuracy of the tool. In line with pre-registered hypotheses, we also examine heterogeneous effects related to partisanship and race.

---

<sup>1</sup>Pre-registered pre-analysis plan available at <https://osf.io/5mrwb/>.

Results indicate that the public is far more favorable toward AI use by local street-level bureaucrats closer to their communities (i.e., their local sheriff’s office versus the FBI), and is even more willing to undertake costly actions like sharing personal data and paying more in taxes to support AI adoption by local law enforcement. However, we find no overall difference in preferences regarding algorithmic targets—that is, whether AI is used internally (i.e., on officers to detect misconduct) or externally (i.e., on the public for predictive policing). We do find some subgroup differences—modest evidence of Democratic preferences for monitoring police vs. Republican preferences for predictive policing, higher support for both uses overall amongst Democrats, as well as some evidence of lower support for predictive policing amongst Black Americans. Finally, agency capacity appears to have little bearing on citizen attitudes, furthering recent evidence that the public is not always sensitive to factors that theoretically should matter (Grimmelikhuijsen 2023, König et al. 2022).

Overall, this study not only investigates a controversial and complex domain for AI adoption in public administration; it also expands the scholarly understanding of the institutional and organizational factors that are theoretically important for AI adoption and implementation. Extending recent related work on how the organizational context impacts predictive policing systems (Meijer, Lorenz and Wessels 2021), we specifically shed light on bureaucratic proximity, algorithmic targets, and agency capacity as drivers of public support for AI in government, and in law enforcement specifically. We find that some of these important factors do indeed matter to the public, but others fail to make a difference, contrary to expectations. The findings encourage further efforts to inventory the institutional and contextual factors that shape citizen preferences for AI in public administration.

## 2 Theory

### 2.1 Citizen Attitudes Towards AI in Government

Scholars of public administration have considered how algorithms and automated decision systems affect government generally, with specific attention devoted to impacts on street-level bureaucrats and to differences between automated and human decision-making. This has included a focus on issues of accountability for AI in public administration (Busuioc 2021), how artificial systems affect discretion at the program- and systems-levels (Bullock 2019, Young, Bullock and Lecy 2019), algorithmic bureaucracy (Vogl et al. 2020), and bias considerations related to automated versus human discretion (Compton et al. 2022). Much of this work has centered on street-level bureaucrats' opinions and attitudes toward AI and automated recommendations, with work additionally addressing how automated systems affect bureaucrats' behavior, such as by inducing pro-automation bias and selective adherence to preferred algorithmic outputs (Alon-Barkat and Busuioc 2022, Selten, Roberer and Grimmelikhuijsen 2023).

However, the opinions and attitudes of the public towards AI use in government are important as well. Most straightforwardly, citizens serve as key 'customers' of government agencies, as emphasized by the literature on e-government adoption (West 2004). This work and extensive research on technology adoption has shown that attitudes like trust influence citizen willingness to engage with government reforms, including technological reforms (Al-Hujran et al. 2015). For example, citizen trust in policing can affect critical outcomes like cooperation with police (Murphy, Mazerolle and Bennett 2014). Yet, members of the public are increasingly viewed as essential not only as clients but also as co-producers and colleagues. This perspective shift, reflecting a critique of the efficiency-oriented approach in New Public Management (Bryson, Crosby and Bloomberg 2014, Osborne 2006), is centered in frameworks such as New Public Service (Denhardt 2002) and the public values orientation (Bozeman 2002). Indeed, recent experimental research on the adoption of new technologies

like AI for government has argued that a sole focus on efficiency or cost reduction in lieu of other public goals may fail to produce key public values and even propagate negative ones (Schiff, Schiff and Pierson 2021).

Furthermore, the public is increasingly called upon to shape, or at least not actively oppose, technological transitions in government (Buhmann and Fieseler 2021, Ouchchy, Coin and Dubljević 2020), and this may be especially important in the area of policing, where public preferences and attention are heightened. As Jakobsen et al. (2019) observe in their review of citizen-state interaction research in public administration, these interactions are highly important across all aspects of the policy cycle, including for design (Neshkova and Guo 2012), implementation (Powers and Thompson 1994), and evaluation (Van Ryzin 2007). Moreover, public opinion, including as measured through polls and surveys, is commonly used as a heuristic for policymakers and elected officials. This can then shape, either directly or through oversight and accountability relationships, the interest in or support for AI integration into government. Finally, citizen attitudes are normatively important for core goals like democratic legitimacy and accountability. An important outstanding question, then, is what drives public support for algorithms and automated decision systems in government? This question is not only important to scholars, but also to bureaucrats and to law enforcement agencies in particular, as they determine where, why, and how to implement AI-based tools.

Recognizing the need to extend these questions towards AI, recent work on public attitudes toward AI in government has addressed issues such as transparency, fairness, explainability, effectiveness, and accessibility (Grimmelikhuijsen 2023, König et al. 2022, Schiff, Schiff and Pierson 2021, Schlicker et al. 2021). These are primarily features of automated decision systems themselves. Yet, we know that conditions external to the algorithms matter as well—the actors implementing new AI tools, trust in those actors, bureaucratic capacity, and other contextual factors impact how AI tools are perceived (Wenzelburger et al. 2022).

This is especially highlighted in the work of Meijer, Lorenz and Wessels (2021) on how context shapes predictive policing systems and is reflected in recent work addressing representative bureaucracy and automated decision making (Gaozhao, Wright II and Gainey 2022, Miller and Keiser 2021), and in work on how state welfare regimes shape receptivity to automated decision systems cross-nationally (Kaun, Larsson and Masso 2023). Notably, in examining the public’s perspectives on AI, some factors with significant theoretical importance—such as transparency or the actual performance of AI systems—have been shown to matter much less than expected (König et al. 2022, Wenzelburger et al. 2022). In lieu of clear or reliable evidence about the design, implementation, and implications of algorithms in given applied contexts, the public may resort to trust in implementers when they do not know whether to trust automated decision systems (Grimmelikhuijsen 2023).

Indeed, as Meijer, Lorenz and Wessels (2021) argue, impacts of AI systems on the operation of public administration cannot be taken deterministically as a result of technology itself. Instead, organizational rearrangements and impacts on bureaucrats due to algorithms are “not determined by the technological features but influenced by social norms and interpretations of the facilities of algorithmic systems” (Meijer, Lorenz and Wessels 2021, p.837). Along the same lines, we strongly suspect that features of the organizational and institutional context have primary importance not only for bureaucrats, but also for citizens and their evaluations of AI in government. The pre-existing institutions and structures into which algorithms are introduced matter, as individuals have baseline trust, expectations, and values associated with these institutions. While additional work is needed to more fully inventory the contextual factors impacting public perceptions of AI in government across a plurality of sectors, use cases, and regions, we take a step forward by addressing three in particular in this study: *bureaucratic proximity*, *algorithmic targets*, and *agency capacity*.



## 2.2 Key Contextual Factors and Hypotheses

### 2.2.1 Bureaucratic Proximity

*Bureaucratic proximity* refers to the institutional and geographic closeness of bureaucrats, and in this case bureaucrats implementing algorithms, to members of the public. Extending work using the lens of representative bureaucracy to examine the *actors* involved in automated decision making in government (Miller and Keiser 2021), bureaucratic proximity additionally incorporates the concept of job and community embeddedness (Mitchell et al. 2001) to consider how bureaucrats’ knowledge of, ties to, and investment in a particular community vary with their job scope. Thus, bureaucratic proximity is most applicable in federalist governmental structures in which some levels of government more directly represent and serve smaller geographic units. Street-level bureaucrats, and police officers in particular, exemplify a high degree of bureaucratic proximity. Yet there is variation in bureaucratic proximity in law enforcement, given that law enforcement operates at local to national scales. In this study, we consider local sheriffs and the FBI as representing opposite ends of this spectrum from the perspective of citizens of a single country (in this case, the United States).

Moving beyond passive, symbolic, or descriptive representation, we suggest that when bureaucrats’ jobs are local rather than national in scope, their public service motivation and/or job responsibilities may lead to higher-quality service or representation through better decision-making due to on-the-ground knowledge and stronger relationships through increased interactions with local stakeholders. These knowledge-based and relationship-based mechanisms may thus also directly foster increased trust between bureaucrats and members of the public. That is, members of the public are more likely to trust government officials who they feel know them better or are more likely to make decisions that are right for their community. We think that these factors may help explain prior research on “relative trust” or “the gap between trust in different levels of government in a federalist system” that typically

finds higher public support for local government than for national government (Leland et al. 2021). These trust gaps are also contingent on how duties and responsibilities are assigned (or attributed) across different levels of government.

As the public generally has a preference for local control in reducing crime (Schneider, Jacoby and Lewis 2011), we expect members of the public to place greater trust in local sheriffs than in the FBI. However, it is also possible that the public may perceive local control of AI tools as rife with opportunity for abuse – especially in the current political climate of large-scale distrust of local law enforcement. Public preferences for AI in government may also track preferences for drone regulation, which indicate higher support for federal and state regulation, as compared to local control, and especially for law enforcement uses (West et al. 2019). Nonetheless, we think on balance that the public will place more trust in local officials over federal officials. We expect that this trust will also extend to these actors’ use of technology, and we therefore hypothesize that members of the public will prefer local implementation of AI-based tools for law enforcement purposes:

**H1 Closer to Home Hypothesis:** *Citizens will be more supportive of AI use by local sheriffs than by the FBI.*

### 2.2.2 Algorithmic Targets

*Algorithmic targets* refers to the individuals or groups that are subject to the decisions or recommendations of automated decision systems. Just as we might expect public attitudes towards AI to vary across policy sectors, we might also expect public perceptions to vary *within* policy sectors, given that AI can be used in diverse ways and can target different actors (Ferguson 2017b). In this study, we specifically explore whether AI tools are aimed externally on the public (i.e., for predictive policing) or internally on the bureaucrats themselves (i.e., for automated review of cases to detect officer misconduct). Prior work in public administration has separately examined internal-facing AI applications such as work surveillance

technologies (Charbonneau and Doberstein 2020) and AI for human resource management (Vrontis et al. 2021), as well as externally-oriented applications like predictive policing (Meijer, Lorenz and Wessels 2021). Wenzelburger et al. (2022) find that citizens’ perceptions of algorithms stem from their sense of personal impact, from the value they attribute to the problem the algorithm is purported to address, and critically, from their trust in the users of the algorithms. Given recently reduced trust in police due to high-profile incidents of police violence, concerns about (human) bias in policing, and increased value attributed to solving the problem of officer misconduct, we hypothesize that the public, on average, to be more in favor of law enforcement officers as algorithmic targets (internal applications) as opposed to members of the public (external). To the extent that algorithmic targets are analogous to policy targets within the Social Construction of Target Populations framework (Schneider and Ingram 1993), the public may also view police officers as less “deserving” of protection and in need of greater (artificial) oversight:

**H2 Watching the Watchdogs Hypothesis:** *Citizens will be more supportive of AI use when applied internally to detect police misconduct as opposed to externally to drive predictive policing of the public.*

Partisanship and race are additionally likely to drive heterogeneous preferences in this context. Compared to their liberal-identifying counterparts, individuals with conservative views are known to hold more globally favorable views of police (Tyler et al. 2019) and to approve of police use of lawful force in the course of their duties (Mourtgos and Adams 2020). Partisanship also moderates youth support for policing specifically, with potential heterogeneous effects across racial groups. For example, Black youths’ perceptions of police are unaffected by partisanship, while Hispanic youth support is instead correlated with Republican identification (Fine et al. 2020). Similarly, partisanship helps explain public support for technology in law enforcement, such as for facial recognition in body-worn camera applications (Bromberg, Charbonneau and Smith 2020) and for governmental surveillance programs

generally (Reddick, Chatfield and Jaramillo 2015). Finally, political ideology and race are known to moderate support for use, management, and regulation of AI systems across various contexts (O’Shaughnessy et al. 2022). For these reasons, we expect partisanship and race to drive support for police use of AI technology as well. In particular, we expect non-white and Democratic respondents to hold less favorable views of law enforcement adoption of AI for policing the public compared to white and Republican respondents:

**H2.1 Partisanship Effects:** *Republican members of the public will be more likely to support AI use targeted at members of the public via predictive policing than will Democratic members of the public.*<sup>2</sup>

Along similar lines, prior research indicates that compared to non-white respondents, white respondents are up to six times more likely to support license plate readers (Merola and Lum 2014) and are more likely to support live facial recognition (Bradford et al. 2020). Non-white individuals may be concerned that these technologies will be used to disproportionately target individuals and communities of color. We therefore expect non-white members of the public to be less accepting of further technological adoption for policing the public, compared to white individuals:

**H2.2 Race Effects:** *Non-white members of the public will be less likely to support AI use targeted at members of the public via predictive policing than will than white members of the public.*

### 2.2.3 Agency Capacity

*Agency capacity* refers to the resources, personnel, and expertise that agencies have to perform certain functions. In the case of this study, it particularly references knowledge and

---

<sup>2</sup>We originally pre-registered a version of this hypothesis focused on ideology rather than partisanship. We recognize that partisanship and ideology are distinct concepts, but as they are highly correlated, especially in the current moment (Abramowitz 2022), we explore partisanship effects and present heterogeneous effects by partisanship in this paper.

ability to implement and use algorithms or automated decision systems effectively and responsibly. Meijer, Lorenz and Wessels (2021) suggests that expertise is a crucial component in the patterns of organizational rearrangement around algorithms. Relatedly, we expect that expertise, and capacity more broadly, may also condition citizen’s preferences surrounding government use of AI. Given the documented association between trust in government and government capacity generally (Denhardt 2002), we expect reduced trust in and support for AI in law enforcement when capacity limitations are made salient:

**H3 Capacity Drives Confidence Hypothesis:** *Citizens will be more supportive of AI use when capacity is not a concern.*

However, some prior work suggests that certain subsets of the public may favor lower levels of government professionalization, and therefore capacity (Kelleher and Wolak 2007, Richardson, Konisky and Milyo 2012). Whether these preferences extend to law enforcement remains an outstanding question with important implications. Moreover, given that some algorithmic features with significant theoretical importance—such as transparency or the actual performance of AI systems—have been shown to matter much less than expected (König et al. 2022, Wenzelburger et al. 2022), it is possible that capacity, an *institutional* feature with significant theoretical importance, will be less salient to members of the public as well.

## 2.3 The Policing Context

In “The Rise of Big Data Policing,” Ferguson (2017a) addresses the use of new technologies considered to be applications of AI in modern policing. As Berk (2021) argues in a broader review of the literature, citizen attitudes toward issues like fairness and transparency in AI systems are key to the legitimacy, adoption, and success of new technologies in policing and for the public’s trust and support for law enforcement generally. Scholars have thus examined questions surrounding public support for law enforcement as a matter of trust in government, investigating when the trust relationship varies, such as after instances of police

violence or based on an individual’s proximity to the law enforcement agency (Boudreau, MacKenzie and Simmons 2019, Silva et al. 2020).

Significant effort has been undertaken to understand the links between police activity and public trust in the broader institution of policing. For example, recent evidence demonstrates that public perceptions of police performance are negatively affected by protests against the police triggered by instances of police violence, such as those following the murder of George Floyd (Wright et al. 2023). The relationships are complex, however, and resist narratives that attempt to linearly predict perceptions of police based only on “good” or “bad” policing outcomes. Police agencies are consistently adopting new operational and communication strategies to improve their public standing (Ho and Cho 2017). For example, agencies adopt community policing operations in the face of public criticism, and these strategies are generally successful in shaping public perceptions of police more positively (McCandless 2018). In other words, the *outcomes* of policing, for good or ill, are not solely responsible for public perceptions of policing. Instead, a complex interplay between police adoption of operational, communication, and technology strategies both affect, and are affected by, public perception of police (Porumbescu, Neshkova and Huntoon 2019).

In pursuit of these questions, criminologists and scholars of technology in public administration have initiated a line of work at the intersection of police legitimacy, public perceptions, and technology. Following the procedural justice model in Tyler (2004), relevant research has investigated public support for police use of license-plate readers (Merola and Lum 2014), body-worn cameras (Mrozla 2021), drones (Sakiyama et al. 2017), and live facial recognition (Bradford et al. 2020). Consistent among studies is the finding that police use of technology both shapes, and is shaped by, public perceptions of police legitimacy. Public support for police expanding use of advanced technology constitutes a *de facto* grant of power to the police. Yet this expanded power is an extension of public trust, and how police enact that power has subsequent effects on generalized views towards police legitimacy, potentially

leading to a positive (or negative) feedback cycle (Neyroud and Disley 2008). For example, recent experimental survey work finds that amongst a sample of respondents with high levels of baseline trust in police, mention of license plate readers caused significantly lower expressions of trust in the police (Merola, Lum and Murphy 2019).

We thus explore our hypotheses within a setting in which context and pre-existing trust may be especially important and for which members of the public have heightened awareness and intensified preferences. We acknowledge that our results may not generalize to all areas of public administration and that much more work needs to be devoted to systematically studying variation across and within policy domains to understand the institutional factors that drive public preferences regarding AI.

### 3 Experimental Design and Analysis

#### 3.1 Experimental Design

We recruited 4,198 respondents between November 24, 2021 and January 4, 2022 through the national polling organization Data for Progress. Respondents were asked to complete an online survey hosted on Qualtrics which included the embedded experimental manipulation described in more detail below. The survey was administered through several waves, and the sample was pooled across waves and weighted to be representative of likely American voters by age, gender, education, race, and voting history. The appendix includes information about our power analysis. The design, hypotheses, models, analysis details, and power calculations were all included in the pre-analysis plan.

Through vignettes displayed to respondents, we experimentally manipulate the contextual and institutional factors surrounding AI use in policing: the AI application (for external predictive policing or for internal automated review of officer misconduct) representing the *algorithmic target*, the level of the law enforcement agency (a local sheriff’s office or the FBI)

representing *bureaucratic proximity*, and the stated resources of the agency (low capacity or no statement about capacity) representing *agency capacity*. This constitutes a 2x2x2 factorial design, detailed in Figure 1, with participants randomly assigned with equal probability to one of eight possible vignettes.

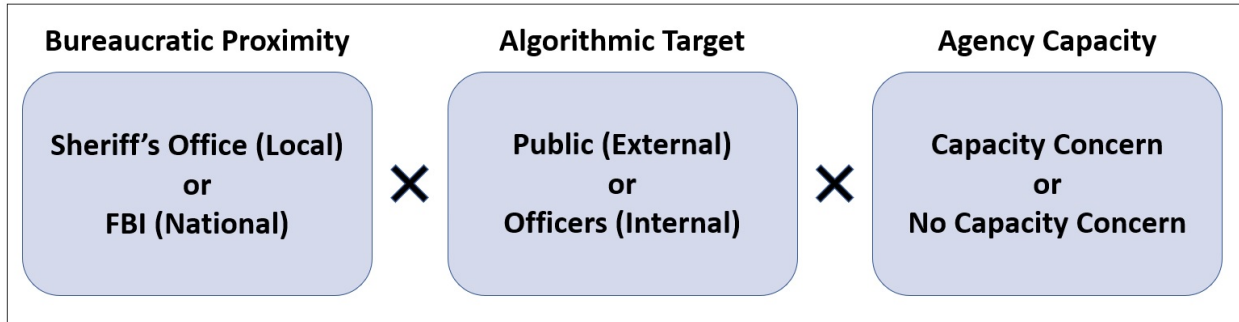


Figure 1: 2x2x2 Factorial Design of Experiment

As an example, the vignette describing the predictive policing application by the sheriff, including a capacity concern, states:

“Imagine that your local sheriff’s office is considering using an artificial intelligence (AI)-based tool to predict where crime is likely to occur, helping them determine where to send officers in order to prevent crime. But some researchers and non-profit organizations argue that these AI systems may be prone to errors and may violate individual privacy without actually reducing crime.

The sheriff’s office has also noted that it doesn’t currently have enough resources and in-house expertise to fully use the new tool, but that it hopes to in the future.”

The statement about researcher and nonprofit concerns regarding errors and privacy is included in all vignettes to present a balanced view of public discourse highlighting both benefits and concerns of AI in law enforcement.<sup>3</sup> The remaining vignettes are included in full in the appendix.

<sup>3</sup>This ‘balanced’ approach for AI vignettes follows a strategy used by O’Shaughnessy et al. (2022) and others.



## 3.2 Outcome Measures

Following the randomly-assigned vignette, we ask survey respondents two standard attitudinal outcome questions assessing respondents’ *support* for the sheriff/FBI plan to use the new AI tool and *trust* in the sheriff/FBI to use the new tool responsibly. Additionally, we include two questions aimed at measuring potential impacts on real-world behaviors. We gauge respondents’ “willingness to pay” in the form of both *taxes* and *data* for technological changes in law enforcement. We ask whether respondents would be willing to “pay more in taxes to help fund” the sheriff/FBI use of the tool and whether they would be willing to “allow access to your personal data to help improve the accuracy of the new tool.” We measure each of the four outcomes using a 5-point Likert scale, and the full wording of the survey questions is included in the appendix.

## 3.3 Analysis Strategy

We use standard OLS regression models with pre-registered specifications and robust standard errors to identify treatment effects for our experimental manipulations. With a 2x2x2 factorial design, and to increase power, we pool across non-relevant treatment groups when assessing differences between the two groups salient for each hypothesis. For example, to address our first hypothesis, the Closer to Home Hypothesis, we pool treatment groups into only two groups based on the hypothetical law enforcement agency—sheriff or FBI—and regress each of our four outcomes, individually, on that treatment indicator:

$$Outcome = \beta_0 + \beta_1 \textit{sheriff} + \gamma \mathbf{X} + \epsilon$$

where *Outcome* references one of the four outcome measures, *sheriff* is an indicator for receiving information about hypothetical local (sheriff) use of AI,  $\mathbf{X}$  refers to the vector of the covariates, and  $\epsilon$  refers to the error. The reference group receives information about hypothetical FBI use of AI. The coefficient on *sheriff* is the average treatment effect. We

include the following demographic covariates provided by Data For Progress or collected directly during our survey pre-treatment: partisanship, gender, race/ethnicity, age, education, household income, employment status, state, and engagement with local media sources.<sup>4</sup> We also construct a measure of survey taker attentiveness along a 0-2 scale based on respondent answers to two attention check questions. Moreover, we include an additional covariate measured post-treatment that assesses general support for policing agencies using a feeling thermometer. Finally, we incorporate survey weights to enable reporting of both weighted and unweighted results.

To address our other main hypotheses, we run similar models and construct similar binary indicator variables to assess treatment effects based on the algorithmic target and based on agency capacity. We further explore joint treatment effects by interacting the capacity treatment indicator with the sheriff treatment indicator. This allows us to examine whether capacity concerns differ for the extensively-resourced FBI versus local sheriff’s offices that may have significant existing capacity challenges. For our exploratory hypotheses investigating heterogeneous effects by partisanship and race, we interact the associated treatments with the relevant demographic variables.

## 4 Discussion of Results

### 4.1 Bureaucratic Proximity

Table 1 presents results assessing bureaucratic proximity as an institutional factor driving citizen attitudes toward AI use in government. As per H1, the Closer to Home Hypothesis, we expected that citizens would feel more favorably toward AI use in policing when the AI tools are used by local government officials more embedded in their community (by

---

<sup>4</sup>As a randomization check and to assess balance on the covariates, we perform F-tests of global significance, regressing each of the eight treatment group indicator variables on the covariates. With p-values ranging from 0.15 to 0.90, balance is supported by failing to reject the null of predictive covariates at the 5% significance level.

sheriffs) rather than by national government officials (FBI agents).<sup>5</sup> We find that citizens are indeed significantly more likely to trust ( $\beta = 0.27$ ,  $p < 0.001$ , standardized effect = 0.24) that sheriffs will use AI tools responsibly and are borderline more likely to support ( $\beta = 0.07$ ,  $p = 0.066$ , standardized effect = 0.06) plans for sheriffs to use AI tools. The large magnitude of the trust effect in particular supports our theory that members of the public place greater trust in community-embedded bureaucrats and that this trust extends to use of emerging technology. Still, this may not directly translate into support gains for policies to implement AI in policing, as average support ratings on the 5-point Likert scale were only 2.76 for FBI use and 2.84 for sheriff use, indicating neutral positions leaning slightly toward opposition.

	Support	Trust	Taxes	Data
	(1)	(2)	(3)	(4)
Sheriff	0.073*	0.267***	0.137***	0.124***
	(0.040)	(0.039)	(0.041)	(0.044)
N	4,198	4,198	4,198	4,198
R <sup>2</sup>	0.066	0.130	0.120	0.103
F Statistic (df = 26; 3733)	10.136***	21.431***	19.490***	16.563***

\*p < .1; \*\*p < .05; \*\*\*p < .01

Notes: With robust SEs, including covariates and weights

Table 1: Preferences for Sheriff vs. FBI Usage of AI Tools

Yet, we do find that citizens indicate they are more willing to take even costlier actions to support sheriff over FBI use of AI. Citizens are significantly more willing to share their own personal data ( $\beta = 0.12$ ,  $p = 0.005$ , standardized effect = 0.09) to facilitate sheriff use of AI, and even to pay higher taxes ( $\beta = 0.14$ ,  $p < 0.001$ , standardized effect = 0.12). To provide some context, while about 42% of respondents who received the vignette about FBI AI use were strongly opposed to providing tax funding and their personal data, only 35% of

<sup>5</sup>While the results presented pool across the external and internal use cases, we do not find significant differences between the two uses cases when interacting the bureaucratic proximity and algorithmic target treatments. Preferences for local control appear for both the predictive policing and internal review applications.

respondents who received the local sheriff vignette were strongly opposed. Increased trust may explain this stated willingness to pay through taxes and data, as there is a strong, positive correlation between trust responses and responses to the taxes and data questions ( $r = 0.59$  and  $r = 0.61$ , respectively).

To probe the robustness of these results, we consider three alternative model specifications for each outcome. Our primary model specification includes demographic covariates and uses weights provided by Data for Progress to make the results for our sample more representative of the US population. The models presented in Figure 2 include the 1) primary covariate-adjusted and weighted model, as well as 2) unadjusted and weighted, 3) adjusted and unweighted, and 4) unadjusted and unweighted models. As indicated in Figure 2, the results are highly robust to these various model specifications and provide consistent evidence that bureaucratic proximity impacts how citizens evaluate government use of AI.

As an exploratory analysis,<sup>6</sup> we investigate heterogeneous effects by partisanship of the bureaucratic proximity treatment. Across all outcomes, we find that respondents with stronger Republican identification feel significantly more favorable toward sheriff over FBI use of AI.<sup>7</sup> For example, while strong Democrats, on average, express slightly reduced trust in sheriff use of AI (difference of means of -0.09 comparing sheriff treatment to FBI treatment), Republicans and strong Republicans express substantially greater trust in sheriff use of AI (differences of +0.64 and +0.58, respectively). In general, except for strong Democrats, the local sheriff treatment produces positive treatment effects across the partisanship spectrum, with the magnitude of the effects increasing in Republican identification (see Figure A2 in the appendix). As the majority of sheriffs in the United States identify as Republican (Farris and Holman 2017), future research should explore how the partisanship of local sheriffs medi-

---

<sup>6</sup>The appendix also includes results for our pre-registered exploratory Local News Effects hypothesis.

<sup>7</sup>We fielded our survey between November 2021 and January 2022, before the FBI search for classified documents at Trump’s estate in Mar-a-Lago in August 2022 that received heavy criticism from Republicans. Therefore, Republican preferences for local control over law enforcement, and for local control over AI use, may have even increased since fielding our survey.

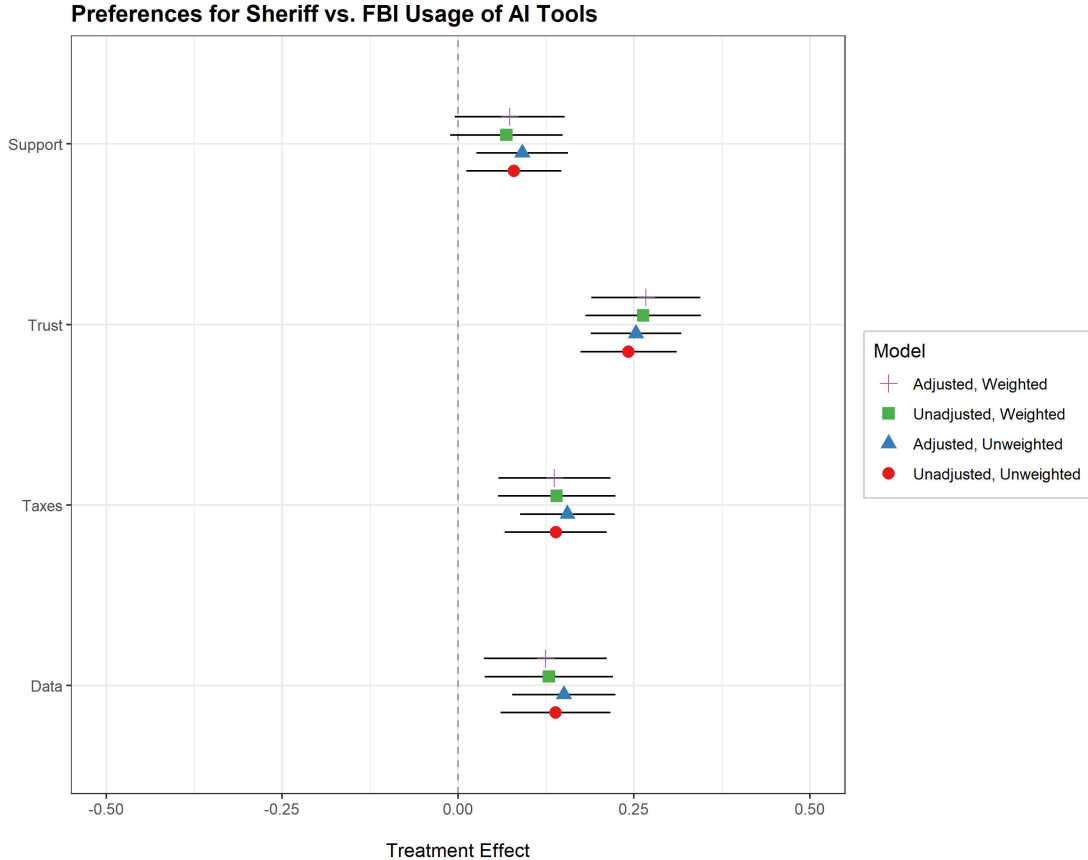


Figure 2: Robustness of Bureaucratic Proximity Results

ates these effects, and, more broadly, how the partisan composition of government interacts with bureaucratic proximity and other contextual factors to drive citizens preferences for AI in government.

## 4.2 Algorithmic Targets

Next, we present results addressing algorithmic targets and investigating H2, the Watching the Watchdogs Hypothesis. We theorized that members of the public would be more supportive of AI tools targeting law enforcement officers—that is, used internally to scrutinize officer misconduct—rather than targeting members of the public through externally-oriented predictive policing applications. Against expectations, we find little difference in public attitudes toward these different applications of AI in policing. Figure 3 reports the regression

results. Not only do the confidence intervals all overlap zero, but the point estimates are all also very close to zero, implying very similar public attitudes, in the aggregate and on average, between the external and internal use cases. This may imply that public attitudes toward AI in policing are somewhat generalized such that members of the public respond to AI as a broad tool. Indeed, some prior work (O’Shaughnessy et al. 2022) has shown that members of the public do not make strong distinctions between different applications of AI. Combined with the evidence above regarding bureaucratic proximity, and the fact that we find no significant differences when examining interactive effects between bureaucratic proximity and algorithmic targets, our results provide some new evidence that the *actors* implementing AI tools matter to the public while algorithms’ particular purposes and targets may not.

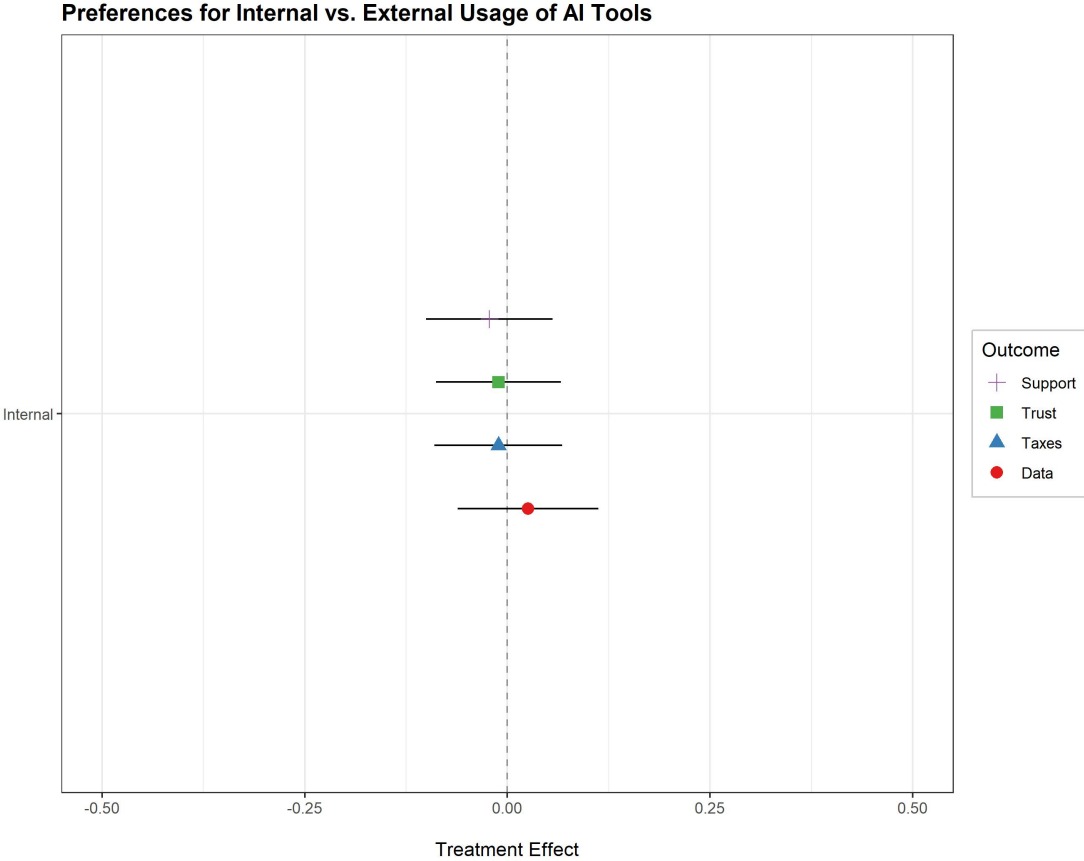


Figure 3: Public Attitudes Toward Algorithmic Targets in Policing

However, it is possible that these overall null effects mask heterogeneous subgroup differences.

We therefore examine the extent to which partisanship drives preferences over algorithmic targets in policing. We code partisanship with Strong Democrats as 1 and Strong Republicans as 5 so that the partisanship scale reflects increasing Republican identification.<sup>8</sup> Figure 4 reveals that Republicans are significantly more likely to support ( $p = 0.008$ ) predictive policing over automated internal review, and are also borderline more willing to contribute their personal data ( $p = 0.051$ ) and even pay higher taxes ( $p = 0.089$ ) to advance predictive policing over automated internal review. Differences in trust are insignificant ( $p = 0.251$ ). These results match our expectations based on the established association between conservative ideology and preferences for stricter policies and policing practices, such as predictive policing, for crime reduction. In contrast, the fact that Democrats are relatively more favorable toward applications of AI that scrutinize police officers can be explained by either concerns over the tactics that the police use on the public (preference *against* members of the public as algorithmic targets) or a growing interest in holding the police accountable for misconduct (preference *for* police officers as algorithmic targets).

Yet, we find another surprising result regarding partisan preferences over algorithmic targets. While Democrats and Republicans differ with respect to their *relative* preferences for internal versus external applications of AI in policing, we find that baseline Democratic support is higher for *both* use cases than Republicans. The average support level for strong Democrats in the automated internal review group was 3.08 compared to 2.58 for strong Republicans. Average support was also higher in the predictive policing group: 2.99 for strong Democrats compared to 2.7 for strong Republicans. Both differences are statistically significant ( $p < 0.001$ ). This appears to stand in contrast to Democrats' lower feelings in general toward the police, approximately 16 points lower on average ( $p < 0.001$ ) on a 100-point feeling thermometer that we used on our survey. While other recent work finds that Democrats typically have stronger support for AI (O'Shaughnessy et al. 2022), it is notable that this

---

<sup>8</sup>Note that this differs from our pre-registration, which considers partisanship as binary Republican identification. We make this adjustment to better illustrate the relationship under study.

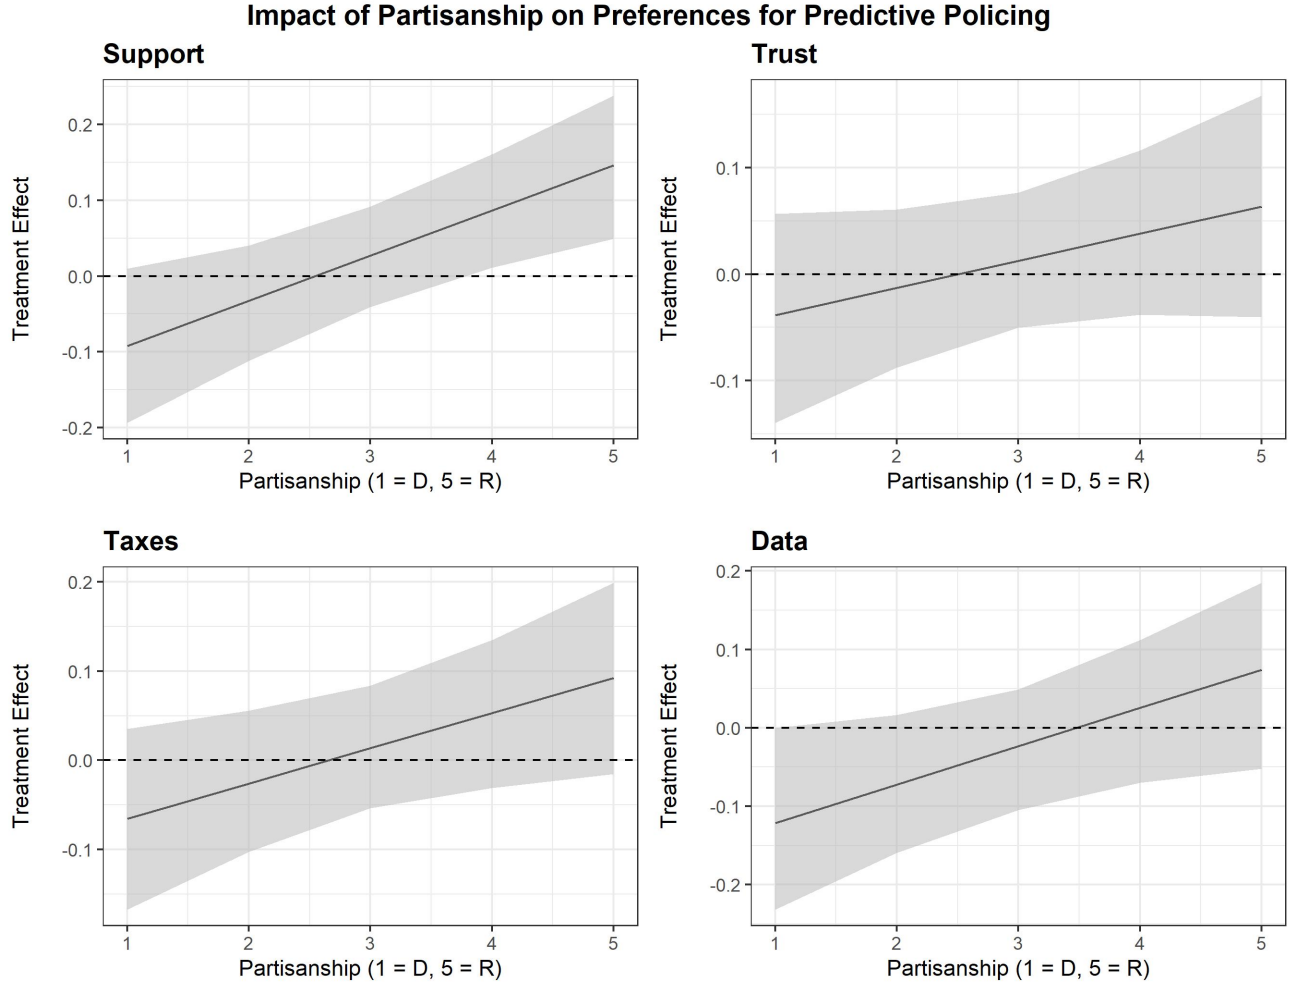


Figure 4: Heterogeneous Effects by Partisanship

extends even to the domain of policing. In this context, Democrats' higher support for AI may override their skepticism of police such that ideological liberals support AI even in arguably controversial settings like predictive policing. Another possibility is that, while predictive policing has come under scrutiny in certain circles, such as within academia and civil society, these criticisms may reflect a relatively narrow set of elite preferences not mirrored by Democrats in the broader public. Finally, a possible explanation or implication is that Democrats may trade off in favor of algorithmic bias over human bias in policing.

Next, we analyze heterogeneous support for predictive policing by race. We hypothesized that individuals from underrepresented minority groups, particularly Black Americans,



would be less supportive of AI tools used for predictive policing, compared to non-Black Americans.<sup>9</sup> We find only modest evidence to support this hypothesis. Figure 5 shows the preferences of Black and non-Black respondents for predictive policing versus automated internal review. Despite significant recent attention to police accountability and nationwide calls for reform, along with concerns about the over-policing of Black communities that may be targeted by predictive policing, Black individuals are not particularly more likely to oppose members of the public as algorithmic targets. While Black support for predictive policing relative to automated internal review is significantly lower ( $p = 0.049$ ), trust, willingness to share data, and willingness to pay taxes are not. Some of the explanations proposed above to explain higher Democratic support for AI in policing than expected may apply here as well.

### 4.3 Agency Capacity

Finally, we hypothesized that citizens would be more supportive of AI in policing when the capacity of law enforcement agencies is not a concern, given the novelty and uncertainty associated with new AI tools and the technical expertise and training needed to use them successfully and responsibly. However, as displayed in Figure 6, we find that stated concerns about existing agency capacity have little impact on citizen attitudes toward AI in policing. Agencies that acknowledge that they do not “currently have enough resources and in-house expertise to fully use” new AI tools garner similar levels of support and trust as those without such stated concerns. While the point estimates indicate slight increases in willingness to provide tax funding and data to low-capacity agencies, the effects are not statistically significant. While a theoretically-important institutional and contextual factor, agency capacity appears to have little bearing on public attitudes.

Moreover, members of the public also do not view agency capacity as a bigger concern for

---

<sup>9</sup>We make another slight alteration to our pre-registration, focusing on Black citizens rather than non-white citizens as the subgroup of interest given the particular relationship between Black Americans and policing institutions.

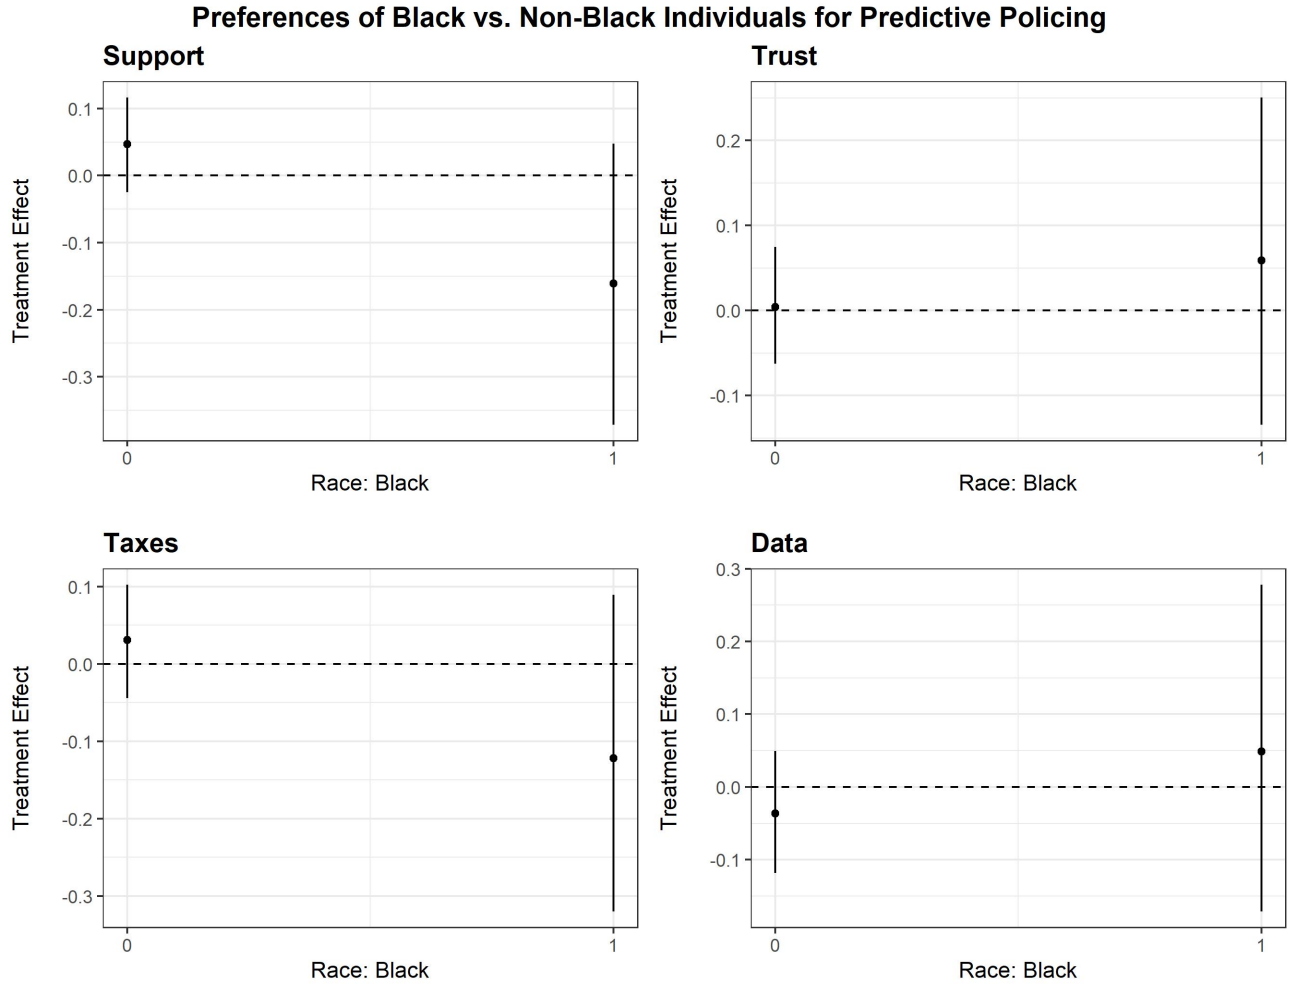


Figure 5: Heterogeneous Effects by Race

(traditionally relatively under-resourced) local sheriffs compared to the FBI. The regression results in Table 2 indicate no significant differences in the impact of capacity concerns on public attitudes toward sheriff versus FBI use of AI tools. We also find no significant differences in reactions to capacity concerns for the intended use cases of predictive policing and automated internal review. Therefore, across actors and use cases, we find that capacity is not especially salient to the public, despite theoretical importance and scholarly and policymaker attention.

This finding suggests a pair of related concerns: it is not clear that citizens are more skeptical of agencies that are less able to implement AI systems, and it is also not evident that

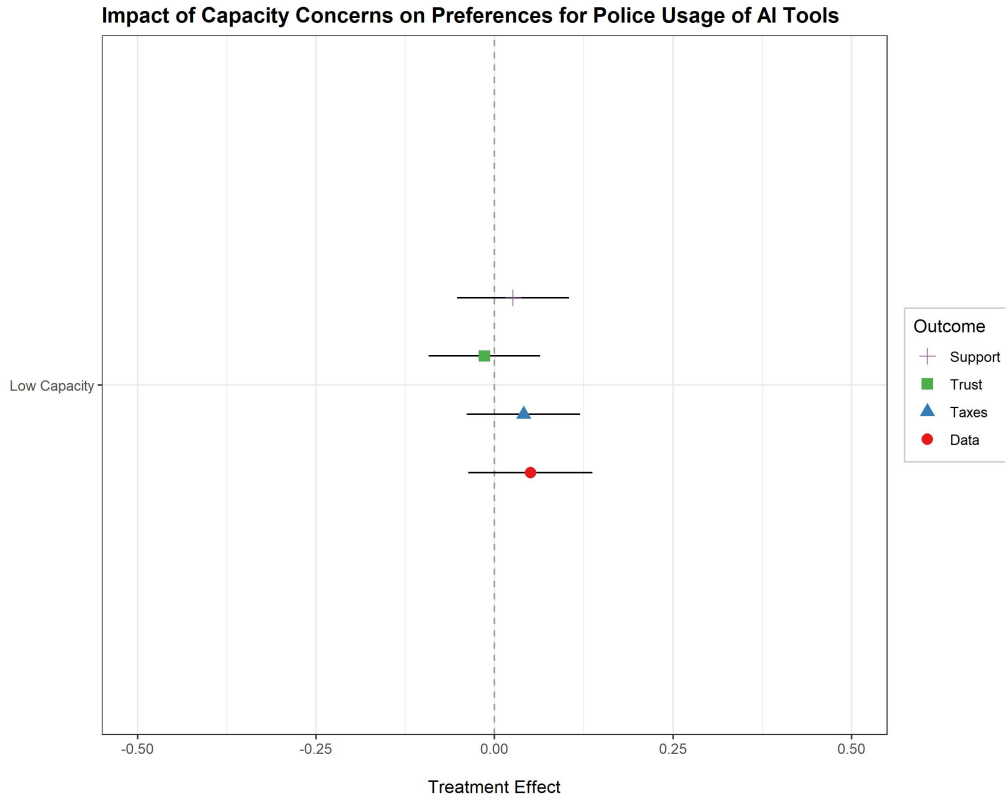


Figure 6: Agency Capacity Concerns and Public Support for AI in Policing

	Support (1)	Trust (2)	Taxes (3)	Data (4)
Low Capacity	0.024 (0.057)	-0.023 (0.059)	0.031 (0.056)	-0.013 (0.062)
Sheriff	0.069 (0.056)	0.249*** (0.056)	0.125** (0.057)	0.057 (0.062)
Low Capacity x Sheriff	0.009 (0.080)	0.034 (0.079)	0.027 (0.081)	0.132 (0.089)
N	4,198	4,198	4,198	4,198
R <sup>2</sup>	0.066	0.130	0.120	0.105
F Statistic (df = 28; 3731)	9.434***	19.901***	18.161***	15.565***

\*p < .1; \*\*p < .05; \*\*\*p < .01

Notes: With robust SEs, including covariates and weights

Table 2: Capacity Concerns and Preferences for Sheriff vs. FBI Use of AI Tools

citizens are more likely to support agencies that are. This disconnect between the ability of law enforcement institutions to implement AI tools responsibly and effectively and citizens' internalization of this key institutional factor suggests a barrier to accountability and a potential limitation on public participation in AI policymaking and governance. Efforts should be directed toward improving public understanding of the organizational, not only the technical, prerequisites necessary for government implementation of AI.

## 5 Conclusion

This study, involving a survey experiment administered to 4,200 American adults, examines public responses to applications of advanced technology in policing practices. While prior work in this emerging domain has largely focused on features of the technology itself, we emphasize three institutional factors that research suggests may influence AI implementation and which implicate the relationship between government and the public: bureaucratic proximity, algorithmic targets, and agency capacity.

Figure 7 presents the findings associated with our main pre-registered hypotheses. First, in line with the literature on trust in local government and our theoretical expectations regarding community embeddedness, we find a strong public preference for local over national use of AI in policing, especially amongst Republicans. Second, we find no average preference for applying AI to scrutinize members of the public versus police officers themselves. However, Republicans have a greater relative preference for predictive policing and Democrats for detecting officer misconduct, and Democrats have greater overall support for both use cases. Finally, we find that agency capacity, or lack thereof, has little bearing on public support, despite its importance for determining whether policing institutions can administer AI systems effectively and responsibly. The results imply that members of the public are particularly attuned to the identity and closeness of the actors and institutions implementing AI systems. The results also suggest two fruitful pathways for future work: 1) examining

what shapes the public’s understanding of government capacity, especially with respect to emerging technologies, and 2) exploring whether and when limited public awareness of institutional capacity in AI implementation poses a threat to accountability and democratic governance.

Our study is nonetheless limited in its ability to answer many key questions in this domain. Due to power considerations, we opted to examine only three institutional factors associated with government use of AI. However, future work should explore more, and it may be especially valuable to evaluate how members of the public compare and trade off between contextual factors and the cost effectiveness benefits of AI tools. While our experimental design offers strong internal validity, we also recognize limitations pertaining to external validity and generalizability. For ethical considerations and in order to cleanly identify treatment effects, we measure responses to hypothetical scenarios. This somewhat reduces the real-world relevance of our treatments, although they are based on actual use cases. Moreover, as is the case with vignette experiments, we examine research questions embedded in a complex context, and the specific choices around treatments and their wording reflect that bundling. For example, Republican support for local use of AI may be driven by political connotations associated with the specific role of sheriffs (versus, say, police chiefs more common in urban contexts). Higher baseline Democratic support for predictive policing could reflect concerns about human bias in policing, leading to a relative preference for automated systems and a tradeoff in favor of potential algorithmic bias.

Therefore, public opinion on these use cases in law enforcement may not extend straightforwardly to other AI use cases or other policy sectors, and future work is needed to inventory and explore additional institutional and contextual factors driving public support for AI in government. For example, when members of the public have prior opinions both about technical factors (e.g., towards AI) and towards institutional ones (e.g., baseline trust in police agencies), a key question is whether potentially complementary or competing priorities act

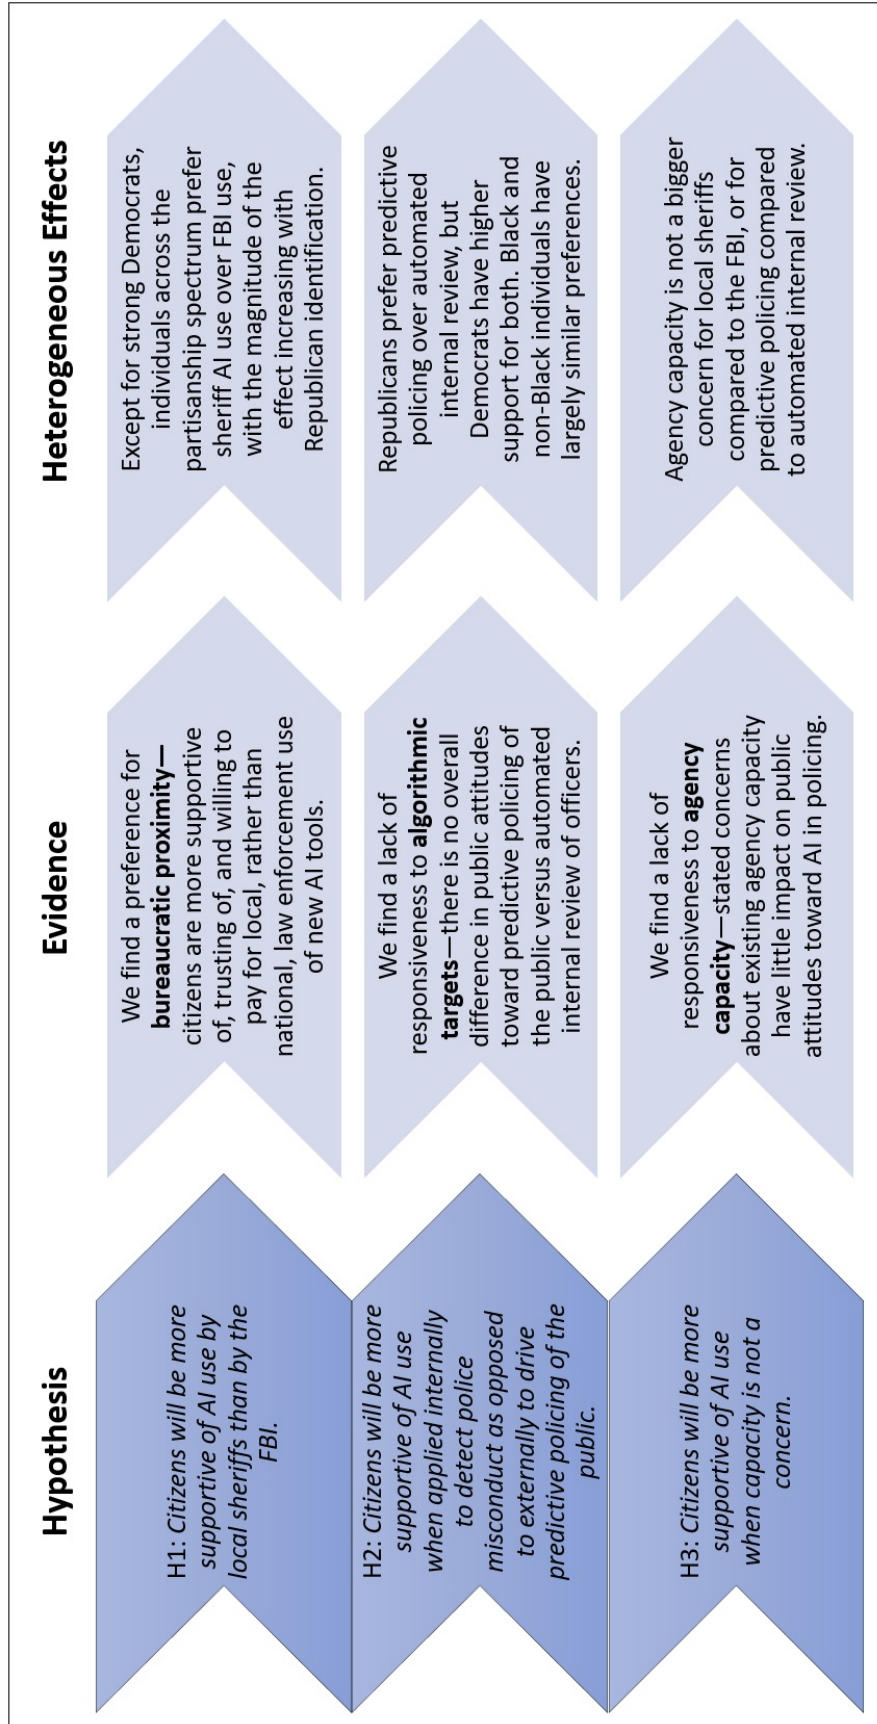


Figure 7: Summary of Findings

in an additive or more complex interactive fashion. That is, technical and contextual factors need to be understood individually and jointly, including with respect to their possible politicization in the public eye. In addition to our interest in bureaucratic proximity, agency capacity, and algorithmic targets, some relevant institutional and contextual factors worth exploring include: the selection process for bureaucrats and government officials (e.g., election, appointment, or career service), agency oversight and accountability structures, and pre-existing levels of human discretion in government (on the government side), as well as perceived personal impacts and symbolic effects, attitudes toward target populations, and concerns about procedural fairness (on the citizen side). Finally, qualitative and mixed-methods research could also supplement this experimental work with in-depth interviews of bureaucrats and members of the public to better understand the foundations for these actors' and stakeholders' preferences for AI in government.

Our results also imply an interesting opportunity and corresponding challenge for practitioners. On the one hand, local governments may be well positioned to serve at the forefront of government use of AI due to broad support across the political spectrum, at least relative to a national-level alternative, and due to the many street-level implementation decisions needed to apply AI systems effectively and carefully. On the other hand, lack of local government capacity and socio-technical expertise to apply AI systems could render local leadership problematic. One implication is that it may be effective to assign a larger portion of the substantial and growing federal funding for AI systems to local rather than federal government actors. That is, while a substantial focus of AI policy discourse has centered on national economic competitiveness, subnational governments have increased their activity and this paper demonstrates their heightened relevance to the public. Federal actors may find they are able to better achieve even national policy goals by paying increased attention to subnational policymakers and local bureaucrats.

In contrast, federal government agencies face an inverse challenge, as they are more likely

to build appropriate expertise in AI but lack public trust and buy-in. This suggests that developing federal governance that is trustworthy should therefore be a priority. To achieve these goals, federal government actors could identify creative strategies to build trust beyond physical proximity. For example, increased communication emphasizing local needs and issues, more frequent contact with local communities through community workshops and constituency visits, and exchange programs with local governments or educational institutions could help build a sense of proximity and trust at the federal level.

The finding that public attitudes are invariant to critical factors like the government's actual capacity to implement AI is concerning, and may reflect that a greater focus in scholarship, civil society, and media to date has been the advantages and limitations associated with *technical* features of AI. The results of this study instead encourage broader attention to institutional factors like capacity, funding, and personnel. These issues are already quite salient to many public administrators but less so to the public, in part because technical aspects of AI (like algorithmic bias or the black box problem) may be perceived as newer or more flashy.

Along these lines, efforts to advance public communication, education, and participation around government use of AI have received significant attention, in part to secure the adoption of AI in society, as well as public input and legitimacy. Our findings suggest that these efforts could benefit from re-centering the importance of influential contextual factors in AI adoption. Building public understanding of these institutional issues can support the public's ability to attend to government needs and challenges around AI adoption, so that decision-makers can more confidently consult public opinion when shaping the direction of AI in government.



## References

- Abramowitz, Alan I. 2022. “The Polarized American Electorate: The Rise of Partisan-Ideological Consistency and Its Consequences.” *Political Science Quarterly* 137(4):645–674.
- Access Now. 2021. *Open letter calling for a global ban on biometric recognition technologies that enable mass and discriminatory surveillance*. Brooklyn, New York: . More than 175 signatories.  
**URL:** <https://www.accessnow.org/ban-biometric-surveillance/>
- Al-Hujran, Omar, Mutaz M. Al-Debei, Akemi Chatfield and Mahmoud Migdadi. 2015. “The Imperative of Influencing Citizen Attitude toward E-Government Adoption and Use.” *Computers in Human Behavior* 53:189–203.
- Alon-Barkat, Saar and Madalina Busuioc. 2022. “Human-AI Interactions in Public Sector Decision-Making: ‘Automation Bias’ and ‘Selective Adherence’ to Algorithmic Advice.” *Journal of Public Administration Research and Theory* .
- Berk, Richard A. 2021. “Artificial Intelligence, Predictive Policing, and Risk Assessment for Law Enforcement.” *Annual Review of Criminology* 4(1):209–237.
- Boudreau, Cheryl, Scott A. MacKenzie and Daniel J. Simmons. 2019. “Police Violence and Public Perceptions: An Experimental Study of How Information and Endorsements Affect Support for Law Enforcement.” *The Journal of Politics* .
- Bozeman, Barry. 2002. “Public-Value Failure: When Efficient Markets May Not Do.” *Public Administration Review* 62(2):145–161.
- Bradford, Ben, Julia A Yesberg, Jonathan Jackson and Paul Dawson. 2020. “Live facial recognition: Trust and legitimacy as predictors of public support for police use of new technology.” *The British Journal of Criminology* 60(6):1502–1522.
- Bromberg, Daniel E, Étienne Charbonneau and Andrew Smith. 2020. “Public support for facial recognition via police body-worn cameras: Findings from a list experiment.” *Government Information Quarterly* 37(1):101415.
- Bryson, John M., Barbara C. Crosby and Laura Bloomberg. 2014. “Public Value Governance: Moving Beyond Traditional Public Administration and the New Public Management.” *Public Administration Review* 74(4):445–456.

- Buhmann, Alexander and Christian Fieseler. 2021. "Towards a deliberative framework for responsible innovation in artificial intelligence." *Technology in Society* 64:101475.
- Bullock, Justin B. 2019. "Artificial Intelligence, Discretion, and Bureaucracy." *The American Review of Public Administration* 49(7):751–761.  
**URL:** <https://doi.org/10.1177/0275074019856123>
- Busuioc, Madalina. 2021. "Accountable Artificial Intelligence: Holding Algorithms to Account." *Public Administration Review* 81(5):825–836.  
**URL:** <https://onlinelibrary.wiley.com/doi/abs/10.1111/puar.13293>
- Charbonneau, Étienne and Carey Doberstein. 2020. "An Empirical Assessment of the Intrusiveness and Reasonableness of Emerging Work Surveillance Technologies in the Public Sector." *Public Administration Review* 80(5):780–791.  
**URL:** <https://onlinelibrary.wiley.com/doi/abs/10.1111/puar.13278>
- Compton, Mallory E, Matthew M Young, Justin B Bullock and Robert Greer. 2022. "Administrative Errors and Race: Can Technology Mitigate Inequitable Administrative Outcomes?" *Journal of Public Administration Research and Theory* p. muac036.  
**URL:** <https://academic.oup.com/jpart/advance-article/doi/10.1093/jpart/muac036/6680693>
- de Boer, Noortje and Nadine Raaphorst. 2021. "Automation and discretion: Explaining the effect of automation on how street-level bureaucrats enforce." *Public Management Review* .
- Denhardt, Robert B. 2002. "Trust as capacity: The role of integrity and responsiveness." *Public Organization Review* 2(1):65–76.
- Denhardt, Robert B. and Janet Vinzant Denhardt. 2000. "The New Public Service: Serving Rather Than Steering." *Public Administration Review* 60(6):549–559.
- Farris, Emily M and Mirya R Holman. 2017. "All politics is local? County sheriffs and localized policies of immigration enforcement." *Political Research Quarterly* 70(1):142–154.
- Ferguson, Andrew G. 2017a. "Policing Predictive Policing." *Washington University Law Review* 94(5):82.
- Ferguson, Andrew Guthrie. 2017b. *The rise of big data policing*. New York University Press.
- Fine, Adam D, Sachiko Donley, Caitlin Cavanagh and Elizabeth Cauffman. 2020. "Youth perceptions of law enforcement and worry about crime from 1976 to 2016." *Criminal justice and behavior* 47(5):564–581.

- Gaozhao, Dongfang, James Wright II and Mylah Gainey. 2022. "Bureaucrat or Artificial Intelligence: People's Preferences and Perceptions of Government Service." *Public Management Review* .
- Grimmelikhuijsen, Stephan. 2023. "Explaining Why the Computer Says No: Algorithmic Transparency Affects the Perceived Trustworthiness of Automated Decision-Making." *Public Administration Review* 83(2):241–262.  
**URL:** <https://onlinelibrary.wiley.com/doi/abs/10.1111/puar.13483>
- Ho, Alfred Tat-Kei and Wonhyuk Cho. 2017. "Government communication effectiveness and satisfaction with police performance: A large-scale survey study." *Public Administration Review* 77(2):228–239.
- Jakobsen, Morten, Oliver James, Donald Moynihan and Tina Nabatchi. 2019. "JPART Virtual Issue on Citizen-State Interactions in Public Administration Research." *Journal of Public Administration Research and Theory* 29(4):e8–e15.
- Johansson, Jörgen, Michel Thomsen and Maria Åkesson. 2022. "Public value creation and robotic process automation: normative, descriptive and prescriptive issues in municipal administration." *Transforming Government: People, Process and Policy* ahead-of-print(ahead-of-print).  
**URL:** <https://doi.org/10.1108/TG-11-2021-0193>
- Kaun, Anne, Anders Olof Larsson and Anu Masso. 2023. "Automating public administration: citizens' attitudes towards automated decision-making across Estonia, Sweden, and Germany." *Information, Communication Society* 0(0):1–19.
- Kelleher, Christine A. and Jennifer Wolak. 2007. "Explaining Public Confidence in the Branches of State Government." *Political Research Quarterly* 60(4):707–721.
- König, Pascal D., Julia Felfeli, Anja Achtziger and Georg Wenzelburger. 2022. "The importance of effectiveness versus transparency and stakeholder involvement in citizens' perception of public sector algorithms." *Public Management Review* 0(0):1–22.
- Leland, Suzanne, Jacqueline Chattopadhyay, Cherie Maestas and Jaclyn Piatak. 2021. "Policy venue preference and relative trust in government in federal systems." *Governance* 34(2):373–393.
- McCandless, Sean. 2018. "Improving community relations: How police strategies to improve accountability for social equity affect citizen perceptions." *Public Integrity* 20(4):370–385.

- Meijer, Albert, Lukas Lorenz and Martijn Wessels. 2021. "Algorithmization of Bureaucratic Organizations: Using a Practice Lens to Study How Context Shapes Predictive Policing Systems." *Public Administration Review* 81(5):837–846.  
**URL:** <https://onlinelibrary.wiley.com/doi/abs/10.1111/puar.13391>
- Merola, Linda M and Cynthia Lum. 2014. "Predicting public support for the use of license plate recognition technology by police." *Police Practice and Research* 15(5):373–388.
- Merola, Linda M, Cynthia Lum and Ryan P Murphy. 2019. "The impact of license plate recognition technology (LPR) on trust in law enforcement: A survey-experiment." *Journal of Experimental Criminology* 15(1):55–66.
- Miller, Susan M and Lael R Keiser. 2021. "Representative Bureaucracy and Attitudes Toward Automated Decision Making." *Journal of Public Administration Research and Theory* 31(1):150–165.  
**URL:** <https://academic.oup.com/jpart/article/31/1/150/5848482>
- Mitchell, Terence R, Brooks C Holtom, Thomas W Lee, Chris J Sablinski and Miriam Erez. 2001. "Why people stay: Using job embeddedness to predict voluntary turnover." *Academy of management journal* 44(6):1102–1121.
- Mourtgos, Scott M and Ian T Adams. 2020. "Assessing public perceptions of police use-of-force: Legal reasonableness and community standards." *Justice Quarterly* 37(5):869–899.
- Mrozla, Thomas. 2021. "Procedural and distributive justice: Effects on attitudes toward body-worn cameras." *International Journal of Police Science & Management* 23(3):317–327.
- Murphy, Kristina, Lorraine Mazerolle and Sarah Bennett. 2014. "Promoting Trust in Police: Findings from a Randomised Experimental Field Trial of Procedural Justice Policing." *Policing and Society* 24(4):405–424.
- Neshkova, Milena I. and Hai (David) Guo. 2012. "Public Participation and Organizational Performance: Evidence from State Agencies." *Journal of Public Administration Research and Theory* 22(2):267–288.
- Neyroud, Peter and Emma Disley. 2008. "Technology and policing: Implications for fairness and legitimacy." *Policing: A Journal of Policy and Practice* 2(2):226–232.
- Osborne, Stephen P. 2006. "The New Public Governance?" *Public Management Review* 8(3):377–387.
- Ouchchy, Leila, Allen Coin and Veljko Dubljević. 2020. "AI in the Headlines: The Portrayal of the Ethical Issues of Artificial Intelligence in the Media." *AI & Society* 35(4):927–936.

- O'Shaughnessy, Matthew R, Daniel S Schiff, Lav R Varshney, Christopher J Rozell and Mark A Davenport. 2022. "What governs attitudes toward artificial intelligence adoption and governance?" *Science and Public Policy* p. scac056. 10.31219/osf.io/pkeb8.
- Porumbescu, Gregory A, Milena I Neshkova and Meghan Huntoon. 2019. "The effects of police performance on agency trustworthiness and citizen participation." *Public Management Review* 21(2):212–237.
- Powers, Kathleen J. and Fred Thompson. 1994. "Managing Coprovision: Using Expectancy Theory to Overcome the Free-Rider Problem\*." *Journal of Public Administration Research and Theory* 4(2):179–196.
- Reddick, Christopher G, Akemi Takeoka Chatfield and Patricia A Jaramillo. 2015. "Public opinion on National Security Agency surveillance programs: A multi-method approach." *Government Information Quarterly* 32(2):129–141.
- Richardson, Lilliard E., Jr., David M. Konisky and Jeffrey Milyo. 2012. "Public Approval of U.S. State Legislatures." *Legislative Studies Quarterly* 37(1):99–116.
- Sakiyama, Mari, Terance D Miethe, Joel D Lieberman, Miliakeala SJ Heen and Olivia Tuttle. 2017. "Big hover or big brother? Public attitudes about drone usage in domestic policing activities." *Security Journal* 30(4):1027–1044.
- Schiff, Daniel S., Kaylyn Jackson Schiff and Patrick Pierson. 2021. "Assessing Public Value Failure in Government Adoption of Artificial Intelligence." *Public Administration* .
- Schlicker, Nadine, Markus Langer, Sonja K. Ötting, Kevin Baum, Cornelius J. König and Dieter Wallach. 2021. "What to expect from opening up 'black boxes'? Comparing perceptions of justice between human and automated agents." *Computers in Human Behavior* 122:106837.  
**URL:** <https://www.sciencedirect.com/science/article/pii/S0747563221001606>
- Schneider, Anne and Helen Ingram. 1993. "Social construction of target populations: Implications for politics and policy." *American political science review* 87(2):334–347.
- Schneider, S. K., W. G. Jacoby and D. C. Lewis. 2011. "Public Opinion Toward Intergovernmental Policy Responsibilities." *Publius: The Journal of Federalism* 41(1):1–30.  
**URL:** <https://academic.oup.com/publius/article-lookup/doi/10.1093/publius/pjq036>

- Selten, Friso, Marcel Robeer and Stephan Grimmelikhuijsen. 2023. “‘Just like I thought’: Street-level bureaucrats trust AI recommendations if they confirm their professional judgment.” *Public Administration Review* n/a(n/a).
- URL:** <https://onlinelibrary.wiley.com/doi/abs/10.1111/puar.13602>
- Silva, Andrea, Diego Esparza, Valerie Martinez-Ebers and Regina Branton. 2020. “Perceived Police Performance, Racial Experiences, and Trust in Local Government.” *Politics, Groups, and Identities* 0(0):1–24.
- Tyler, Tom R. 2004. “Enhancing police legitimacy.” *The annals of the American academy of political and social science* 593(1):84–99.
- Tyler, Tom R, Robert J Boeckmann, Heather J Smith and Yuen J Huo. 2019. *Social justice in a diverse society*. Routledge.
- Van Ryzin, Gregg G. 2007. “Pieces of a Puzzle: Linking Government Performance, Citizen Satisfaction, and Trust.” *Public Performance & Management Review* 30(4):521–535.
- Vogl, Thomas M., Cathrine Seidelin, Bharath Ganesh and Jonathan Bright. 2020. “Smart Technology and the Emergence of Algorithmic Bureaucracy: Artificial Intelligence in UK Local Authorities.” *Public Administration Review* 80(6):946–961.
- URL:** <https://onlinelibrary.wiley.com/doi/abs/10.1111/puar.13286>
- Vrontis, Demetris, Michael Christofi, Vijay Pereira, Shlomo Tarba, Anna Makrides and Eleni Trichina. 2021. “Artificial intelligence, robotics, advanced technologies and human resource management: a systematic review.” *The International Journal of Human Resource Management* 33:1–30.
- Wenzelburger, Georg, Pascal D. König, Julia Felfeli and Anja Achtziger. 2022. “Algorithms in the Public Sector. Why context matters.” *Public Administration* .
- URL:** <https://onlinelibrary.wiley.com/doi/abs/10.1111/padm.12901>
- West, Darrell M. 2004. “E-Government and the Transformation of Service Delivery and Citizen Attitudes.” *Public Administration Review* 64(1):15–27.
- West, Jonathan P, Casey A Klofstad, Joseph E Uscinski and Jennifer M Connolly. 2019. “Citizen support for domestic drone use and regulation.” *American Politics Research* 47(1):119–151.
- Wirtz, Bernd W., Paul F. Langer and Carolina Fenner. 2021. “Artificial Intelligence in the Public Sector - a Research Agenda.” *International Journal of Public Administration* 0(0):1–26.

Wright, James E, Dongfang Gaozhao, Kenneth Dukes and Da'Shay Templeton. 2023. "The power of protest on policing: Black Lives Matter protest and civilian evaluation of the police." *Public Administration Review* 83(1):130–143.

Young, Matthew M, Justin B Bullock and Jesse D Lecy. 2019. "Artificial Discretion as a Tool of Governance: A Framework for Understanding the Impact of Artificial Intelligence on Public Administration." *Perspectives on Public Management and Governance* p. gvz014.

**URL:** <https://academic.oup.com/ppmg/advance-article/doi/10.1093/ppmgov/gvz014/5602198>

Zuiderwijk, Anneke, Yu-Che Chen and Fadi Salem. 2021. "Implications of the use of artificial intelligence in public governance: A systematic literature review and a research agenda." *Government Information Quarterly* 38(3):101577.

# A Supplementary Information for “Institutional Factors Driving Citizen Perceptions of AI in Government: Evidence from a Survey Experiment on Policing”

## Contents

A.1	Survey Instrument . . . . .	2
A.1.1	Vignette Treatments . . . . .	2
A.1.2	Outcomes . . . . .	3
A.2	Power . . . . .	4
A.3	Additional Results . . . . .	5



## A.1 Survey Instrument

### A.1.1 Vignette Treatments

Participants are randomly assigned with equal probability to one of the following:

1. Imagine that your local sheriff's office is considering using an artificial intelligence (AI)-based tool to predict where crime is likely to occur, helping them determine where to send officers in order to prevent crime. But some researchers and non-profit organizations argue that these AI systems may be prone to errors and may violate individual privacy without actually reducing crime.

The sheriff's office has also noted that it doesn't currently have enough resources and in-house expertise to fully use the new tool, but that it hopes to in the future.

2. Imagine that your local sheriff's office is considering using an artificial intelligence (AI)-based tool to predict where crime is likely to occur, helping them determine where to send officers in order to prevent crime. But some researchers and non-profit organizations argue that these AI systems may be prone to errors and may violate individual privacy without actually reducing crime.

3. Imagine that your local sheriff's office is considering using an artificial intelligence (AI)-based tool to help them evaluate and detect police misconduct. But some researchers and non-profit organizations argue that these AI systems may be prone to errors and may violate individual privacy without actually reducing police misconduct.

The sheriff's office has also noted that it doesn't currently have enough resources and in-house expertise to fully use the new tool, but that it hopes to in the future.

4. Imagine that your local sheriff's office is considering using an artificial intelligence (AI)-based tool to help them evaluate and detect police misconduct. But some researchers and non-profit organizations argue that these AI systems may be prone to errors and may violate individual privacy without actually reducing police misconduct.

5. Imagine that the Federal Bureau of Investigation (FBI) is considering using an artificial intelligence (AI)-based tool to predict where crime is likely to occur, helping them determine where to send agents in order to prevent crime. But some researchers and non-profit organizations argue that these AI systems may be prone to errors and may violate individual privacy without actually reducing crime.

The FBI has also noted that it doesn't currently have enough resources and in-house expertise to fully use the new tool, but that it hopes to in the future.

6. Imagine that the Federal Bureau of Investigation (FBI) is considering using an artificial intelligence (AI)-based tool to predict where crime is likely to occur, helping them determine where to send agents in order to prevent crime. But some researchers and non-profit organizations argue that these AI systems may be prone to errors and may violate individual privacy without actually reducing crime.

7. Imagine that the Federal Bureau of Investigation (FBI) is considering using an artificial intelligence (AI)-based tool to help them evaluate and detect agent misconduct. But some researchers and non-profit organizations argue that these AI systems may be prone to errors and may violate individual privacy without actually reducing agent misconduct.

The FBI has also noted that it doesn't currently have enough resources and in-house expertise to fully use the new tool, but that it hopes to in the future.

8. Imagine that the Federal Bureau of Investigation (FBI) is considering using an artificial intelligence (AI)-based tool to help them evaluate and detect agent misconduct. But some researchers and non-profit organizations argue that these AI systems may be prone to errors and may violate individual privacy without actually reducing agent misconduct.

### A.1.2 Outcomes

To what extent do you support the [Sheriff's office's | FBI's] plan to use this new tool? [Strongly oppose to Strongly support]

To what extent do you trust that the [Sheriff's office | FBI] will use this new tool responsibly? [Strongly distrust to Strongly trust]

To what extent would you be willing to pay more in taxes to help fund the [Sheriff's office's | FBI's] use of this tool. [Definitely would not to Definitely would]

The [Sheriff's office's | FBI's] proposed tool will require data from civilians, such as your history of traffic stops, age, race, or gender, to make the AI system as accurate as possible. To what extent would you be willing to allow access to your personal data to help improve the accuracy of the new tool? [Definitely would not to Definitely would]

## A.2 Power

To assess power and sample size needs, we perform 1000 simulations for each of a range of possible sample sizes (500 to 4000) in order to estimate our power to detect significant treatment effects. We do so for one of our main outcomes (e.g., the support outcome) and pool treatment groups into only two groups (corresponding to the main analyses in H1, H2, and H3). We use an error rate of  $\alpha = .05$ , and conservatively assume a small standardized effect of  $f = 0.1$ . We find that we have sufficient power (80%) with approximately 3,100 subjects. In partnership with Data for Progress, we were able to recruit 4,198 respondents. However, we lack power to fully address our exploratory hypotheses, e.g., hypotheses that require splitting the sample into four or more treatment groups to make comparisons.

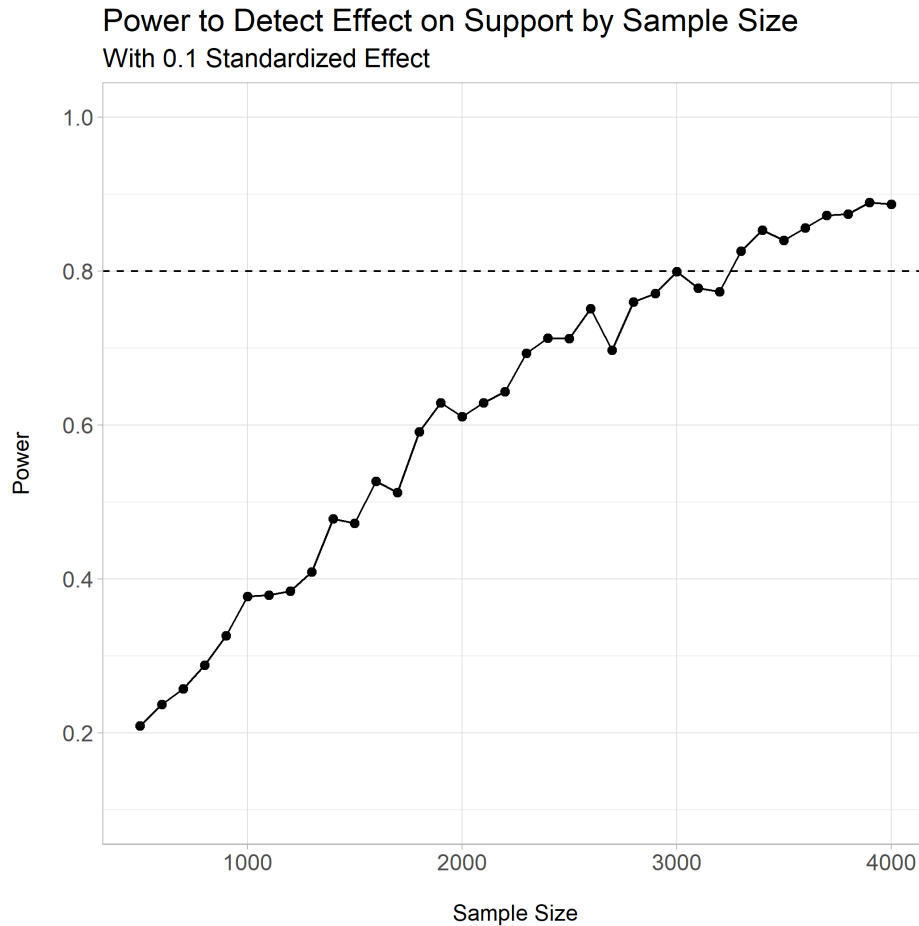


Figure A1: Power

### A.3 Additional Results

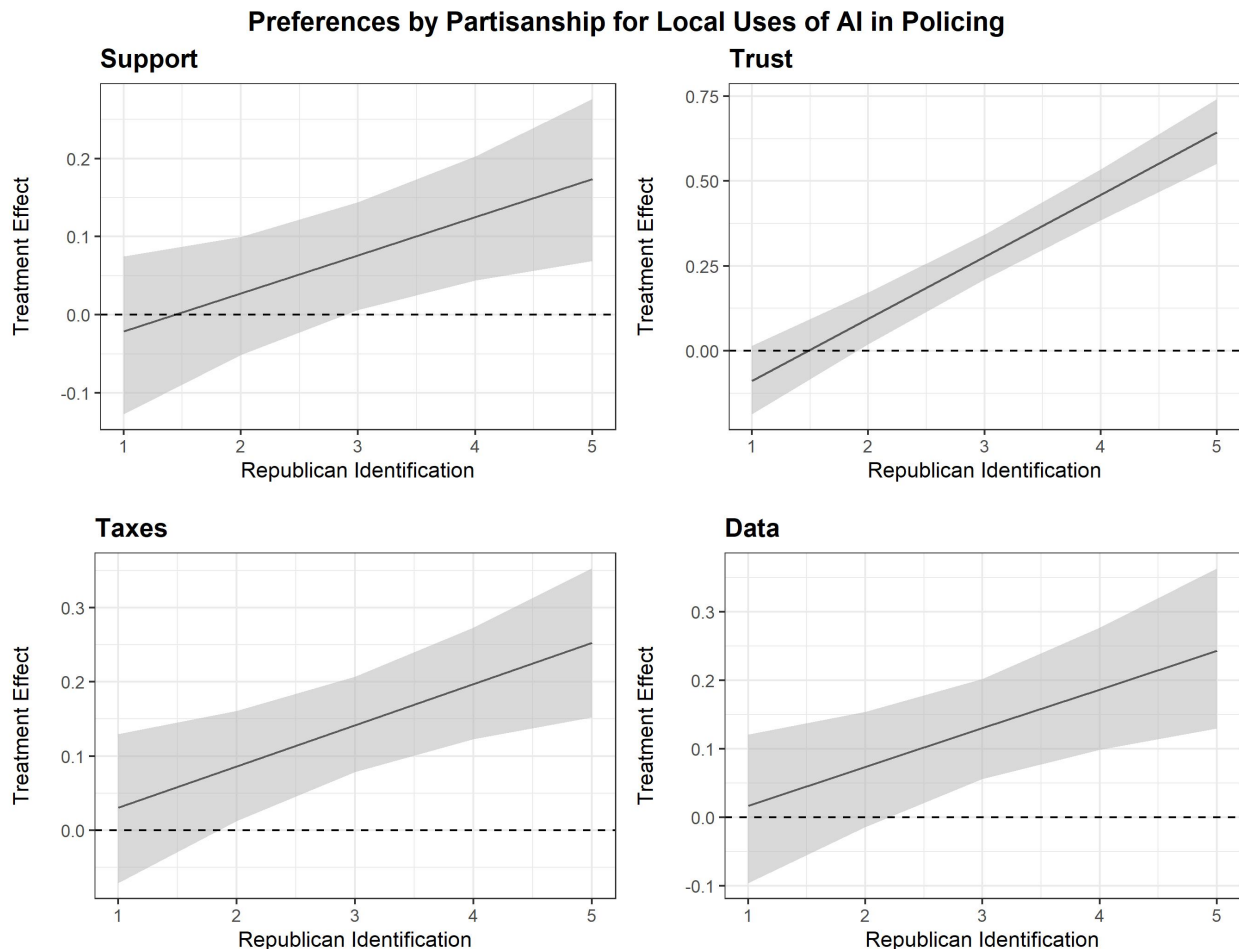


Figure A2: Heterogeneous Effects of the Bureaucratic Proximity Treatment by Partisanship

We also hypothesized that local news consumption would moderate trust in local sheriff use of AI. We expected that citizens who are more attentive to local news would be more knowledgeable about the local government, on average, and would therefore be more supportive of local control of new technology in policing. However, our results, if anything, point in the opposite direction. As depicted in Figure A3, those who watch local news are significantly less likely ( $p < 0.001$ ) to trust sheriff use of new AI tools in policing and are borderline less likely ( $p = 0.086$ ) to be willing to share their data with sheriffs.<sup>10</sup> Local news viewership is also associated with lower support for sheriff AI use and willingness to pay taxes, but these differences are not statistically significant. One potential explanation for these results is that local news viewers are exposed to social and political problems at the local level, such as crime, political scandals, instances of police officer misconduct, and other failures of government, producing more informed skepticism

<sup>10</sup>We control for demographic factors that may be associated with watching local news, including partisanship, gender, race, age, education, income, employment, region, and feelings toward the police.

for local news viewers. To the extent that the local news media act to hold government accountable and cover critical stories, this could explain why local news viewers are relatively more skeptical of sheriff use of AI.

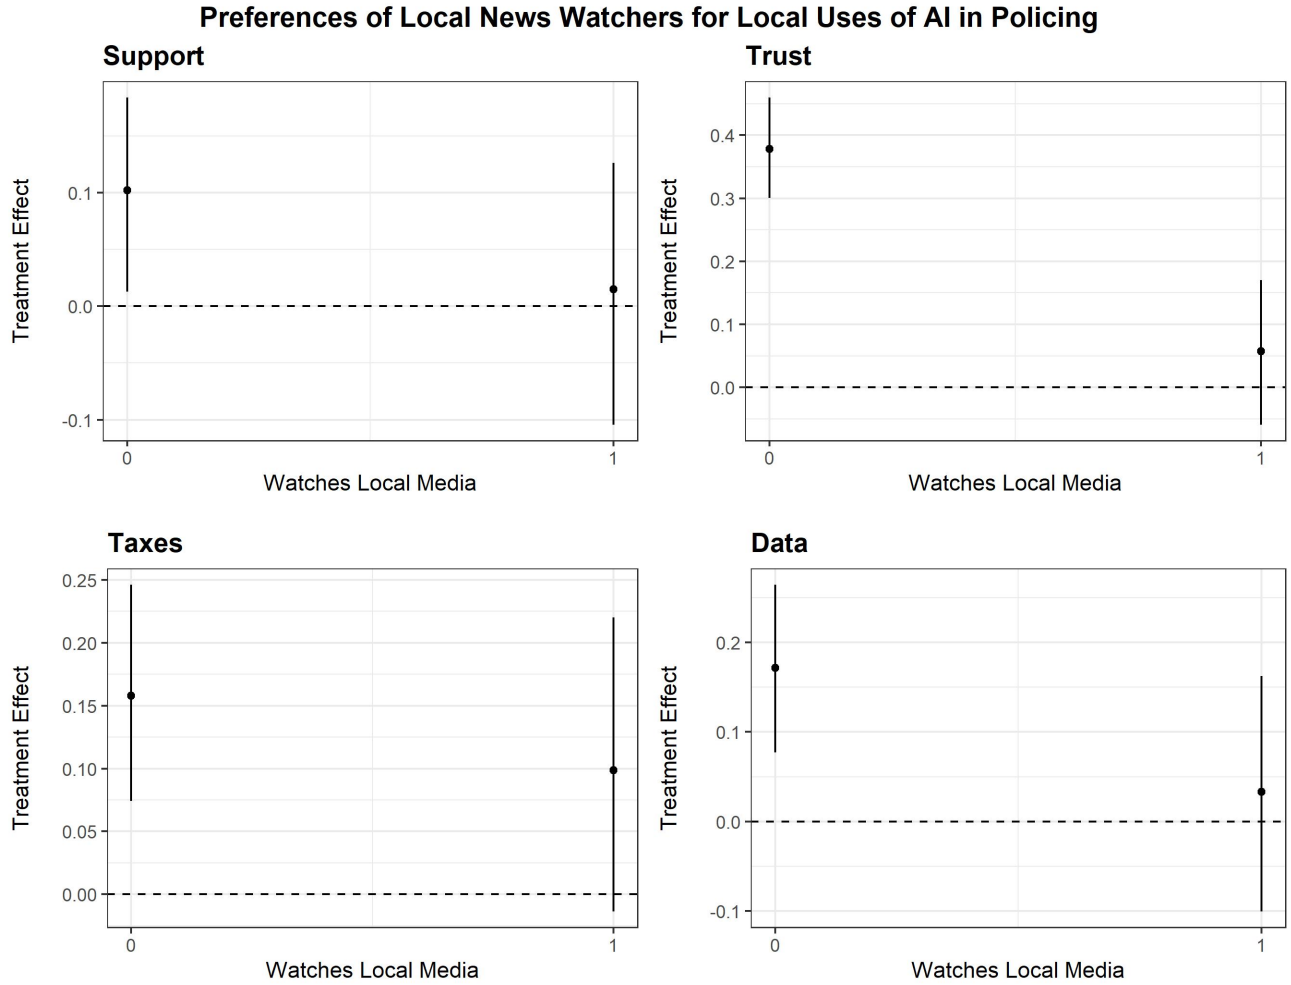


Figure A3: Heterogeneous Effects by Local News Consumption