From Silence to Salience –
How Critical Events Shape Protests

Tim Baule
University of Bayreuth
Faculty of Law, Business and Economics
tim.baule@uni-bayreuth.de

Henrik Guhling (corresponding author)
University of Bayreuth
Faculty of Law, Business and Economics
henrik.guhling@uni-bayreuth.de

This version: August 12, 2026

Abstract: We propose a novel attention-based mechanism for protest participation, based on salience theory (Bordalo *et al.*, 2013). While existing models emphasize thresholds, networks, or informational dynamics, our model incorporates high-profile events that alter the context of protest decisions, shifting salience towards the benefits of protesting while underweighting associated costs. We challenge our model empirically by examining the killing of Black individuals by police officers as high-profile events and assessing their impact on attention and protest activity. Our findings show that such events significantly increase protest mobilization, with a substantial portion of the effect mediated by shifts in salience. This mechanism offers a behavioral component to existing protest models and expands our understanding of protest dynamics in the aftermath of attentional shocks.

Keywords: Attention; Black Lives Matter; High-Profile Events; Protests; Salience Theory.
JEL classification: D71; D74; D91.

Acknowledgments: The authors thank the participants of the following research conferences: Jahrestagung des Ausschusses für Wirtschaftssysteme und Institutionenökonomik 2025, 29th Annual SIOE Conference 2025, Public Choice Conference 2025, Radein Seminar 2025, Bayreuth Graduate Seminar 2024 and 2025, and 38th BGPE Research Workshop 2025. The authors thankfully acknowledge helpful suggestions and feedback from Martin Leschke, Fabian Herweg, and Thomas Apolte. The authors would like to thank RW Alumni Bayreuth and the Institute for Humane Studies for providing travel support.

1. Introduction

High-profile events that attract substantial media coverage have the potential to trigger instant, far-reaching public uprisings in the form of protests. For instance, the murder of George Floyd on May 25, 2020, provoked large-scale protests against police violence and racial discrimination in the US and beyond. Like other protest movements during the last decade, protest mobilization went hand in hand with a wide online presence of the topic, steering societal attention towards the protest cause. These dynamics raise essential questions about how attention shocks triggered by high-profile events influence protest outcomes related to the affected issue. Exposure to a protest's cause through social media and other platforms alters the context in which individuals make participation decisions. By leading individuals to focus on the benefits of protesting and neglecting the associated personal costs, such exposure may increase mobilization–particularly in contexts marked by heightened public awareness following high-profile events, such as the murder of George Floyd.

Identifying the theoretical and empirical mechanisms through which high-profile events influence protest mobilization remains a significant challenge. From a theoretical perspective, while such events can change the context in which individuals decide whether to protest, modeling this context dependence is particularly challenging when neither the private benefits nor the private costs change in the event's immediate aftermath. From an econometric point of view, such events occur rarely, and thus examining them empirically poses challenges due to the typically small number of available observations. An increase in the number of included events almost certainly introduces either substantial heterogeneity among the considered events or across regional units of observation. Additionally, high-attention events are usually defined ex post, where the events have already undergone a selection process. This introduces a selection bias, as events are often included based on retrospective influence on protest outcomes rather than ex ante criteria.

To address the theoretical challenge, we develop an experience-based salience model that incorporates the attentional context into protest decisions, building on Bordalo *et al.* (2013, 2020a). Individuals decide whether to protest based on two attributes: the benefits (e.g., expressive utility, network effects) and the costs (e.g., opportunity costs, risk of violence). Over time, these individuals encounter experiences related to the protest or its underlying issue through channels such as social media, news outlets, or personal conversations. These experiences vary in how strongly they shape perceptions of the protest, with specific contexts amplifying their influence. For instance, extensive media coverage can make an incident significantly more impactful than a similar, less-publicized event. We refer to abrupt changes in this attentional environment as context shocks, which can shift attention toward the benefits of protesting and increase participation even without changes in the objective benefits or costs.

The proposed model connects high-profile events to protest mobilization by modeling an attention mechanism through salience. In doing so, it highlights the importance of contextual and attentional features in individual decision-making. We claim that the mechanism has substantial explanatory power for rapid protest mobilization patterns after high-profile events. Several instances clearly illustrate the mobilization dynamics of interest. Such events include, for example, the murder of George Floyd, which led to sustained anti-police violence protests; the Chinese extradition plans in 2019, which led to mass protests in Hong Kong; and the self-immolation of Mohamed Bouazizi in 2010, which started the Arab Spring.[1]

To overcome the aforementioned econometric difficulties, we shift the analysis to the local level in the US and define fatal police encounters with Black individuals as local context shocks. Especially in the context following George Floyd's death, such events received increased media attention and thus changed the context in which individuals make their protest decisions. The investigation at the local level allows for assessing the effects of context shocks on the salience level and protest outcome as a direct comparison to units sufficiently distant but also sufficiently similar. Hence, we control for nationwide trends by comparing units with context shocks (treatment group) vis-à-vis units without context shocks (control group). We thereby circumvent the econometric difficulties of investigating nationwide context shocks by (a) defining events ex ante and (b) having suitable control units.

We test the model's empirical predictions at the US county level for 2022, extending the analysis period from 2017 through 2022 for specific specifications. From an identification standpoint, the primary empirical challenge lies in establishing the plausibility of the conditional independence assumption. We address this potential endogeneity by implementing a data-driven double machine learning framework via LASSO covariate selection. This procedure operates through a three-step sequence: (1) we aggregate a

[1] Several excellent studies investigate protest dynamics in these contexts: among others, Moreno-Medina *et al.* (2025) investigate biased media attention in police violence reporting, Williamson *et al.* (2018) assess the effect of police violence on protest outcomes in the context of BLM protests. Cantoni *et al.* (2019) argue that protests are strategic substitutes in the Hong Kong anti-authoritarian movement, and Cantoni *et al.* (2022) investigate which personality traits make individuals more likely to protest. In the Arab Spring literature, there are studies about its economic causes (Malik and Awadallah, 2013) and the role of social media (Khondker, 2011).

high-dimensional conditioning space comprising a rich set of county characteristics that could plausibly confound policing behavior and protest activity; (2) we estimate separate penalized regressions to project both treatment assignment (shootings) and the outcome (protests) onto the covariate space, identifying variables with predictive power in either equation; and (3) we include the union of all variables retaining non-zero coefficients in either first-stage step into the final causal estimation. By forcing the inclusion of variables that predict treatment status, the approach minimizes the risk of omitted variable bias and strengthens the robustness of our cross-sectional estimates. In addition, we also perform Oster bound calculations to quantify the potential influence of unobserved confounders (Oster, 2019). Across the tested specifications, the bounds exceed the conventional threshold of $\delta = 1$: any unobserved confounder would have to be more important than the entire set of LASSO-selected controls to explain away the estimated effects, making unobserved confounding an unlikely driver of our results.

To support the validity of this unconfoundedness assumption, our covariate set explicitly absorbs the long-run structural, socioeconomic, and institutional determinants of county-level variance. Specifically, the baseline vector controls, among other variables, for political leanings and local governance metrics (e.g., county-level Democratic support and political fractionalization scores), alongside baseline socioeconomic strains captured via county Gini coefficients, median household incomes, and racial income gaps. We also incorporate granular demographic profiles (including population density, age-cohort distributions, and racial shares) and explicitly account for localized civic and social capital dynamics by integrating measures of charitable organization density and collective efficacy. Conditional on these extensive controls, we apply a suite of treatment evaluation methods–specifically regression adjustment, propensity score methods, and matching estimators–to isolate the average treatment effects of these context shocks on salience levels and protest outcomes. Potential violations of the Stable Unit Treatment Value Assumption (SUTVA) due to regional spatial spillovers are assessed using distance-based sample trimming following Clarke (2017). We utilize Google Trends search data as an active behavioral proxy for explicit issue salience, and implement a distributed lag model to capture the temporal decay and dynamic predictions of these attention shocks.

The empirical evidence provides robust support for the major propositions derived from the model. Fatal police encounters involving Black individuals, interpreted as context shocks, are found to increase protest mobilization for Black Lives Matter (BLM) and anti-police violence demonstrations compared to a control group. The majority of this effect is attributable to the proposed salience mechanism. Furthermore, the mechanism is identifiable for the years 2018, 2020, and 2022, but differs substantially in size. After the George Floyd incident, the effects increased substantially, indicating that the nationwide context shock interacts with the local context shocks. In terms of sustained mobilization over the years, the salience mechanism does not seem to be the main driver. Instead, the salience mechanism appears to have substantial explanatory power for short-run mobilization processes, while the persistence of protest activity over time likely hinges on alternative mechanisms beyond attentional dynamics.

This paper makes several contributions to the literature. First, it introduces a novel attention-based channel that helps explain the emergence of protests. Second, it formalizes the role of high-profile events in decision-making by defining context shocks. Third, examining mobilization events at the local level allows for an assessment of their impact on protests by establishing an appropriate control group. This also makes it possible to classify such events by their magnitude and influence on protest activity. Finally, the paper offers an empirical approach to measuring context and salience by interpreting media coverage of a police shooting as context and Google Trends data as salience.

The paper's structure is as follows. After discussing the related literature in the remainder of this section, Section 2 presents the proposed model and testable hypotheses resulting from the model. In Section 3, we introduce our empirical strategy to assess the model's predictions. Section 4 presents the results and assesses their significance for the proposed model and the protest literature.

Related Literature The paper contributes to different strands of the literature. First, we introduce a novel attention mechanism to the theoretical literature, which typically attempts to explain protests by relying on models that focus on thresholds, networks, or information. The threshold model of Granovetter (1978) describes how an individual's decision to engage in collective action depends on the number or proportion of others already participating. Yin (1998) applies this framework to the domain of protests, showing that individuals may abstain from participation until a critical point is reached, after which joining becomes the rational choice. Such tipping points can potentially be reached by a small minority of "certain actors," which leads to widespread mobilization in the aftermath, as modeled by Wiedermann *et al.* (2020). While threshold models explain participation as a response to the number of others already protesting, our proposed attention mechanism focuses on why an individual decides to protest in the first place.

Within the literature on threshold models, our proposed attention mechanism contributes to understanding the transition between different equilibria by explicitly modeling how they switch. Threshold models often rely on an "exogenous shock" to explain why a first group of individuals begins to protest, thereby triggering a cascade that leads to widespread participation. However, this framing leaves the nature of the initial activation mechanism underspecified. Our proposed attention-based mechanism offers a formal explanation for how individuals may choose to protest even when their personal cost-benefit calculus remains unchanged – what shifts is the context in which they make their participation decision. In doing so, we replace the vague notion of an exogenous shock with a concrete, endogenous mechanism that accounts for the initial mobilization.

Another related strand of literature emphasizes the strategic interdependence of protest decisions. Shadmehr (2018) develops a model in which individuals decide whether to protest based on uncertain payoffs and the anticipated actions of others. They show that protest participation can exhibit either strategic substitutability or complementarity: when protest goals are modest, individuals act as strategic substitutes, while ambitious goals foster complementarity. Empirical evidence from Hong Kong supports the substitutability channel: Cantoni *et al.* (2019) find that students reduced their likelihood of participating in protests after learning that turnout would be higher than they initially expected. In contrast, in settings characterized by strategic complementarity, coordination becomes central. De Mesquita (2010) argues that public signals can shift beliefs about the broader anti-regime sentiment, thereby facilitating collective action. In contrast, our approach focuses on the decision of an individual independent of others' protest behavior, without relying on the protesting decisions of others, with heightened attention due to context shocks serving as the potential trigger.

Information-based models typically start with the assumption that individuals have imperfect information about government quality and/or the attitudes of fellow citizens. This category of models focuses on how information diffuses throughout society and how this mechanism affects protest mobilization. For example, Lohmann (1994, 2000) models protesting as a decentralized information aggregation mechanism and applies the theory to anti-regime mobilization dynamics in former East Germany. Little (2016) builds on that research by modeling how social media lowers the costs of communicating personal grievances with regimes and sharing information about protest coordination. Personal interactions are modeled in Barbera and Jackson (2019) in which individuals can update their beliefs according to communications with others. Belief updating about the government quality is explicitly modeled in Meirowitz and Tucker (2013) with a focus on the quality distribution of potential governments. Such informational changes are consistent with our proposed attention channel insofar as they alter the context in which individuals make decisions. Our analysis focuses on the attention mechanism, where context shocks influence individuals' perceptions.

Furthermore, this paper contributes to the empirical literature on protest occurrences, specifically to the literature on attention markets like social media platforms. Both González-Bailón *et al.* (2011) and Larson *et al.* (2019) identify a significant role of Twitter in the 2011 mobilization waves in Spain and the 2015 Charlie Hebdo protests in Paris. Enikolopov *et al.* (2020) demonstrate a similar effect of the VK platform, a major social media platform in Russia, on protests and their size. Regarding effect timing, Kann *et al.* (2023) use Twitter data and Black Lives Matter protests to show that social media increases interest in the topic in the short term, but that there is no sustainable increase in collective identity. In our analysis, we focus on the role of social media content in shaping the perceptual context in which individuals decide to protest, as captured by our attention mechanism framework.

Finally, we contribute to the strand of literature that extends the salience model of Bordalo *et al.* (2012, 2013) to topics in political economy. In this respect, Bordalo *et al.* (2020b) built on salience to explain why voters sometimes overweight extreme positions. In addition, Bilgin and Destan (2024) examine the impact of decoy candidates, who are not intended to be elected but can nonetheless influence the election outcome. Bonomi *et al.* (2021) show how the salience of a social divide leads voters to identify more strongly with their economic or cultural group, thereby pulling beliefs more strongly toward the stereotype of that group. Our analysis applies salience theory to the domain of protesting, using field data to measure both context and salience levels and to test hypotheses derived from the model.

2. Salience Model

To model our proposed attention mechanism, we adopt a salience-based framework. According to Bordalo *et al.* (2012, 2013), salience refers to the tendency of individuals to focus on the options, attributes, or outcomes that attract the greatest attention, i.e., the most salient ones. In salience theory, the context is typically defined by the choice set, with all available options forming a reference that shapes the perception of each option relative to the whole.

In contrast, our model focuses on a single choice option, namely protesting, in which perceptions are shaped by past experiences with similar protests or related topics. These past experiences include, among other aspects, information from social media, news outlets, personal conversations, and other sources. Thus, following Bordalo *et al.* (2020a), we model the context as an experience-dependent parameter, where the reference is formed not from all available options, but from the collection of past experiences with a single option, i.e., the protest.

In the following, we consider a decision maker (DM) who decides in each period $t = 0, 1, \ldots$ whether or not to take part in a protest $P = (b, c)$. We model P as a two-attribute good with $(b, c) \in \mathbb{R}_0^+ \times \mathbb{R}_0^+$, where b reflects individual potential benefits (e.g., expressive utility, network effects) and c the individual potential costs (e.g., opportunity costs, risk of violence). The DM's outside option is given by $\underline{P} = (0, 0)$.

At time t, the DM retrieves all past experiences regarding the protest or the protest-related topic and aggregates them into a reference protest $\overline{P}_t = (\bar{b}_t, \bar{c}_t)$. Thus, the perception of the protest P at time t depends on $\overline{P}_t$. Each experience e_s, $s \leq t$, until time t is characterized by

$$e_s = (b_s, c_s, k_s) \in \mathbb{R}_0^+ \times \mathbb{R}_0^+ \times \mathbb{R}_0^+ \tag{1}$$

where b_s and c_s reflect the endogenously perceived benefits or costs of the protest at time s, respectively. Importantly, these perceived attributes need not coincide with the true attributes b and c. k_s denotes the exogenously given context at time s.[2]

Context is broadly defined to encompass the situational environment in which the DM perceives an experience. Such environments may include reading a newspaper, browsing social media, or attending a protest event. The way in which protest-related information is presented within these contexts can influence its perceived importance and relevance. For example, reading about a political development on the front page of a newspaper is likely to shape perceptions differently than encountering the same article in a less prominent section. Similarly, context can reflect the overall media coverage of a protest-related topic: frequent reporting can make the issue appear more important, while the choice of media outlet can frame the topic in a distinct way.

[2] In general, we assume that different situations can be compared with each other. The context k_s captures this comparability by assigning different numbers to different situations. Initially, whether a value is high or low is irrelevant; however, the smaller the numerical distance between two values, the more similar the corresponding situations are assumed to be.

2.1. Similarity and Salience

All experiences until time t are organized in a memory database M_t, which is updated at each point in time. Then, the current experience at time t stimulates recall of similar past experiences.[3] Based on Bordalo *et al.* (2020a), similarity is measured along all three dimensions of experiences.

Definition 1 (Similarity). *The similarity of the experience $e_s = (b_s, c_s, k_s)$, $s \leq t$, to the current experience $e_t = (b_t, c_t, k_t)$ is given by the multiplicatively separable distance:*

$$S(e_s, e_t) = S_1(|b_t - b_s|)S_2(|c_t - c_s|)S_3(|k_t - k_s|) \tag{2}$$

where $S_k : \mathbb{R}_0^+ \to \mathbb{R}_0^+$ is decreasing for $k = 1, 2, 3$.

Thus, $S(e_s, e_t)$ takes the highest value if the experience e_s contains the same parameter values as e_t. For example, $S(e_s, e_t)$ could have an exponential form, i.e.,

$$S(e_s, e_t) = exp(-\big[|b_t - b_s| + |c_t - c_s| + |k_t - k_s|\big]) \tag{3}$$

such that it takes values between 0 and 1.

Building on this concept of similarity, the reference protest is defined as the similarity-weighted average of perceived benefits and costs, where each experience in M_t is compared to the current one e_t, and experiences with greater similarity receive higher weights.[4]

Definition 2 (Memory weights and reference protest). *The memory weight of experience $e_s = (b_s, c_s, k_s)$, $s \leq t$, at time t is given by*

$$w(e_s, e_t) = \frac{S(e_s, e_t)}{\sum_{e_u \in M_t} S(e_u, e_t)}. \tag{4}$$

Then, for the reference protest at time t, $\overline{P}_t = (\bar{b}_t, \bar{c}_t)$, it holds

$$\bar{b}_t = \sum_{e_u \in M_t} b_u \cdot w(e_u, e_t), \quad \bar{c}_t = \sum_{e_u \in M_t} c_u \cdot w(e_u, e_t). \tag{5}$$

[3] Recall mechanisms are usually described by a cue that contains one or more dimensions of an experience and thereby evokes memories (Bordalo *et al.*, 2020a). However, in our case, a context at time t is always associated with perceived b_t and c_t, so that the cue always contains all three dimensions and simply corresponds to the current experience e_t at time t. Therefore, we omit the definition of a cue.

[4] Unlike Bordalo *et al.* (2020a), we adopt a discrete-time framework because, in our model, sudden jumps in perceived benefits should have a substantial effect on the reference in order to capture context shocks. In continuous time and when the Riemann integral exists, such jumps could be measured as zero, as they may only represent single points. Using discrete time avoids this issue.

In our salience model, the protest's attributes b and c are weighted based on discrepancies between their true values and the corresponding perceptions in the reference. Thus, at time t, the protest P is evaluated according to the utility function

$$U_t(b,c) = \frac{2 \cdot \delta^{-\sigma(b,\bar{b}_t)}}{\delta^{-\sigma(b,\bar{b}_t)} + \delta^{-\sigma(c,\bar{c}_t)}} \cdot b - \frac{2 \cdot \delta^{-\sigma(c,\bar{c}_t)}}{\delta^{-\sigma(b,\bar{b}_t)} + \delta^{-\sigma(c,\bar{c}_t)}} \cdot c \tag{6}$$

where $\delta \in (0,1]$ measures the degree of salient thinking. For $\delta = 1$, there is no salience bias such that both attributes receive the same weight. The smaller δ, the stronger the perception bias. The salience function $\sigma(\cdot,\cdot)$ captures the key features of sensory perception (Bordalo *et al.*, 2013) summarized in Definition 3.[5]

Definition 3 (Salience function)**.** *For an attribute $a \in \{b,c\}$, the salience function $\sigma(a,\overline{a})$ is a symmetric, continuous, and non-negative function that satisfies the following properties:*

1. ***Ordering**: Let $\mu = sign(a - \overline{a})$. Then, for any $\epsilon, \epsilon' \geq 0$ with $\epsilon + \epsilon' > 0$ it holds*

$$\sigma(a + \mu\epsilon, \overline{a} - \mu\epsilon') > \sigma(a, \overline{a}) \tag{7}$$

2. ***Diminishing sensitivity**: For any $a, \overline{a} \geq 0$ and all $\epsilon > 0$, it holds*

$$\sigma(a + \epsilon, \overline{a} + \epsilon) < \sigma(a, \overline{a}) \tag{8}$$

2.2. Context Shocks and Reference Protest

In our attention mechanism, we propose that high-profile events, such as the George Floyd incident, can alter the context substantially, acting as a context shock. This shock creates an experience with extreme k_t and b_t, which in turn shifts the reference protest. This shift changes salience, potentially directing more attention toward the protest's benefit b and prompting the decision-maker to participate in the protest.

To incorporate this mechanism in our model, we assume dependencies between the context k_t and the perceived benefit b_t. Furthermore, for tractability, we fix the cost perception bias.

Assumption 1. *We assume the following two conditions:*

[5] We adapt the rank-based salience approach of Bordalo *et al.* (2013) by introducing continuous attribute weights. This allows us to assess not only which attribute is more salient, but also the degree to which one exceeds the other in salience. Setting one attribute's salience to 0 and the other's to 1 recovers the original rank-based model.

(i) The perceived benefit at time t, b_t, depends on the context k_t, such that $b_t = f(k_t)$ with f continuous and injective.

(ii) The costs c and $\Delta_c = |c_t - c|$ are assumed to be constant over time, such that there is a constant cost perception bias.

The first assumption specifies a functional form for how the decision-maker's perception of benefits responds to changes in context, with larger changes in context leading to greater changes in perceived benefits.

The second assumption fixes the salience of costs, allowing all salience relationships between benefits and costs to be captured solely through changes in the salience level of b. This simplification involves no loss of generality, although perceived costs could, in principle, also be made context-dependent.

Next, we introduce a general definition of context shocks. Mathematically, a context shock occurs when the context reaches a level not previously achieved in a specific period of time.

Definition 4 (Context shock)**.** *A context at time t, k_t, is called a* ***context shock on*** *$\{s, ..., t\}$, $s < t$, if and only if the following property is fulfilled:*

$$\max_{u,w \in \{s,...,t-1\}} |k_u - k_w| < \max_{u,w \in \{s,...,t\}} |k_u - k_w| \tag{9}$$

If $s = 0$, then k_t is called a ***global context shock****.*

That is, we observe a context shock in t if there is a time period $\{s, ..., t\}$ such that the maximum fluctuation of the context becomes larger by adding the time t. In the simple case of a constant context, every change represents a context shock. Importantly, it does not matter if k_t gets bigger or smaller. It is only vital that it changes.

Based on our mathematical definition, every context change can represent a context shock. However, this does not imply that every context shock also affects an individual's perception of protests. Only shocks that entail a change in the reference benefit $\bar{b}$ can change perceptions by influencing salience. The following lemma derives sufficient criteria for the relationship between context shocks and the reference benefit.[6]

Lemma 1 (Context shocks and reference benefit)**.** *Consider discrete time with $s \in \{0, ..., t\}$. Then a context shock at time t, k_t, induces a higher (lower) reference benefit, i.e., $\bar{b}_t >$ $(<)$*

[6]A proof of Lemma 1 can be found in Appendix A.

$\bar{b}_{t-1}$, *if and only if the following inequality is fulfilled:*

$$\sum_{s=0}^{t} S(e_s, e_t) \cdot [b_s - \bar{b}_{t-1}] > (<)\ 0 \tag{10}$$

A global context shock k_t *inducing* $b_t = max_{s \leq t} b_s$ ($b_t = min_{s \leq t} b_s$) *leads to a higher (lower) reference benefit if one of the following conditions is fulfilled:*

(i) It holds $\bar{b}(k_{t-1}) \leq (\geq)\ \frac{1}{t} \sum_{s=0}^{t} b_s$.

(ii) For $\alpha_s = \frac{S(e_s, e_t)}{S(e_s, e_{t-1})}$, *it holds* $\alpha_i \geq \alpha_j \Leftrightarrow b_i \geq (\leq)\ b_j, \quad \forall i, j \in \{0, ..., t-1\}$.

Condition 10 is always fulfilled for a constant context before time t. Here, a global context shock always results in a higher (lower) reference benefit.[7] However, the reference benefit, and thus salience, remains unchanged if Condition 10 holds with equality. In the following, we define salience-changing context shocks as salience shifts.

Definition 5 (Salience shift)**.** *A context shock on* $\{s, ..., t\}$ *is called*

(i) ***positive salience shift*** *if and only if*

$$\sigma(b, \bar{b}_t) > \max_{u \in \{s, ..., t-1\}} \sigma(b, \bar{b}_u). \tag{11}$$

(ii) ***negative salience shift*** *if and only if*

$$\sigma(b, \bar{b}_t) < \min_{u \in \{s, ..., t-1\}} \sigma(b, \bar{b}_u). \tag{12}$$

A change in context that leads to the benefits' salience reaching a new level is referred to as a salience shift. When the advantages of protests are brought into greater focus, resulting in the benefit becoming more salient, this constitutes a positive salience shift. Conversely, if a context shock causes the salience of the benefits to diminish, thereby indirectly highlighting the disadvantages or costs of protests, this is termed a negative salience shift.

A salience shift can only occur through changes in the reference benefit of the form specified in Lemma 1. If the reference benefit already exceeds the true benefit b, the ordering property implies that even a slight upward change suffices to trigger a positive salience shift. If it lies below b, a slight downward change is sufficient. Theoretically,

[7] Please note that the example function $S(e_s, e_t) = exp(-[|b_t - b_s| + |c_t - c_s| + |k_t - k_s|])$ from Equation 3 always fulfills Condition (ii) of Lemma 1 such that a higher b_t leads to an increased $\bar{b}_t$.

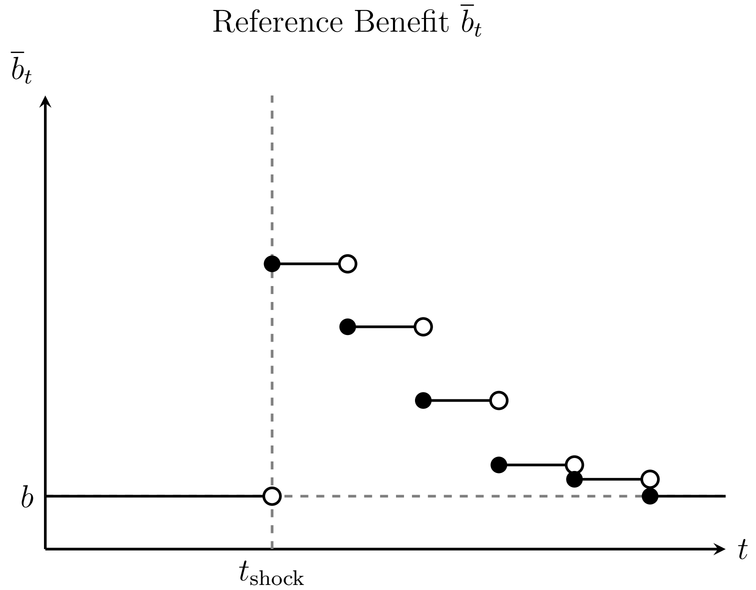

Figure 1: Shift in reference benefit $\bar{b}_t$ due to context shock in t_{shock} for $b_t = k_t$ and $S(e_s, e_t) = exp(-|b_t - b_s|)$.

sufficiently large changes that move the reference benefit from above b to below it, or vice versa, can also induce a positive salience shift. The reverse logic holds for a negative salience shift.

Figure 1 illustrates how a context shock can influence the reference protest and, consequently, the perception of the benefit. With context and perceived benefit held constant, any change constitutes a context shock. According to Lemma 1, an upwards change of b_t in t_{shock} raises the reference benefit to a new, higher level. With $b_t = b$ for $t < t_{shock}$, this context shock heightens the benefits' salience, producing a positive salience shift. In the example shown, the context gradually returns to its original level, causing b_t and thus the reference benefit to normalize to its prior level.

2.3. Protest Participation

Applying the model to the decision problem, a DM of type δ protests at time t if and only if

$$U_t(b, c) \geq U_t(0, 0) = 0. \tag{13}$$

Rearranging that leads to

$$\delta^{\sigma(c,\bar{c}_t)-\sigma(b,\bar{b}_t)} \geq \frac{c}{b}. \tag{14}$$

If both attributes are equally salient, Condition 14 shows that for $c > b$, nobody protests irrespective of type δ, whereas everybody protests for $b \geq c$. In this case, the protest decision only depends on the true attributes b and c.

In general, the participation of individuals in our model depends on which attribute is more salient. The more salient the benefit is compared to the costs, the more types of δ should participate. By rearranging Condition 14, the following proposition gives an overview of protesting individuals based on the attributes' salience.

Proposition 1. *If the participation costs exceed the benefits, i.e., $c > b$, an individual of type $\delta \in (0,1]$ decides to protest at time t if and only if $\delta \leq \delta_t^*$ where*

$$\delta_t^* = \begin{cases} 0 & \text{if } \sigma(b,\bar{b}_t) \leq \sigma(c,\bar{c}_t), \\ \left(\frac{b}{c}\right)^{\frac{1}{\sigma(b,\bar{b}_t)-\sigma(c,\bar{c}_t)}} \in (0,1) & \text{if } \sigma(b,\bar{b}_t) > \sigma(c,\bar{c}_t). \end{cases} \tag{15}$$

If the benefits exceed the participation costs, i.e., $b \geq c$, an individual of type $\delta \in (0,1]$ decides to protest at time t if and only if $\delta \geq \delta_t^*$ where

$$\delta_t^* = \begin{cases} 0 & \text{if } \sigma(b,\bar{b}_t) \geq \sigma(c,\bar{c}_t), \\ \left(\frac{c}{b}\right)^{\frac{1}{\sigma(c,\bar{c}_t)-\sigma(b,\bar{b}_t)}} \in (0,1) & \text{if } \sigma(b,\bar{b}_t) < \sigma(c,\bar{c}_t). \end{cases} \tag{16}$$

The proposition provides a threshold value δ_t^*, which depends on the attributes' salience at time t and pinpoints the critical type, i.e., the type indifferent between protesting and not. Two trivial extreme cases arise by considering δ_t^*. First, if costs exceed benefits and are more salient, no one will protest. This makes sense, as attention is focused on the higher costs. Second, if benefits exceed costs and are more salient, all individuals will protest regardless of type δ. This also makes sense, since attention is on the higher benefits. In the remaining two cases, only a proportion of individuals protest, determined by their type δ.

If costs exceed benefits but the benefits are more salient, all individuals whose salient thinking exceeds a certain minimum, i.e., who focus sufficiently on the more salient attribute, will protest. As the salience of the benefit increases, the threshold value δ_t^* rises, leading to more types of δ participating.

Conversely, if benefits exceed costs but the costs are more salient, all individuals whose salient thinking is below a certain level, i.e., who do not focus too heavily on the more salient attribute, will protest. The smaller the salience gap between benefits and costs, the lower the threshold δ_t^* and the higher the participation rate.

Next, we examine how context shocks that induce a salience shift influence protest participation. The following corollary summarizes the results.

Corollary 1 (Context Shocks and Protest Participation)**.** *Consider discrete time with* $s \in \{0, ..., t\}$*.*

1. *For* $c > b$*, it holds:*
   - *(i) If* $\sigma(b, \overline{b}_{t-1}) - \sigma(c, \overline{c}_{t-1}) \geq 0$*, then every* ***positive salience shift*** *on* $\{s, ..., t\}$ *leads to increased protest participation.*
   - *(ii) If* $\sigma(b, \overline{b}_{t-1}) - \sigma(c, \overline{c}_{t-1}) < 0$*, then only sufficiently large* ***positive salience shifts*** *on* $\{s, ..., t\}$ *lead to increased protest participation.*
   - *(iii) If* $\sigma(b, \overline{b}_{t-1}) - \sigma(c, \overline{c}_{t-1}) > 0$*, then every* ***negative salience shift*** *on* $\{s, ..., t\}$ *leads to decreased protest participation.*
2. *For* $c < b$*, it holds:*
   - *(iv) If* $\sigma(b, \overline{b}_{t-1}) - \sigma(c, \overline{c}_{t-1}) \leq 0$*, then every* ***negative salience shift*** *on* $\{s, ..., t\}$ *leads to decreased protest participation.*
   - *(v) If* $\sigma(b, \overline{b}_{t-1}) - \sigma(c, \overline{c}_{t-1}) > 0$*, then only sufficiently large* ***negative salience shifts*** *on* $\{s, ..., t\}$ *lead to decreased protest participation.*
   - *(vi) If* $\sigma(b, \overline{b}_{t-1}) - \sigma(c, \overline{c}_{t-1}) < 0$*, then every* ***positive salience shift*** *on* $\{s, ..., t\}$ *leads to increased protest participation.*

Parts 1 and 2 of Corollary 1 cover the cases $c > b$ and $b > c$, respectively. If $b = c$, participation changes only when a context shock reverses the weak salience ordering between benefits and costs, i.e., when $\sigma(b, \overline{b}_t) - \sigma(c, \overline{c}_t)$ crosses zero.

The results indicate that not all context shocks that alter salience necessarily affect protest participation. When costs exceed benefits and are also more salient, only a substantial increase in the salience of benefits can spur participation. Presumably, the George Floyd incident can be understood as such an event that reweights the attention towards the benefit as a consequence of a positive salience shift.

Conversely, when benefits exceed costs and are more salient, a sharp decline in benefits' salience is needed for costs to become salient again. Episodes of severe repression in autocratic regimes, such as the Tiananmen massacre, exemplify large negative salience shifts that redirect attention toward costs. However, multiple smaller negative shifts can also erode participation over time. For instance, waning attention to the protest-related topic can cause movements to shrink and eventually vanish, unless a new context shock triggers a positive salience shift that recaptures attention.

Discussion of Protest Movements over Time Finally, we turn to how our model can explain the life cycle of protest movements. Before protests emerge, mobilization is typically low, consistent with $c \geq b$ and costs being more salient; in our framework, this may also occur with $b > c$ when the benefits' salience is sufficiently low. In both situations, positive salience shifts can increase participation, though in the first case, the shift must be sufficiently large. Major salience shifts, such as the George Floyd incident, often serve as trigger events for protest movements. In some cases, several smaller shifts can cumulatively spark mobilization, albeit more gradually. In many cases, however, a single major event can be identified as the catalyst.[8] After such a shock, additional smaller shocks can further boost participation.

Over time, attention typically dissipates or shifts to other issues, reducing the salience level. As salience falls, participation declines, eventually returning to its initial low level, leaving only those individuals who were consistently motivated. Strong negative salience shifts, such as repression, can accelerate this decline and bring the protest "wave" to an earlier end.

In principle, the ratio of b to c can change from $c > b$ to $b > c$ and back, for example, through increasing network effects and a growing protest movement. Yet even in this case, a decline in the salience of benefits reduces participation, leaving only a small core of participants. However, the case starting with $c > b$ is particularly interesting: if benefits are more salient and the ratio changes to $b > c$, a subsequent drop in benefits' salience means that the movement's initiators, i.e., those with extremely salient thinking, are also the first to withdraw. This pattern is observable in movements that radicalize over time, as initiators disengage when the protest shifts in a direction they no longer support.

Figure 2 presents a numerical example of protest participation for different distributions of δ. In this example, a positive salience shift produces an immediate jump in participation, but over time, negative salience effects dominate, leading to a decline. For $\delta \sim \mathcal{U}(0, 1]$, the participation share matches the critical type δ_t^*. Under $\delta \sim Beta(\alpha, \beta)$, there is less probability mass on the tails of the distribution, accelerating the decline. In the left-skewed case ($Beta(3, 7)$), the initial surge in participation following the context shock is much larger, while in the right-skewed case ($Beta(7, 3)$), the effect is smaller. With a symmetric $Beta(7, 7)$ distribution, the initial increase is larger than in the uniform

[8] Examples include the George Floyd incident; the self-immolation of Mohamed Bouazizi, which heralded the Arab Spring; the shooting of Neda Agha-Soltan, which ignited the Green Movement in Iran; and the death of Mahsa Amini, which triggered the Women's Rights Movement in Iran.

case, but the subsequent decline is also faster.[9]

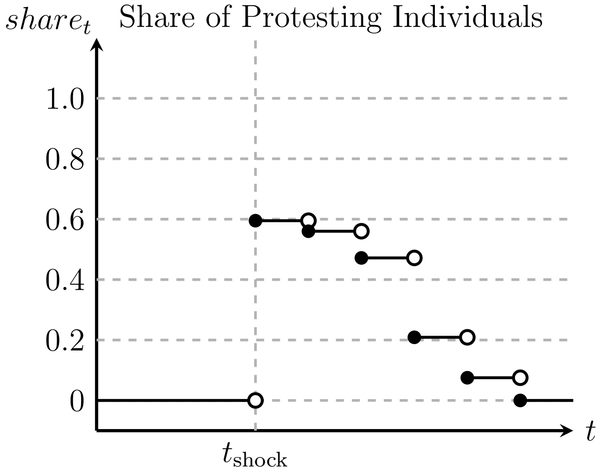

(a) $\delta \sim \mathcal{U}(0,1]$

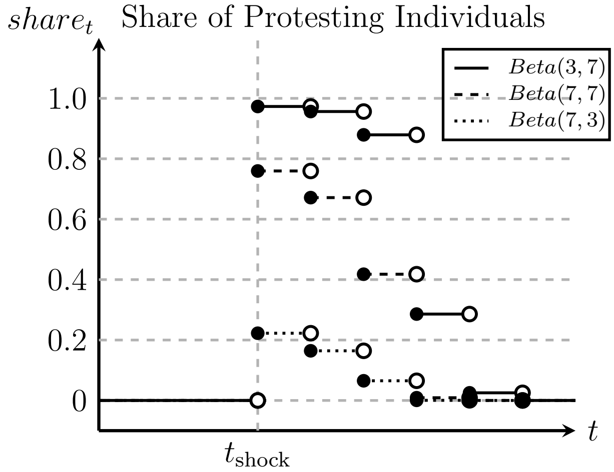

(b) $\delta \sim Beta(\alpha, \beta)$

Figure 2: Share of protesting individuals for $b = 0.7$, $c = 1$, and $\sigma(b, \bar{b}) = \frac{|b-\bar{b}|}{|b|+|\bar{b}|}$.

2.4. Hypotheses

Based on our theoretical model, we derive empirically testable hypotheses. For the empirical analysis, we treat fatal police encounters with Black individuals as context shocks and focus on mobilization patterns in BLM and anti-police protests between 2018 and 2022. For simplicity, we refer to the salience of the protest's benefit as the salience of the protest, as in our theoretical model, variation in protest behavior arises solely from changes in the salience of the benefit, with higher salience directly increasing the likelihood of participation. Furthermore, due to limited data quality in terms of the number of protest participants, we rely on the number of protest events at the US county level as a dependent variable.[10]

Our first hypothesis reflects the core idea that certain acts or events can serve as context shocks, thereby influencing protest outcomes.

Hypothesis 1 (Context shock effect)**.** *We expect context shocks to impact protest outcomes, i.e., police shootings significantly increase the number of anti-police and Black Lives Matter protest events.*

[9] The corresponding *Beta* densities are shown in Appendix B, Figure 6.

[10] Although the used database (Chenoweth and Pressman, 2017) reports estimates of protest participation, these values vary widely between low and high estimates and are missing for a substantial share of events. Therefore, we rely on the number of protest events as a more consistent and informative measure for protest mobilization.

The second hypothesis aims to test the role of salience in shaping protests. In our model, a context shock constitutes a positive salience shift if it raises salience, which in turn increases the number of protest events.

Hypothesis 2 (Salience mechanism)**.** *We expect that context shocks impact salience, i.e., the salience of protest movements is shifted significantly upwards by police shootings (Mechanism 1). Furthermore, this increased salience leads to more protest events, i.e., the number of protest events increases significantly by a shifted salience (Mechanism 2).*

The third hypothesis tests how differently sized context shocks influence protests. If larger shocks lead to more salience, they should also produce a stronger increase in protest events. Since we classify shocks by the number of Google search results for a given shooting, greater media coverage should correspond to higher salience and, consequently, more protests.

Hypothesis 3 (Context shock size)**.** *If bigger context shocks lead to bigger positive salience shifts, we expect to observe more protest events when the context is shocked more drastically.*

The final hypothesis examines the dynamic implications of our model. If the context returns to its pre-shock level, the impact of a context shock should fade as salience also normalizes.

Hypothesis 4 (Dynamic effects)**.** *If context normalizes over time, we expect to observe a decreasing effect of context shocks on salience and the number of protest events.*

3. Empirical Strategy

The following section describes the data collection procedure and empirical strategy to test the theory-derived hypotheses. In the first step, the most important variables for context shocks, salience, and protest outcomes are collected and introduced. Then, three empirical estimation strategies are presented: (1) binary treatment evaluation, (2) continuous treatment evaluation, and (3) distributed lag model (DLM). These estimation strategies are applied in different parts of the results section to test the hypotheses. The final subsection addresses challenges in the empirical strategy and presents sensitivity tests to demonstrate the robustness of the estimated effects.

3.1. Measuring Context Shocks, Salience, and Protests

The three key variables to empirically test the model's predictions are fatal police encounters as context shocks, salience, and protest outcomes. The question of interest is whether there is evidence in the data that a context shock impacts the salience of protesting, which in turn is hypothesized to impact protest mobilization. As protest coordination necessarily involves social groups, the US county level is selected as the observational unit. For the cross-sectional analysis, we choose 2022 as the base year to maintain some distance from the atypical year 2020, marked by nationwide protests following George Floyd's murder. Throughout the paper, we also apply estimation techniques to exploit the time variation in the variables in the period from 2018 to 2022.

Context shocks are measured as fatal encounters of Black individuals with the police (from here on referred to as *fatal police encounters*). The *Mapping Police Violence* database provides detailed data on all fatal encounters of the police with civilians. The data are filtered for fatal encounters with black civilians only and subsequently aggregated on the US county level.[11] These events constitute (micro-level) context shocks, which are supposed to facilitate protest coordination. Having the database in place, the racial disparities in exposure to police violence become visible. While Black individuals' population share is 12.6%, their share in fatal encounters with the police is 23.5%. The states with the highest per capita police killings are North Carolina (0.25 per 1,000 residents), New York (0.23 per 1,000 residents), Nevada (0.20 per 1,000 residents), and New Mexico (0.10 per 1,000 residents).

To capture the extent to which a fatal police encounter led to a context shock, we rely on the Google SerpAPI, which enables us to measure context shocks also on a continuous scale. Based on this, we determine for each police shooting the number of Google results that contain the exact name of the victim and at least one of the keywords "police shooting," "police killing," "fatal police encounter," or "officer-involved shooting." In addition, only results that were published after the date of the victim's death are taken into account. This measure of media coverage about the fatal police encounters serves two purposes. First, it is a more precise measure of how much the context was shocked in a given county. Second, it aggregates the number of results when there was more than one fatal police encounter in a county in a given year (for 2022, for example, there were

[11] In 2022, there were 53 counties in which more than one fatal police encounter happened. In the first step, we do not account for this heterogeneity and instead classify all counties with at least one fatal police encounter as treated counties. Later, we specify media coverage and aggregate a new variable at the county level, which accounts for heterogeneities between treatment intensities.

53 counties in which more than one fatal police encounter happened).

We measure the attentional response to context shocks as salience, proxied by Google Trends data. These data reflect the relative frequency of searches for the term "police shooting." The intuition behind this measure is that the more often the keywords "police shooting" are googled, the more attention people pay to police brutality and problematic race-based discrimination. Importantly, we measure the extent of a context shock by the media coverage of the event, whereas salience is captured by individuals' active search behavior, which we interpret as an indicator of attention. For a discussion on measuring issue salience with Google Trends data, see Mellon (2013). In other scientific contexts, Google Trends data is frequently used to obtain measures of economic uncertainty (Castelnuovo and Tran, 2017; Choi and Varian, 2012). Using Google Trends, we obtain the relative frequency of how often the term was googled for a yearly cross-section on a regional level. The region with the highest Google searches receives a value of 100, and all other values are relative to this benchmark. The regional units of the Google Trends dataset contain multiple counties, and their values are assigned to the lower-level county units. For robustness, all estimations are also conducted on the higher-level Google Trends regional units. The results are generally robust to either specification.

Obtaining a valid measure for salience over time is more elaborate because the Google API sets the maximum value of each region automatically to 100 and then displays all other values relative to the maximum. Cross-county comparisons are, therefore, infeasible because each county, at some point, has a maximum value of 100. We follow Stephens-Davidowitz (2014) to overcome this problem by using reference words that are as uncorrelated to time and geography as possible and measure the salience relative to these reference words. Importantly, the reference words set the maximum value, and the relative frequency of the salience measure is thus comparable over time and across all counties. The selected reference words include "recipies" (misspelled on purpose), "Restroom", "German", and "Back Pain". Figure 3 illustrates how salient the term "Police shooting" was throughout the period of investigation.

Protest data are retrieved from the *Crowd Counting Consortium* database (Chenoweth and Pressman, 2017) – a dataset based on news data. Within the covered years from 2017 to 2022, a natural language processing algorithm retrieves information about protest events from a large number of media outlets. The result is an event database with geocoded protest events across the US. The protest events are aggregated on the US county level and filtered for protest events whose motivations are classified as "race re-

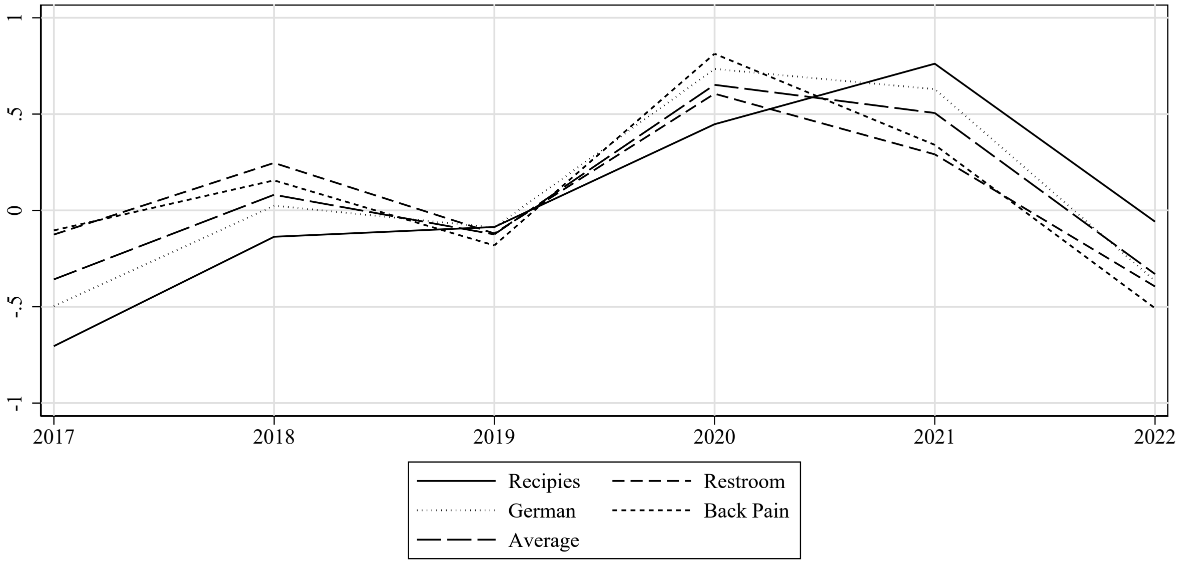

Figure 3: Salience of “Police Shooting” over time across all counties. Note that “Recipies” is misspelled on purpose.

lations” or “against police.” These two motivations coincide with a large share of protest events, and as soon as one category applies, the event is included in the aggregation.

Table 1 describes the main variables of interest. There was a fatal police encounter in 5.9% of US counties in 2022. The average salience of “Police Shooting” was 40. This is measured relative to the maximum of 100 in the county with the highest salience level. Furthermore, in 11% of the counties, there was at least one protest event about race relations or police violence. The remaining empirical section describes techniques by which we aim to establish a causal link from the context shock (fatal police encounter) to protest events via salience.

Key Variables	N	Mean	Std. Dev.	Min	Pctl. 25	Pctl. 75	Max
Fatal Police Encounters	3,222	0.059	0.24	0	0	0	1
Salience “Police Shooting”	3,045	40	9.9	18	34	44	100
BLM Protest Events	3,222	1.4	18	0	0	0	573
BLM Protest Events (binary)	3,222	0.11	0.31	0	0	0	1

Table 1: Summary Statistics for 2022. “BLM Protest Events” include all protest events to which the category “against police” or “race relations” applies. As explained in the text, fatal police encounters include only incidents with black victims.

3.2. Estimation Strategy

We employ a set of cross-sectional estimators to test the implications of Hypotheses 1-3. As a first approach, we assess the effect of context shocks on the number of BLM and anti-police protests on the US county level (Hypothesis 1). Then, we test the mediation channel via salience by splitting the channel into two mechanisms, assessing first the effect of context shocks on salience, and then of salience on protest outcomes (Hypothesis 2). The third specification evaluates heterogeneities towards the size of context shocks by relying on the number of Google results (Hypothesis 3). Finally, the dynamic effects of the salience mechanism are assessed (Hypothesis 4).

I. Binary Treatment

The target parameter for estimation is the average treatment effect on the treated (ATET). This parameter gives an indication of how much a context shock has increased the salience or protest outcome in counties in which a fatal police encounter occurred. The average treatment effect is of less interest as there are some counties with certain characteristics (e.g., almost exclusively white population) that make fatal police encounters (of Black individuals) so unlikely that it is questionable to extrapolate from the treated observations to these counties.[12] In line with Abadie and Cattaneo (2018), Equation 17 presents the target parameter where Y denotes the dependent variable, X denotes the covariates, and W denotes the binary treatment variable.

$$\tau_{\text{ATET}} = E[E[Y \mid X, W=1] - E[Y \mid X, W=0] \mid W=1] \tag{17}$$

The identifying assumptions underlying the estimators are conditional unconfoundedness and common support (Abadie and Cattaneo, 2018). We consider conditional unconfoundedness credible in our setting for three reasons. First, rather than relying solely on a researcher-chosen set of controls, we select covariates through the data-driven LASSO double-selection procedure of Belloni *et al.* (2014). From 76 candidate variables covering demographic composition, income and inequality, labor market conditions, educational attainment, political preferences and fractionalization, immigration history, social capital, civic infrastructure, and population density, the procedure retains variables with predictive power for either treatment assignment or the outcome. By forcing the inclusion of all variables that predict where fatal police encounters occur, the double selection reduces the risk of omitted-variable bias that would arise if treated counties differed systematically

[12] This fact becomes visible in Figure 7. There, a large number of untreated counties have a treatment probability of close to zero.

from control counties along dimensions correlated with protest potential. Second, the candidate set includes the structural determinants of both policing outcomes and protest capacity emphasized in the literature, such as racial composition, economic deprivation and inequality, partisan leaning, and collective efficacy, so that remaining unobserved confounders would have to capture dimensions not adequately represented in this conditioning space to threaten identification. Third, the resulting specifications achieve covariate balance across estimators, and the overidentification test of Imai and Ratkovic (2014) provides no evidence against the imposed balancing conditions, supporting the comparability of treatment and control groups conditional on the selected covariates. All relevant details about the LASSO selection procedure are presented in the Appendix.[13]

LASSO double selection is employed as a first line of defense to make the conditional unconfoundedness assumption plausible. However, because this assumption remains fundamentally untestable, we assess the sensitivity of our estimates to selection on unobservables using Oster bounds (Oster, 2019). The approach exploits the joint movement of the treatment coefficient and the explained variance when moving from an uncontrolled to the fully controlled specification, thereby considering a hypothetical confounder whose relationship to treatment and outcome is proportional to that of the observed covariates. The resulting statistic δ measures how strong selection on unobservables would need to be, relative to selection on the observed controls, to reduce the treatment effect to zero, with $\delta \geq 1$ serving as the conventional robustness threshold. We conduct this exercise for the IPW estimates, as these can be reproduced as weighted least squares regressions and thus map directly into the regression-based framework the procedure requires. Following Oster (2019), we set the maximum attainable explanatory power to $R_{\max} = 1.3\tilde{R}^2$, computed separately for each specification since the attainable explained variance differs with the signal-to-noise ratio of the dependent variable.

Besides the technical assumption that the treatment probability cannot be 1 (Wooldridge, 2010), we assess the overlap of the treatment and control groups to ensure that there are valid units for comparison in the opposite group. The propensity score is estimated as a logit regression of the treatment variable on a set of covariates. The propensity score density function is displayed in Figure 7 in Appendix D.1. Much of the density of the control group is centered at 0, with very low propensity scores. As the target parameter is the ATET, there is no reason to worry because we want to assess whether there are

[13] Appendix C.1 offers an extensive description of the applied procedure. Appendix C.2 presents an exhaustive list of all 76 potential control variables. Finally, Appendix C.3 lists the sets of covariates used, which will be referenced in all subsequent regression tables.

comparison units for the treatment group. These comparison units are given as the control group shows substantial density for medium to high values of the propensity score. In case there remain concerns about extreme propensity scores, we follow Crump *et al.* (2009) in trimming the sample by excluding all observations with propensity scores lower than 0.05 or higher than 0.95 in the sensitivity analysis.

There are various estimators to obtain our target parameter τ_{ATET} as presented in Abadie and Cattaneo (2018). Hence, we apply five estimators to demonstrate our results' robustness: (1) regression adjustment, (2) inverse probability-weighting, (3) inverse probability-weighting regression adjustment (doubly-robust estimation), (4) covariate matching, and (5) propensity score matching. All of these estimators ensure that under the conditional independence assumption, the target parameter is statistically identified. This objective is either achieved by weighting or matching on propensity scores obtained by a prior logit regression or by directly adjusting the comparison unit using a regression adjustment approach. Additionally, we employ a covariate matching estimator that compares units according to how similar they are based on the selected covariates. The included covariates are selected using the previously described LASSO selection procedure. A detailed account of the estimators is delegated to Appendix D.2.

To test Hypothesis 3, we apply the method described in this section and successively trim the sample by excluding observations with Google results below a given threshold.

II. Continuous Treatment

Investigating the effect of salience on protest outcomes requires adapting the estimation strategy for continuous treatments (Mechanism 2 of Hypothesis 2). For this, we rely on Imbens (2000) and Hirano and Imbens (2004). The authors developed a method based on a generalized propensity score (GPS) that accommodates a continuous treatment variable.

The generalized conditional unconfoundedness assumption implies that the potential outcome for all possible treatment dosages must be independent of the actual received treatment dosage, conditional on a set of control variables. Imbens (2000) demonstrates that, similar to the propensity score method discussed in the previous subsection, conditioning on the GPS suffices to eliminate all bias arising from the non-random treatment dosage assignment. The GPS is a scalar function that indicates the likelihood of receiving the actual treatment dosage based on a set of covariates.

The estimation procedure follows Hirano and Imbens (2004) and their application of the GPS for assessing continuous treatments. In the first step, we regress the continuous treatment on a constant and a set of control variables using OLS. From this model, we obtain predicted values for each observation and the standard deviation of the residuals. The GPS for each observation is derived by assuming a normal distribution of the disturbances. Therefore, we use the conditional Gaussian density to estimate the GPS ($\hat{r}(w, X_i)$) for each observation.

With the GPS at our disposal, we follow Hirano and Imbens (2004) in modeling the treatment effect using a quadratic conditional mean model. If conditional unconfoundedness holds, the model fully accounts for all potential confounding variations in the covariate set. It includes a quadratic term and an interaction term between the treatment dosage and the GPS to impose weaker assumptions about the functional form of the treatment effect. The model is estimated using OLS regression, as illustrated in Equation 18. Due to the uncertainty in the first-stage estimation of the GPS ($\hat{r}(w, X_i)$), standard errors are obtained using the Delta method.

$$Y_i = \alpha_0 + \alpha_1 w + \alpha_2 w^2 + \alpha_3 \hat{r}(w, X_i) + \alpha_4 \hat{r}(w, X_i)^2 + \alpha_5 w \hat{r}(w, X_i) + \epsilon_i \tag{18}$$

Once all estimations are conducted, we obtain a dose-response function and plot it, including its confidence bands. Based on this function, we can assess the treatment effect of salience on protest outcomes at each level of the salience distribution.

Identification in this empirical setting similarly relies on the conditional unconfoundedness assumption, which in this context requires that the salience level is independent of potential protest outcomes conditional on the observed covariates. To address this challenge of covariate selection, we mirror our previous identification strategy by implementing the LASSO double-selection methodology. This systematic procedure yields a new set of control variables that are uniformly applied across all specifications estimating the causal impact of salience on protest outcomes.

III. Distributed Lag Model

To assess the dynamic persistence of context shocks, we estimate a distributed lag model (DLM) presented in Equation 19. This contributes to assessing the intertemporal dynamics of context shocks as theorized in Hypothesis 4. We are interested in the lagged effects of context shocks on (1) protest mobilization and (2) salience level. Therefore, W is a binary variable indicating whether there was a fatal police encounter in a given year or not. The outcome variable Y is the (transformed) protest variable, and, in the second specification, the salience level. Equation 19 is then estimated using an OLS regression with heteroskedasticity-robust standard errors. Since salience is observed at a higher regional level, the results for the effect of context shocks on the current salience level are displayed for both the county level and the Google Trends regional level.

$$Y_i = \alpha_0 + \alpha_1 W_{y,i} + \alpha_2 W_{y-1,i} + \alpha_3 W_{y-1,i} + ...\beta X_i + \epsilon_i \tag{19}$$

The overall effect of context shocks on both the protest outcomes and salience level is assessable by computing the long-run propensity (LRP). The LRP is the sum of all lagged coefficients for which we are able to compute standard errors, too. These are obtained in a regression of the outcome variable on a constant, the contemporaneous treatment, the first difference of all past treatments, and a set of controls. The contemporaneous treatment in this regression represents the LRP, including a valid standard error to assess the joint significance of all treatments (Wooldridge, 2016).

3.3. Robustness and Identification Considerations

To complement the main estimation results, we conduct a series of robustness checks, with detailed results presented in the Appendix and referenced throughout the results section. First, we evaluate covariate balance and placebo treatment effects to assess the credibility of the identification strategy. Second, we outline our approach to account for potential spatial spillovers, which may bias treatment effects if left unaddressed. To address empirical concerns, we implement alternative specifications designed to mitigate two key issues: the influence of outliers and heterogeneity across margins of the dependent variable. The strategies employed to address these concerns are discussed in this subsection. We omit further discussion of confounding and overlap, as these issues have already been addressed through LASSO-based covariate selection and sample trimming procedures described earlier.

(1) Covariate Balance. Table 8 in the Appendix D.3 presents balance diagnostics and placebo tests for the baseline specification. Standardized mean differences (SMDs) are used to assess the extent to which each estimator achieves covariate balance between the treatment and control groups. Overall, most covariates exhibit satisfactory balance across estimators. However, notable discrepancies remain for the *Gini coefficient* and *collective efficacy*, which show residual differences in means between groups. To formally assess balance, we employ the overidentification test proposed by Imai and Ratkovic (2014), which fails to reject the null hypothesis of covariate balance, thereby supporting the comparability of treatment and control groups under the preferred specifications.[14] In addition,

[14]This test is only applicable for weighting estimators. The corresponding p-values of the overidentification test for IPW and IPW-RA are 0.80 and 0.89, respectively.

results from the placebo treatment test suggest that the IPW, IPW-RA, and propensity score matching estimators outperform the RA and covariate matching estimators in mitigating confounding bias. Consequently, estimates derived from the RA and covariate matching approaches should be interpreted with caution, while those from the weighting and propensity score matching estimators appear more credible and robust.[15]

(2) Spatial Spillover. There are two kinds of spatial spillovers that affect protest mobilization and salience levels across counties: nationwide and regional spillovers. The first category of nationwide spillovers does not pose a challenge to the empirical estimation strategy. The most striking event that caused spillovers across the nation and beyond has been the murder of George Floyd by a police officer. Such nationwide spillovers are interesting as they shift the context but do not pose a problem for effect identification, as both the control and the treatment groups are affected by them to the same extent. Hence, the level of protest events due to the nationwide context shock increases similarly in the treated and untreated counties and, thereby, cancels out in the estimation of the ATET.

Regional spillovers present a challenge for empirical estimation because a fatal police encounter in one county can influence neighboring counties through social connections, regional media, or social media, undermining their validity as a control group. This violates the Stable Unit Treatment Value Assumption, as the outcomes in one county may be affected by events in nearby areas. To address this concern, the study follows a sample trimming approach from Clarke (2017), excluding counties within a certain distance of treated counties to form a "clean" control group unlikely to be impacted by spillovers. For robustness, additional regressions exclude counties within a 75- and 100-mile radius to test the result's sensitivity to spatial spillovers.

(3) Outliers. A source of empirical imprecision, not necessarily bias, is extreme values in the outcome variable. Possibly, the observed effects could only hold for a few treated counties that coincide with the outliers of protest mobilization. To assess whether the described mechanisms hold in general, we conduct robustness checks for a transformed dependent variable. Therefore, all estimations with the number of protest events as a dependent variable are also estimated with the inverse hyperbolic sine transformation (IHS, $arsinh\,(x) = ln(x + \sqrt{x^2+1})$) of the dependent variable. This transformation is an

[15] We attribute the residual significance of the placebo treatment for the IPW, IPW-RA, and propensity score matching estimators in 2021 to the time persistence of protest events across years. Specifically, the correlation of protest events in 2022 with those in 2021, 2020, and 2019 is 0.88, 0.69, and 0.62, respectively.

approximation of the log transformation, but is also valid for zero values. Furthermore, the transformation allows for an (approximate) interpretation of the results in percentages.

(4) Extensive Margin. Another strategy to reduce the influence of outliers and assess the extensive margin is to transform the dependent variable into a binary variable. The binary variable indicates one if there was at least one protest event in the given county and zero otherwise. The estimation, thus, assesses the change in the likelihood that at least one protest occurs as a consequence of treatment. This helps us to evaluate whether the results are only driven by the intensive margin or whether the treatment also has an impact on the extensive margin.

4. Empirical Results

The results are presented in the following manner, progressing from the most general to the most specific hypothesized relationships. We begin by estimating the overall impact of context shocks on protest outcomes using the baseline static model (Hypothesis 1). Next, we examine whether salience serves as a mediating mechanism between context shocks and protest activity, thereby testing Hypothesis 2. Building on these findings, we evaluate Hypothesis 3, which posits that larger context shocks exert stronger effects on both mobilization and salience of protesting. Time dependencies are assessed relying on the DLM estimation strategy (Hypothesis 4). A comprehensive set of robustness checks in Subsection 4.5 addresses potential concerns related to model specification and identification.

4.1. Context Shocks and Protest Outcomes

We begin by discussing the evidence collected to answer Hypothesis 1, which theorizes that a context shock should impact protest mobilization. As outlined in Subsection 3.2, we investigate whether a context shock caused BLM or anti-police protest events. Additionally, a sensitivity analysis is performed in the appendix to demonstrate robustness.

The results presented in Table 2 indicate that there is a substantial effect of a context shock on protest mobilization. For those counties that are treated, the number of protest events increased by 8.5-10.5 units as a consequence of a context shock. Given that the mean of BLM and anti-police protest events in all counties for 2022 is 1.4, this is a substantial increase and highlights the importance of the surrounding context for protest coordination.

Table 2: Binary Treatment Evaluation: Context Shock on Protest Outcomes 2022

	RA	IPW	IPW RA	C. Match	P.S. Match
Protest Count	10.463**	8.768*	7.913*	9.044**	9.728**
	(4.364)	(4.474)	(4.560)	(4.525)	(4.292)
Observations	3007	3007	3007	3007	3007
Inverse hyper. Y	0.734***	0.469***	0.300**	0.648***	0.244
	(0.110)	(0.133)	(0.122)	(0.126)	(0.155)
Observations	2968	2968	2968	2968	2968
Binary Y	0.242***	0.129***	0.083**	0.193***	0.058
	(0.034)	(0.041)	(0.042)	(0.046)	(0.036)
Observations	3007	3007	3007	3007	3007

Heteroskedasticity robust SE in parentheses; * p <0.1, ** p <0.05, *** p <0.01

Note: Estimation results of the Average Treatment Effect on the Treated (ATET) from the context shock on protest events for the year 2022. For the protest count Covariate Set 1, the inverse hyperbolic Y Covariate Set 2, and the binary Y Covariate Set 3 are applied.

A first robustness check assesses how sensitive the estimates are to selection on unobservables by applying the Oster bounds introduced in Subsection 3.2. With the specification-specific $R_{\max} = 1.3\tilde{R}^2$, the bounds are $\delta = 1.49$ for the protest count, $\delta = 1.32$ for the inverse hyperbolic sine transformation, and $\delta = 1.21$ for the binary outcome. Hence, unobserved confounders would have to be 1.2 to 1.5 times as important as the LASSO-selected controls to fully explain the estimated treatment effects. All three values exceed the threshold of $\delta = 1$ (Oster, 2019), indicating that unobserved confounding is unlikely to drive the estimated treatment effects.

The results demonstrate considerable robustness to model selection choices, as becomes evident from the alternative specifications. Estimates using the transformed dependent variable suggest that a context shock increases the number of protest events by 32.4-73.4 percent and raises the probability of at least one protest occurring in a county by 8.3-24.2 percentage points (all effects are estimated on the treated units). These findings remain robust to sample trimming procedures designed to address potential spillover effects and violations of the overlap assumption, as demonstrated in Table 9 in Appendix E.1. The larger coefficient in the spatial spillover specification indicates the presence of spillover effects. In particular, counties adjacent to treated units appear to experience more protest as a consequence of their neighboring counties being treated, which may inflate protest levels in the control group. These spatial spillovers likely bias the baseline estimates downward by attenuating the difference between treated and control units.

The results, along with the robustness checks, support Hypothesis 1, providing confidence that fatal police encounters qualify as context shocks for BLM and anti-police protests. The hypothesis states that context shocks generally impact protest mobilization. We observe significant results across several specifications and robustness tests, reinforcing Hypothesis 1.

4.2. Disentangling Salience from Alternative Mechanisms

After establishing the initial link between a context shock and protest mobilization in the previous subsection, we focus on the explicit salience mechanism in the following paragraphs. As modeled in Section 2, context shocks are theorized to affect the protest outcome by changing the salience of protesting. Hypothesis 2 is thus divided into two mechanisms. Mechanism 1 implies that a fatal police encounter works as a positive salience shift by increasing the salience in a given location. The subsequent Mechanism 2 hypothesizes that larger salience shifts cause more protesting. Both mechanisms are empirically assessed by relying on binary and continuous treatment evaluation techniques presented in Subsection 3.2.

Mechanism 1 is robustly identified across a variety of estimation techniques. The standard binary treatment technique in the first line of Table 3 indicates that a context shock increases the salience level by 2.6-3.5 units. This corresponds to a salience increase of 26-35% of the standard deviation of the salience measure.[16] The effect of context shocks on salience remains robust to the transformation of the dependent variable, accounting for spillovers, addressing overlap, and shifting the observational unit to the Google Trends regions (Table 10 in Appendix E.1).

Table 3: Static Treatment Evaluation: Context Shock on Salience

	RA	IPW	IPW RA	C. Match	P.S. Match
Context Shock	2.612***	3.388***	3.528***	2.704***	2.979***
	(0.830)	(0.868)	(0.866)	(0.858)	(0.757)
Observations	2924	2924	2924	2924	2924

Heteroskedasticity robust SE in parentheses; * p <0.1, ** p <0.05, *** p <0.01

Note: Estimation results of the Average Treatment Effect on the Treated (ATET) from the context shock on salience for the year 2022. LASSO selection Covariate Set 4 is applied.

Mechanism 2 is also identified relying on the estimation technique for continuous treatments presented in Subsection 3.2. Figure 8 shows the dose-response function of each salience level for its effect on protest mobilization. At the mean level (salience equals 38), a one-unit increase in salience leads to a 1.9 increase in protest events. The effect becomes insignificant for salience values larger than 46, which corresponds to the top 17th percentile. The effect is robust to transformations of the dependent variable, as depicted

[16]The salience distribution is illustrated in Figure 9 in Appendix E.2.

in Figure 8 in Appendix E.1. For these specifications, the effect is statistically significant at the 5% level for salience values between 20 and 60, which corresponds to 92% of all salience values.

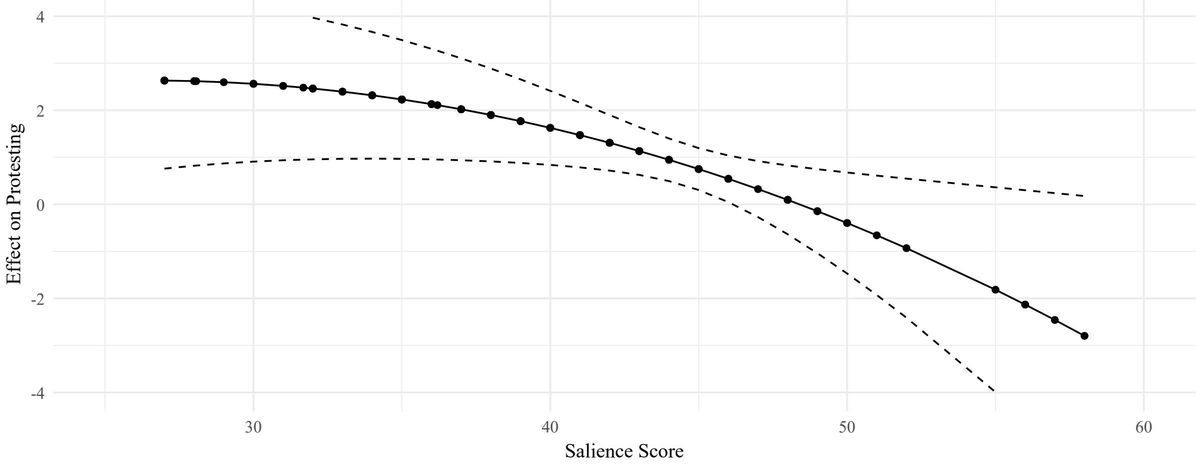

Figure 4: The graph illustrates the estimation results of the continuous salience treatment on protest events. The dashed lines indicate the 95% confidence interval. The salience score ranges from 20 to 60, which includes 92% of all salience values. Covariate Set 2 is applied.

By combining the estimated effects of Mechanisms 1 and 2, we derive an overall estimate for the entire channel, which we can compare to the results of the complete channel presented in Subsection 4.1. Specifically, taking the average coefficient of Mechanism 1 (3.02) and multiplying it by the density-weighted coefficient[17] of Mechanism 2 (2.21) yields a total effect of 6.69. This is about 72% of the total effect size of the average effect obtained in the previous subsection for the whole channel (9.3). Thus, we conclude that the salience mechanism is the dominant effect for short-run protest mobilization. In subsequent sections, we return to the question of what drives the remaining gap between the estimated indirect effect via salience and the total effect.

In conclusion, the results support Hypothesis 2 as they provide evidence that context shocks work as positive salience shifts on protest outcomes. This effect holds for a variety of robustness tests. We continue this analysis in the following subsection, focusing on heterogeneity in the magnitude of context shocks and their implications for protest outcomes.

[17]Only statistically significant values are considered, as we want to avoid extrapolating for high values of salience where almost no observed values are.

4.3. Differently Sized Context Shocks

To evaluate Hypothesis 3, which states that larger context shocks lead to stronger protest outcomes by increasing salience, we examine the estimation results presented in Table 4. We operationalize the magnitude of a context shock using the volume of media coverage associated with a specific fatal police encounter, measured by article counts. In this framework, greater media attention is expected to enhance salience, which, in turn, is predicted to lead to increased protest mobilization.

Distinguishing between context shocks, those that received significant media attention appear to have led to more protest events. This tendency is evident in the increasing coefficients for the regression outcomes in lines 1-3 in Table 4. For the regression adjustment specification, a small context shock increases protest by 9 units, a medium-sized shock by 14 units, and a large shock by 24.9 units on treated units. The size of the context shock, therefore, seems to substantially influence protest coordination. It should be noted that the estimates for these results are imprecisely estimated with relatively large standard errors. The salience shift appears to be highest for large context shocks, for which the salience level increases by 5.2-5.7 units. However, the coefficient sizes for small and medium context shocks are roughly the same, especially when accounting for the standard errors.

The empirical findings are broadly consistent with the hypothesis, although they are not statistically significant. In both tables, we observe a positive association between the size of the context shock and the magnitude of its effect: larger context shocks tend to correspond to larger estimated coefficients for both protest activity and salience level. While overlapping confidence intervals of the point estimates prevent definitive conclusions about statistical differences between these coefficients, the overall pattern suggests a consistent relationship, particularly between high-salience events and mobilization.

Importantly, the article count is used to categorize context shocks based on the strength of the *external shock* on individuals' attention. If the analysis involved regressing article count on protest intensity or salience, the estimates would likely be biased because heightened protest activity or salience could drive increased media coverage. Instead, our approach involves excluding a subset of context shocks, specifically those that received little media attention, and examining whether the estimated treatment effects on protest incidence or salience differ from those of more widely covered shocks.

In sum, the results provide suggestive evidence that context shocks in the top quartile of media attention have the strongest effect on subsequent protest mobilization, with a

Table 4: Context Dependence: Context Shocks on Protest Outcomes

	RA	IPW	IPW RA	C. Match	P.S. Match
Dependent Variable: Protest Outcome					
Article C. ≥ 3	9.019*	6.358	5.018	7.083	3.883
	(4.858)	(5.139)	(5.533)	(5.125)	(5.913)
Observations	2905	2905	2905	2905	2905
Article C. ≥ 6	13.992*	10.201	8.860	12.414*	12.397*
	(7.171)	(7.531)	(7.978)	(7.194)	(6.974)
Observations	2875	2875	2875	2875	2875
Article C. ≥ 27	24.937*	21.277	19.491	21.214	18.250
	(14.063)	(14.198)	(15.152)	(14.372)	(14.789)
Observations	2845	2845	2845	2845	2845
Dependent Variable: Salience Level					
Article C. ≥ 3	3.443**	4.256***	4.531***	3.481**	3.345**
	(1.338)	(1.409)	(1.410)	(1.486)	(1.579)
Observations	2824	2824	2824	2824	2824
Article C. ≥ 6	2.998*	3.973**	4.345***	3.046*	2.793
	(1.639)	(1.679)	(1.675)	(1.769)	(2.165)
Observations	2794	2794	2794	2794	2794
Article C. ≥ 27	5.198**	5.659**	4.569*	5.155**	5.619***
	(2.279)	(2.422)	(2.438)	(2.567)	(1.555)
Observations	2764	2764	2764	2764	2764

Heteroskedasticity robust SE in parentheses; * p <0.1, ** p <0.05, *** p <0.01

Note: Estimation results of the Average Treatment Effect on the Treated (ATET) from the context shock on protest events and salience for the year 2022. The threshold values represent the top 75th, 50th, and 25th percentiles of the article count distribution. LASSO selection Covariate Set 1 is applied for the first 3 sets of estimators and Covariate Set 4 for the latter 3 sets.

substantial portion of this effect mediated through increased salience. The estimation of these effects is relatively imprecise, potentially due to the reduced number of context shocks at higher media-attention thresholds. The difference in response between small and medium-sized shocks remains uncertain.

4.4. Dynamic Effects

The final empirical estimation examines the dynamics of the salience mechanism. Following Hypothesis 4, a return of the context to its original value after a high-profile event should be accompanied by a decrease in protest mobilization. These dynamics are assessed by examining the publication gap between fatal police encounters and their media coverage, as well as a DLM that evaluates the intertemporal effects of context shocks on the salience level and protest outcomes.

The first piece of evidence is presented in Figure 5, which illustrates the timing of media coverage for fatal police encounters. Within the first 30 days, more than one-third of articles are published about these incidents. This indicates that the context, as defined in the model, is shocked most heavily in the immediate aftermath of a fatal police encounter. Over time, media interest in such incidents stabilizes at a low level until after four years, when more than 85% of the articles are published. The premise of Hypothesis 4, that the context normalizes over time, is therefore consistent with the data.

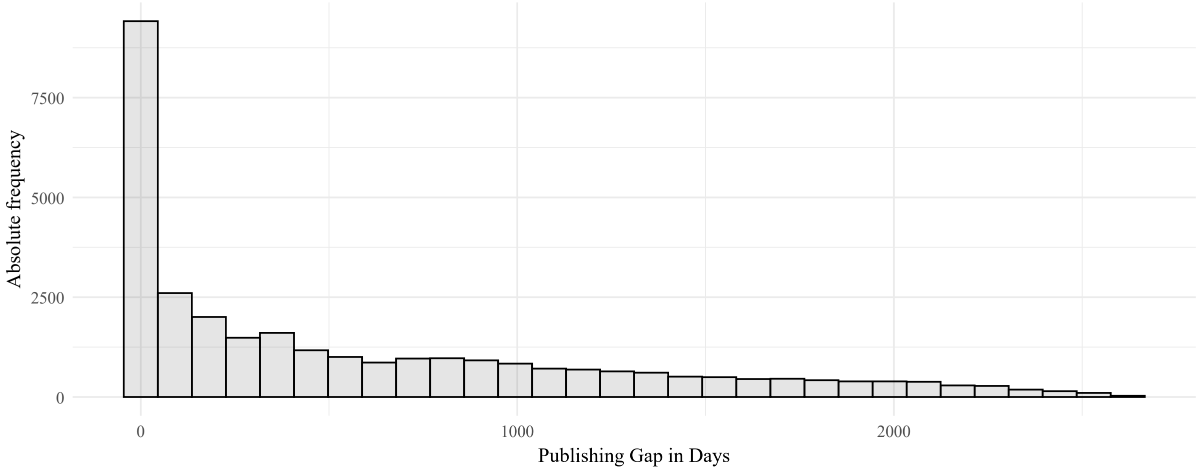

Figure 5: Publishing gap illustration as the difference in days between a fatal police encounter involving a Black individual and the publication of related articles. The analysis includes all incidents occurring between 2017 and 2022. Given this timeframe, the maximum observable publication lag for the most recent incidents is three years. 35.8% of articles are published within the first month after the fatal police encounter.

Assessing the timing of the salience mechanism, the presented evidence supports the proposed mechanism's substantial explanatory power for short-term mobilization. To evaluate the persistence of salience-based protesting, we estimate a DLM, which is presented in Table 5. The contemporaneous effects identified in the previous regressions are replicated, as the context shock in 2022 significantly increases both the protest outcome and salience level. Regarding sustained mobilization, the context shocks from previous years continue to influence protest mobilization in the current period, but decrease in magnitude in the more precise second specification. The effect of context shocks on the salience level disappears after one period and even turns negative for one lag, speaking in favor of something like an "attention-hangover" after the treatment in 2020. We conclude that the salience mechanism plays a significant explanatory role in the short run, after which salience levels rapidly revert to their pre-treatment state. Consequently, the effect theorized in Hypothesis 4 materializes within the first year after a context shock, after which no statistically significant shift in salience is observed.

So, if the lagged effects of context shocks are not driven by the salience mechanism,

Table 5: Distributed Lag Model: Context Shocks on Protest Outcomes.

Dependent Variable	*Protest Events*		*Salience Level*	
	Level	Inverse hyp.	County Level	Regional Level
Context Shock 2018	5.421	0.202**	1.155	1.715
	(3.297)	(0.094)	(0.795)	(1.757)
Context Shock 2019	3.559	0.267***	-0.913	-1.988
	(2.571)	(0.090)	(0.815)	(1.807)
Context Shock 2020	4.603	0.324***	-1.829**	-5.479**
	(2.809)	(0.096)	(0.725)	(2.384)
Context Shock 2021	5.772*	0.344***	0.091	1.514
	(3.266)	(0.096)	(0.769)	(1.695)
Context Shock 2022	3.473*	0.399***	2.751***	6.631***
	(2.075)	(0.091)	(0.865)	(1.705)
LRP	22.828**	1.536***	1.255	2.393
	(8.977)	(0.159)	(1.274)	(2.084)
Observations	3007	3007	2924	201

Heteroskedasticity robust SE in parentheses; * p <0.1, ** p <0.05, *** p <0.01

Note: Distributed Lag Model for the effects of past and current context shocks on protest outcomes and salience level in 2022. LRP denotes the long-run elasticity. Covariate Set 1 is applied for protest events and Covariate Set 4 for salience as dependent variable.

what else drives these effects? A number of candidates are discussed in the literature. For one, networks may form during the early mobilization period. Thus, as soon as a context shock occurs, the salience mechanism mobilizes individuals to protests, where they form protest networks. Such networks may be capable of sustaining protest mobilization within a smaller circle of individuals without relying on increased attention in the larger regional environment. This mechanism aligns with the findings of Bursztyn *et al.* (2021), who show that initial protest participation, when reinforced by social network effects within close peer groups, can generate sustained mobilization over time. A complementary mechanism may be that organizational capacities were built during the early protest stage, facilitating later mobilization, as networks could.

In summary, we presented a variety of empirical evidence supporting the existence of a salience mechanism. This mechanism has immediate effects and socio-economic significance. However, regarding sustained mobilization in the aftermath of high-profile events, the explanatory power of the proposed mechanism is low, indicating that other approaches are needed for explanations. Given the current dynamics in attention markets, such as social media, this finding should not be too surprising. The only thing faster than the emergence of hype is its inevitable end, which also seems to be true for attention-motivated protesting.

4.5. Robustness

The regression results in the previous subsection draw a clear picture of the proposed salience mechanism from a context shock to increased protesting via salience. In this sensitivity analysis, we present supplementary empirical evidence to demonstrate the robustness of the salience mechanism. To this end, we estimate robustness checks for the period between 2018 and 2022.

To demonstrate the robustness of the salience mechanism across different years, we assess the binary treatment model and estimate it for the years 2018, 2020, and 2022. Generally, the effect of context shocks on protest outcomes varies substantially across these years, as the estimates in Table 11 in Appendix E.3 illustrate. A context shock is found to increase protest events in treated counties by 18.1 to 25 units in 2020. This effect is smaller for 2022; here, a context shock increases protest events by 8.5-10.5 units. The smallest effects are identified for 2018, where the context shock led to an increase of 0.5-0.6 units only. The results, therefore, indicate a much greater sensitivity of protest outcomes to context shocks since 2020, which likely is the result of an interaction effect between the context shock and the murder of George Floyd in the same year.

Besides the results presented in Table 11 in Appendix E.3, there are further robustness checks presented in Appendix E.3 and E.4. In Table 12 in Appendix E.3, the baseline regressions for 2019 and 2021 generally show mixed results, with some specifications in line with the previous findings, while others lack statistical identification. For the years 2018 and 2020, we run the complete set of robustness checks (Table 12 in Appendix E.3). The results generally demonstrate the robustness of the previous findings. For 2018, the specification where we trimmed extreme propensity scores loses significance, potentially due to smaller statistical power as a consequence of fewer observations and a generally small point estimate that we find prior to 2020. For 2020, the extensive margin is not statistically identified, potentially due to the widespread mobilization of protests during this year. As many counties showed protests anyway during this year, it is likely to become infeasible to identify the extensive margin for this year. However, the eight alternative robustness checks for both years are in line with the previous results and thereby demonstrate substantial robustness of the empirical identification. The effect of salience on protest outcomes is estimated using the same continuous treatment model as in Subsection 4.2 and presented in Figure 10 in Appendix E.4. The effect size of salience on protests is generally much higher in 2020 compared to 2022 and decreases in size for 2018, where the effect is barely statistically significant.

5. Conclusion

In this paper, we develop an experience-based salience model that incorporates the attentional context in protest decisions. In this framework, high-profile events such as fatal police encounters with Black individuals act as context shocks by altering the perception of protests and drawing attention to the benefits of protesting, thereby increasing their salience. This shifted salience can change protest decisions and increase participation, even when the objective benefits or costs remain unchanged. We assess the model's predictions using a set of treatment evaluation methods in the context of BLM and anti-police violence protests in the US. Overall, we find that fatal police encounters significantly affect protest outcomes, with a large portion of this effect mediated through salience. We distinguish between the sizes of the fatal police encounters by categorizing them according to their article count and evaluating their intertemporal dynamics using a DLM. The results indicate that larger context shocks exert a stronger influence on protest mobilization than smaller ones. Moreover, attention-driven protest activity tends to emerge rapidly following such shocks but dissipates within approximately one year.

The paper contributes to a rich literature on protest models. The attentional channel has the potential to be incorporated into many of the previous protest models. For instance, incorporating critical events and their attentional consequences into a threshold model or a model in which protests are strategic complements could help to explain fast-paced protest mobilization. The salience mechanism should be understood as complementary to previous protest models, as it explicitly incorporates subconscious perceptual aspects of protest mobilization.

One promising path for future research is to more thoroughly assess the dynamic aspects of attention-based protesting. The results of this study indicate that attention-motivated protesting is a short-term phenomenon that cannot fully explain sustained protesting. Other factors, such as the formation of networks, organizational capacities, or information cascades, are therefore more likely to shed light on intertemporal protest dynamics. A broader perspective that combines explanatory elements from multiple theories may offer a more accurate and precise account of intertemporal protest dynamics.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process During the preparation of this manuscript, the authors used ChatGPT to assist in formulating and refining text. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the final content of the article.

References

Abadie, A. and Cattaneo, M. D. (2018). Econometric methods for program evaluation. *Annual Review of Economics*, **10** (1), 465–503.

Bailey, M., Cao, R., Kuchler, T., Stroebel, J. and Wong, A. (2018). Social connectedness: Measurement, determinants, and effects. *Journal of Economic Perspectives*, **32** (3), 259–280.

Barbera, S. and Jackson, M. O. (2019). A model of protests, revolution, and information. *Quarterly Journal of Political Science*, **15** (3), 297–335.

Belloni, A., Chernozhukov, V. and Hansen, C. (2014). Inference on treatment effects after selection among high-dimensional controls. *Review of Economic Studies*, **81** (2), 608–650.

Bilgin, G. E. and Destan, C. G. (2024). *Voting Under Salience Bias and Strategic Extremism*. Tech. rep., University of Bonn and University of Mannheim, Germany.

Bonomi, G., Gennaioli, N. and Tabellini, G. (2021). Identity, beliefs, and political conflict. *The Quarterly Journal of Economics*, **136** (4), 2371–2411.

Bordalo, P., Gennaioli, N. and Shleifer, A. (2012). Salience theory of choice under risk. *The Quarterly journal of economics*, **127** (3), 1243–1285.

—, — and — (2013). Salience and consumer choice. *Journal of Political Economy*, **121** (5), 803–843.

—, — and — (2020a). Memory, attention, and choice. *The Quarterly journal of economics*, **135** (3), 1399–1442.

—, Tabellini, M. and Yang, D. Y. (2020b). *Issue salience and political stereotypes*. Tech. rep., National Bureau of Economic Research.

Bursztyn, L., Cantoni, D., Yang, D. Y., Yuchtman, N. and Zhang, Y. J. (2021). Persistent political engagement: Social interactions and the dynamics of protest movements. *American Economic Review: Insights*, **3** (2), 233–250.

Cantoni, D., Heizlsperger, L.-J., Yang, D. Y., Yuchtman, N. and Zhang, Y. J. (2022). The fundamental determinants of protest participation: Evidence from hong kong's antiauthoritarian movement. *Journal of Public Economics*, **211**, 104667.

—, Yang, D. Y., Yuchtman, N. and Zhang, Y. J. (2019). Protests as strategic games: experimental evidence from hong kong's antiauthoritarian movement. *The Quarterly Journal of Economics*, **134** (2), 1021–1077.

Castelnuovo, E. and Tran, T. D. (2017). Google it up! a google trends-based uncertainty index for the united states and australia. *Economics Letters*, **161**, 149–153.

Chenoweth, E. and Pressman, J. (2017). Crowd Counting Consortium. [Data set].

Chernozhukov, V., Hansen, C. and Spindler, M. (2016). hdm: High-dimensional metrics. *R Journal*, **8** (2), 185–199.

Choi, H. and Varian, H. (2012). Predicting the present with google trends. *Economic record*, **88**, 2–9.

Clarke, D. (2017). Estimating difference-in-differences in the presence of spillovers.

Congress, U. (2018). The geography of social capital in america. [Data set].

Crump, R. K., Hotz, V. J., Imbens, G. W. and Mitnik, O. A. (2009). Dealing with limited overlap in estimation of average treatment effects. *Biometrika*, **96** (1), 187–199.

De Mesquita, E. B. (2010). Regime change and revolutionary entrepreneurs. *American Political Science Review*, **104** (3), 446–466.

Enikolopov, R., Makarin, A. and Petrova, M. (2020). Social media and protest participation: Evidence from russia. *Econometrica*, **88** (4), 1479–1514.

González-Bailón, S., Borge-Holthoefer, J., Rivero, A. and Moreno, Y. (2011). The dynamics of protest recruitment through an online network. *Scientific reports*, **1** (1), 1–7.

Granovetter, M. (1978). Threshold models of collective behavior. *American journal of sociology*, **83** (6), 1420–1443.

Hirano, K. and Imbens, G. W. (2004). The propensity score with continuous treatment. *Applied Bayesian modeling and causal inference from incomplete-data perspectives, edited by Gelman A. Meng XL. West Sussex: Wiley InterScience*, pp. 73–84.

Huber, M. (2023). *Causal analysis: Impact evaluation and Causal Machine Learning with applications in R.* MIT Press.

Imai, K. and Ratkovic, M. (2014). Covariate balancing propensity score. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, **76** (1), 243–263.

Imbens, G. W. (2000). The role of the propensity score in estimating dose-response functions. *Biometrika*, **87** (3), 706–710.

James, G., Witten, D., Hastie, T., Tibshirani, R. *et al.* (2021). *An introduction to statistical learning*, vol. 112. Springer.

Kann, C., Hashash, S., Steinert-Threlkeld, Z. and Alvarez, R. M. (2023). Collective identity in collective action: evidence from the 2020 summer blm protests. *Frontiers in Political Science*, **5**, 1185633.

Khondker, H. H. (2011). Role of the new media in the arab spring. *Globalizations*, **8** (5), 675–679.

Larson, J. M., Nagler, J., Ronen, J. and Tucker, J. A. (2019). Social networks and protest participation: Evidence from 130 million twitter users. *American Journal of Political Science*, **63** (3), 690–705.

Lecy, J. (2023). Nccs legacy core series. [Data set].

Little, A. T. (2016). Communication technology and protest. *The Journal of Politics*, **78** (1), 152–166.

Lohmann, S. (1994). The dynamics of informational cascades: The monday demonstrations in leipzig, east germany, 1989–91. *World politics*, **47** (1), 42–101.

— (2000). Collective action cascades: An informational rationale for the power in numbers. *Journal of Economic Surveys*, **14** (5), 655–684.

Malik, A. and Awadallah, B. (2013). The economics of the arab spring. *World Development*, **45**, 296–313.

Meirowitz, A. and Tucker, J. A. (2013). People power or a one-shot deal? a dynamic model of protest. *American Journal of Political Science*, **57** (2), 478–490.

Mellon, J. (2013). Where and when can we use google trends to measure issue salience? *PS: Political Science & Politics*, **46** (2), 280–290.

MIT (2018). County presidential election returns 2000-2020.

Moreno-Medina, J., Ouss, A., Bayer, P. and Ba, B. A. (2025). Officer-involved: The media language of police killings*. *The Quarterly Journal of Economics*, **140** (2), 1525–1580.

Neyman, J. (1959). Optimal asymptotic tests of composite hypotheses. *Probability and statsitics*, pp. 213–234.

Oster, E. (2019). Unobservable selection and coefficient stability: Theory and evidence. *Journal of Business & Economic Statistics*, **37** (2), 187–204.

Rosenbaum, P. R. and Rubin, D. B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika*, **70** (1), 41–55.

Shadmehr, M. (2018). Protest puzzles: tullock's paradox, hong kong experiment, and the strength of weak states. *Quarterly Journal of Political Science, Forthcoming.*

Stephens-Davidowitz, S. (2014). The cost of racial animus on a black candidate: Evidence using google search data. *Journal of Public Economics*, **118**, 26–40.

Wiedermann, M., Smith, E. K., Heitzig, J. and Donges, J. F. (2020). A network-based microfoundation of granovetter's threshold model for social tipping. *Scientific reports*, **10** (1), 11202.

Williamson, V., Trump, K.-S. and Einstein, K. L. (2018). Black lives matter: Evidence that police-caused deaths predict protest activity. *Perspectives on Politics*, **16** (2), 400–415.

Wooldridge, J. M. (2007). Inverse probability weighted estimation for general missing data problems. *Journal of econometrics*, **141** (2), 1281–1301.

— (2010). *Econometric analysis of cross section and panel data*. MIT press.

— (2016). *Introductory Econometrics: A Modern Approach 6rd ed.* Cengage learning.

Yin, C.-C. (1998). Equilibria of collective action in different distributions of protest thresholds. *Public Choice*, **97** (4), 535–567.

Appendix

A. Mathematical Appendix

Proof of Lemma 1: For discrete-time $s \in \{0, ..., t\}$ the reference benefit at time t is given by

$$\bar{b}_t = \sum_{b_s \in M_t} b_s \cdot w(e_s, e_t) = \sum_{s=0}^{t} b_s \cdot w(e_s, e_t).$$

Then it holds $\bar{b}_t > (<)\ \bar{b}_{t-1}$ if and only if

$$\begin{aligned}
& \frac{\sum_{s=0}^{t} S(e_s, e_t) \cdot b_s}{\sum_{s=0}^{t} S(e_s, e_t)} - \bar{b}_{t-1} && \gtrless 0 \\
\Leftrightarrow \quad & \frac{\sum_{s=0}^{t} S(e_s, e_t) \cdot b_s - \bar{b}_{t-1} \cdot \sum_{s=0}^{t} S(e_s, e_t)}{\sum_{s=0}^{t} S(e_s, e_t)} && \gtrless 0 \\
\Leftrightarrow \quad & \frac{\sum_{s=0}^{t} S(e_s, e_t) \cdot [b_s - \bar{b}_{t-1}]}{\sum_{s=0}^{t} S(e_s, e_t)} && \gtrless 0 \\
\Leftrightarrow \quad & \sum_{s=0}^{t} S(e_s, e_t) \cdot [b_s - \bar{b}_{t-1}] && \gtrless 0
\end{aligned}$$

As $f(k_s) = b_s$ is a continuous and injective function, a global context shock k_t leads to the perceived benefit b_t reaching a new extreme level, i.e., $b_t > max_{s<t} b_s$ or $b_t < min_{s<t} b_s$. Thus, if k_t induces a new maximum (minimum) level it holds

$$S(e_i, e_t) > S(e_j, e_t) \Leftrightarrow b_i > (<)\ b_j, \quad \forall i, j \in \{0, ..., t\}$$

since $S(e_s, e_t)$ is strictly decreasing in $|k_t - k_s|$. Then, by Chebyshev's sum inequality, it holds

$$\begin{aligned}
& t \cdot \sum_{s=0}^{t} S(e_s, e_t) \cdot [b_s - \bar{b}_{t-1}] \gtrless \left(\sum_{s=0}^{t} S(e_s, e_t) \right) \cdot \left(\sum_{s=0}^{t} b_s - \bar{b}_{t-1} \right) \\
\Leftrightarrow \quad & \frac{\sum_{s=0}^{t} S(e_s, e_t) \cdot [b_s - \bar{b}_{t-1}]}{\sum_{s=0}^{t} S(e_s, e_t)} \gtrless \frac{1}{t} \cdot \sum_{s=0}^{t} b_s - \bar{b}_{t-1} \\
\Leftrightarrow \quad & \frac{\sum_{s=0}^{t} S(e_s, e_t) \cdot b_s}{\sum_{s=0}^{t} S(e_s, e_t)} - \bar{b}_{t-1} \cdot \frac{\sum_{s=0}^{t} S(e_s, e_t)}{\sum_{s=0}^{t} S(e_s, e_t)} \gtrless \frac{1}{t} \cdot \sum_{s=0}^{t} b_s - \frac{1}{t} \cdot \sum_{s=0}^{t} \bar{b}_{t-1} \\
\Leftrightarrow \quad & \frac{\sum_{s=0}^{t} S(e_s, e_t) \cdot b_s}{\sum_{s=0}^{t} S(e_s, e_t)} \gtrless \frac{1}{t} \cdot \sum_{s=0}^{t} b_s
\end{aligned}$$

such that for $\bar{b}_{t-1} \leq (\geq)\ \frac{1}{t} \cdot \sum_{s=0}^{t} b_s$, it always holds $\bar{b}_t > (<)\ \bar{b}_{t-1}$ and Condition (i) is shown.

In general, Inequality 10 is always fulfilled if

$$\alpha_i \geq \alpha_j \Leftrightarrow b_i \geq (\leq)\ b_j, \quad \forall i, j \in \{0, ..., t-1\}$$

with $\alpha_s = \frac{S(e_s,e_t)}{S(e_s,e_{t-1})}$ since then, by the weighted Chebyshev's sum inequality, it holds

$$\begin{aligned}
&\frac{\sum_{s=0}^{t-1} \alpha_s \cdot S(e_s, e_{t-1}) \cdot [b_s - \bar{b}_{t-1}]}{\sum_{s=0}^{t-1} S(e_s, e_{t-1})} \underset{\leq}{\geq} \left(\frac{\sum_{s=0}^{t-1} \alpha_s \cdot S(e_s, e_{t-1})}{\sum_{s=0}^{t-1} S(e_s, e_{t-1})} \right) \cdot \left(\frac{\sum_{s=0}^{t-1} S(e_s, e_{t-1}) \cdot [b_s - \bar{b}_{t-1}]}{\sum_{s=0}^{t-1} S(e_s, e_{t-1})} \right) \\
\Leftrightarrow \quad & \sum_{s=0}^{t-1} \alpha_s \cdot S(e_s, e_{t-1}) \cdot [b_s - \bar{b}_{t-1}] \underset{\leq}{\geq} \left(\frac{\sum_{s=0}^{t-1} \alpha_s \cdot S(e_s, e_{t-1})}{\sum_{s=0}^{t-1} S(e_s, e_{t-1})} \right) \cdot \left(\sum_{s=0}^{t-1} S(e_s, e_{t-1}) \cdot [b_s - \bar{b}_{t-1}] \right) \\
\Leftrightarrow \quad & \sum_{s=0}^{t-1} S(e_s, e_t) \cdot [b_s - \bar{b}_{t-1}] \underset{\leq}{\geq} 0
\end{aligned}$$

The last equivalence transformation holds since we have

$$\begin{aligned}
& \bar{b}_{t-1} = \frac{\sum_{s=0}^{t-1} S(e_s, e_{t-1}) \cdot b_s}{\sum_{s=0}^{t-1} S(e_s, e_{t-1})} \\
\Leftrightarrow \quad & \frac{\sum_{s=0}^{t-1} S(e_s, e_{t-1}) \cdot b_s - \bar{b}_{t-1} \cdot \sum_{s=0}^{t-1} S(e_s, e_{t-1})}{\sum_{s=0}^{t-1} S(e_s, e_{t-1})} = 0 \\
\Leftrightarrow \quad & \sum_{s=0}^{t-1} S(e_s, e_{t-1}) \cdot [b_s - \bar{b}_{t-1}] = 0
\end{aligned}$$

As $b_t > (<)\ \bar{b}_{t-1}$, Inequality 10 is fulfilled and Condition (ii) is shown.

q.e.d

B. Beta Densities

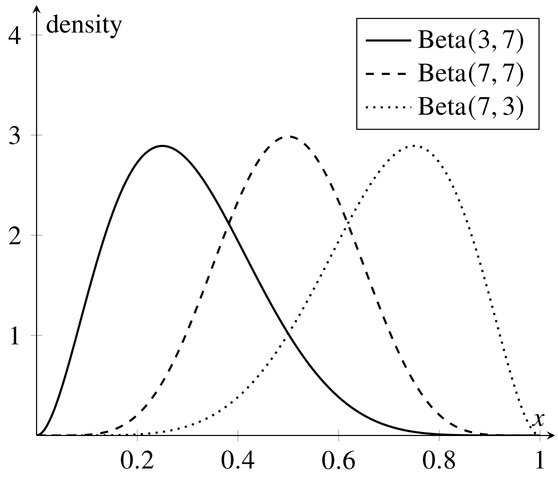

Figure 6: Densities for different $Beta(\alpha, \beta)$ distributions.

C. LASSO Double Selection

C.1. LASSO Double Selection

In order to obtain causal estimates, the estimation strategy relies on the selection-on-observables assumption (or conditional independence assumption). This assumption holds that after controlling for a set of covariates, the selection into treatment is random and, thereby, not endogenous to the outcome. One approach for justifying this assumption is to include a large set of covariates that are frequently used in the literature. However, this procedure has its downsides as it is subject to researchers' prior conceptions of fitting control variables and may lead to inflated standard errors due to multicollinearity.

LASSO (least absolute shrinkage and selection operator) double selection is a suitable method for establishing the selection-on-observables assumption in a way that avoids subjectivity and provides valid test statistics for causal inference (Belloni *et al.*, 2014; Huber, 2023). Subjective perceptions about suitable control variables have a smaller influence as the researcher only selects a large set of potential covariates. The necessary control variables for establishing the selection-on-observables assumption in the causal analysis are selected from this set in a data-driven way, reducing the researcher's involvement. Additionally, LASSO double selection increases the confidence in a model's test statistics. Treatment effect estimators (and ordinary regression methods) are generally not robust to human-made rules for variable selection, potentially leading to invalid test statistics (Huber, 2023). Procedures, such as LASSO double selection, alleviate such concerns, as the inclusion of covariates is determined in a data-driven way. These advantages, of course, only materialize if the set of potential control variables is sufficiently large. To ensure this, 76 potential covariates are selected in this study as presented in Table 6 and 7 in Appendix C.2.

Following Belloni *et al.* (2014), the selection procedure and the post-LASSO estimations rely on two assumptions: Neyman-orthogonality and approximate sparsity. Neyman-orthogonality (Neyman, 1959) ensures that the estimation of treatment effects is robust to approximation errors in the nuisance parameters (Huber, 2023). That is, the target parameter is insensitive to perturbations in the LASSO estimation due to regularization. Furthermore, approximate sparsity ensures that the number of important covariates is small relative to the sample size n. Given these assumptions, a large set of variables can be reduced using LASSO double selection, and causal analysis may be conducted via post-LASSO OLS estimation. This procedure, thanks to the given assumptions, is $\sqrt{n}$-consistent.

LASSO double selection relies on the LASSO regularization technique stemming from the predictive machine learning literature (James *et al.*, 2021). The model is presented in Equation A.1. The left-hand side of the curly bracket is the OLS minimization problem for which the squared residuals are minimized. The right-hand side introduces a penalty term for absolute coefficient size. The penalty term introduces (regularization-) bias and reduces variance. The tuning parameter λ is chosen to balance bias and variance so as to maximize predictive performance in the validation set. All coefficients that do not contribute to the predictive performance of the model (reducing the OLS term) are shrunken to 0 (to reduce the penalty term).

$$\text{minimize}_{\boldsymbol{\beta}} \left\{ \sum_{i=1}^{n} \left(y_i - \beta_0 - \sum_{j=1}^{p} \beta_j x_{ij} \right)^2 + \lambda \sum_{j=1}^{p} |\beta_j| \right\} \tag{A.1}$$

Causal inference with LASSO double selection is conducted in three steps (Belloni *et al.*, 2014). The first is a LASSO regression of the outcome variable on a constant and all potential covariates. The second is a LASSO regression of the treatment variable on a constant and all potential covariates. Each variable with a positive absolute coefficient in one of these LASSO regressions is included in a list of selected control variables. Finally, a causal OLS regression of the outcome variable on a constant, the treatment variable, and all selected variables is run. The treatment variable's coefficient is causally interpretable as all potentially confounding variance is controlled through the selected control variables.

LASSO double selection is applied to the data at hand. For the first empirical specification, two LASSO regressions are run: (1) protest events on a constant and all potential control variables, and (2) focal points on a constant and all potential control variables. The potential control variables are summarized in Table 6 and 7. All variables with a positive absolute coefficient in at least one of these regressions are included as control variables for the different estimation techniques in the following. The LASSO estimation is conducted in R using the *High-Dimensional Metrics* package (Chernozhukov *et al.*, 2016).

C.2. Included Variables for LASSO Selection Procedure

Table 6: Included Variables for LASSO Selection Procedure (1/2)

Variable	Description
population	Population number
log population	Logarithm population number
ages 0 19	Population share 0 to 19 year olds
ages 20 29	Population share 20 to 29 year olds
ages 30 39	Population share 30 to 39 year olds
ages 40 49	Population share 40 to 49 year olds
ages 50 59	Population share 50 to 59 year olds
ages 60 69	Population share 60 to 69 year olds
ages 70 79	Population share 70 to 79 year olds
ages 80 older	Population share 80 years and older
white	Polulation share white individuals
black	Population share black individuals
native	Population share native individuals
asian	Population share asian individuals
native pacific	Population share native pacific individuals
other	Population share other race individuals
two or more	Population share two or more race individuals
average household size	Average household size
married	Population share of married individuals
not married	Population share of not married individuals
highschool degree	Population share highschool degree
no degree	Population share no educational degree
university degree	Population share university degree
employed	Employment rate
unemployed	Unemployment rate
employed youth	Employment rate for 16 to 24 year olds
unemployed youth	Unemployment rate for 16 to 24 year olds
k10	Population share with income of 0 to 10,000 USD.
k30	Population share with income of 20,000 to 30,000 USD.
k40	Population share with income of 30,000 to 40,000 USD.
k50	Population share with income of 40,000 to 50,000 USD.
k60	Population share with income of 50,000 to 60,000 USD.
k75	Population share with income of 60,000 to 75,000 USD.
k100	Population share with income of 75,000 to 100,000 USD.
k125	Population share with income of 100,000 to 125,000 USD.
k150	Population share with income of 125,000 to 150,000 USD.
k20	Population share with income of 10,000 to 20,000 USD.
k200	Population share with income of 150,000 to 200,000 USD.
kmore200	Population share with income of more than 200,000 USD.

Note: Set of potential control variables, the final set of control variables is determined by LASSO double selection. The variables are retrieved from the US Census if not indicated otherwise.

Table 7: Included Variables for LASSO Selection Procedure (2/2)

Variable	Description
median household income	Median household income
log median household income	Logarithm median household income
median household income white	Median household income white individuals
log median HI white	Logarithm median household income white individuals
median household income black	Median household income black individuals
log median HI black	Logarithm median household income black individuals
median household income race gap	Median household income race gap
no other income	Population share with labor income only
other income	Population share with non-wage income
per capita income	Income per capita
log per capita income	Logarithm of income per capita
gini	Gini coefficient
rent percent householde income	Rent as percentage of household income
median gross rent	Median gross rent
log median gross rent	Logarithm median gross rent
im 1990 1999	Share of population of immigrants that came to the US beween 1990 and 1999
im 2000 2009	Share of population of immigrants that came to the US beween 2000 and 2009
im 2010 later	Share of population of immigrants that came to the US beween 2010 and later
im before 1990	Share of population of immigrants that came to the US before 1990
health insurance	Share of population with health insurance
no health insurance	Sahre of population without health insurance
election clos	Votes 2nd Cand. / Votes Winner (MIT, 2018)
pol fractionalization ELF	Political fractionalization ELF score of voter groups (MIT, 2018)
election clos poppoll	Votes 2nd cand. / votes winner (MIT, 2018)
election clos states poppoll	Votes 2nd cand. / votes winner (states) (MIT, 2018)
election clos counties poppoll	Votes 2nd cand. / votes winner (counties) (MIT, 2018)
online network	Number of Facebook connections/ possible number of Facebook connections (logarithmized, 2021 only) Bailey *et al.* (2018)
family unity	Weighted measure of births to unwed women, children living in single-parent households, and percentage of women aged 35-44 who are married (Congress, 2018)
community health	Index containing share of (1) adults who volunteered, (2) discussed community affairs, or/and (3) worked for improving the community (Congress, 2018)
institutional health	Adherence to state institution measure (Congress, 2018)
social capital	Indicator composed of different dimensions of social capital like trust, networks, or shared values (Congress, 2018)
collective efficacy	Violent crime rate (Congress, 2018)
charitable organizations	Number of charitable organizations per capita (Lecy, 2023)
area	Area of land
density	Population number divided by area size (own calculation)
center 1	Dummy variable that equals 1 if density $>$ 40,000 (own calculation)
center 2	Dummy variable that equals 1 if density $>$ 20,000 (own calculation)

Note: Set of potential control variables, the final set of control variables is determined by LASSO double selection. The variables are retrieved from the US Census if not indicated otherwise.

C.3. LASSO Selected Covariate Sets

The following are covariate sets that are obtained by applying the LASSO double selection procedure. The notes below all regression tables indicate which of the covariate sets are used for the specific estimations.

Covariate Set 1: log population, the share of age group 30-39, the share of white individuals, the share of individuals with university degree, the Gini coefficient, the share of immigrants (from after 2010), binary variable whether democrats have majority support in the county, and a measure of collective efficacy.

Covariate Set 2: log population, the share of age group 30–39, the share of white individuals, the share of Asian individuals, the share of married individuals, the share with a university degree, log per capita income, the Gini coefficient, the share of immigrants (from 1990–1999), the share of immigrants (from 2010 or later), a binary variable indicating whether Democrats won, a measure of collective efficacy, a measure of charitable organizations, and population density.

Covariate Set 3: log population, the share of age group 30–39, the share of white individuals, the share of married individuals, the share with a high school degree, the share with a university degree, the Gini coefficient, the share of immigrants (from 2010 or later), a binary variable indicating whether Democrats won, and a measure of collective efficacy.

Covariate Set 4: log population, the share of age group 30–39, the share of age group 50–59, the share of white individuals, the share with a university degree, the Gini coefficient, the share of immigrants (from 2010 or later), the share of immigrants (from before 1990), the share of individuals without health insurance, a binary variable indicating whether Democrats won, a measure of political fractionalization, a measure of institutional health, and a measure of collective efficacy.

Covariate Set 5: log population, the share of age group 30–39, the share of white individuals, the share of Black individuals, the share with a university degree, the Gini coefficient, the share of immigrants (from 2010 or later), the share of immigrants (from before 1990), the share of individuals without health insurance, a binary variable indicating whether Democrats won, a measure of political fractionalization, a measure of institutional health, and a measure of collective efficacy.

D. Binary Treatment

D.1. Overlap Assumption

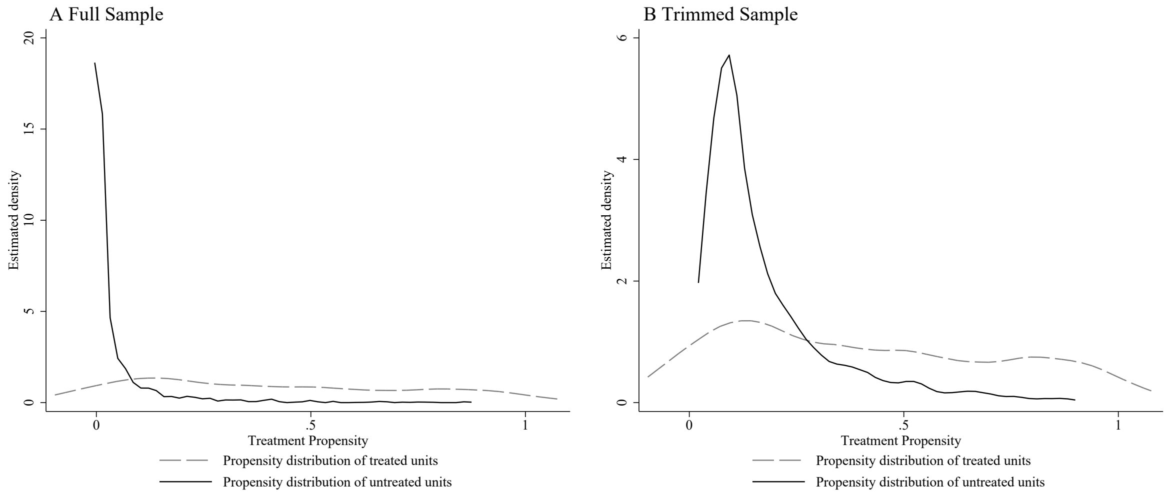

Figure 7: Smoothed density of the propensity score for the treatment and control group. B presents a trimmed sample in which all propensity scores below 0.05 were dropped for the control group to make the overlap of the remaining units visible. The smoothing is conducted relying on the Epanechnikov kernel.

D.2. Estimator Description

(1) A regression adjustment (RA) estimator is the first assessment of the target parameter from Equation 17 (Wooldridge, 2010). The estimator is depicted in Equation A.2 where w_i is a dummy equal to 1 for all treated units and $\hat{m}$ is a linear model predicting the unobserved outcome. The estimator computes the average difference between the treated outcome and the hypothetical untreated outcome of all treated units in the sample. The hypothetical outcome is a predicted value from the linear model.

$$\hat{\tau}_{\mathrm{RA}} = \left(\sum_{i=1}^{N} W_i \right)^{-1} \left\{ \sum_{i=1}^{N} W_i \left[\hat{\boldsymbol{m}}_1 \left(\mathbf{X}_i \right) - \hat{\boldsymbol{m}}_0 \left(\mathbf{X}_i \right) \right] \right\} \tag{A.2}$$

(2) The second static estimator is an inverse probability weighting (IPW) estimator that weights observations differently depending on how likely it is that the observation had the opposite treatment status. The estimator is presented in Equation A.3 (Wooldridge, 2010). Observations closer to the counterfactual group are weighted more strongly in the estimation. In the first stage, the propensity score function $P(X)$ is obtained using a logit model. In the second stage, the observations are weighted by the inverse of the propensity score. Given the obtained inverse propensity scores, the ATET is estimated according to Equation A.3.

$$\widehat{\tau}_{\mathrm{IPW}} = N^{-1} \sum_{i=1}^{N} \frac{[W_i - P(X_i)]\, Y_i}{\frac{N_1}{N}\,[1 - P(X_i)]} \tag{A.3}$$

(3) Following Wooldridge (2007), inverse probability-weighting can be enhanced by using a regression adjustment technique, making the estimator doubly-robust. In the procedure, the inverse probability weights are computed using a probit model as for IPW estimation. Given the inverse-probability weights, a weighted regression is fit for each treatment status as in the RA estimation. This regression is used to obtain the hypothetical outcomes, whose average differences are the ATE, and the average differences among the treated are the ATET of interest. The inverse probability-weighting regression adjustment (IPWRA) estimator thus combines the estimators from Equation A.2 and A.3.

(4) A fourth estimation strategy is propensity score matching (PSM) depicted in Equation A.4 and A.5. In the first step, a propensity score for the likelihood of treatment is calculated using a probit model and the selected covariates, resulting in $P(X)$. Conditioning on this probability is equivalent to controlling for the covariates themselves (Rosenbaum and Rubin, 1983). Instead of controlling for the propensity score in an OLS regression, a matching estimator is applied. This matching estimator matches the three nearest untreated observations in terms of the difference in the propensity score to a treated observation and compares the differences in the outcome variable.

$$\widehat{\tau}_{\mathrm{PSM}} = \frac{1}{n_1} \sum_{i=1}^{n} W_i \left(Y_i - Y_{j(i)}\right) \tag{A.4}$$

where

$$Y_{j(i)} = \sum_{j:D_j=0} I\left(|P(X_j) - P(X_i)| = \min_{l:D_l=0} |P(X_l) - P(X_i)|\right) Y_j \tag{A.5}$$

(5) A final estimation strategy is direct covariate matching (COM). Instead of using the propensity score, all covariates are used to find the three nearest untreated neighbors of each treated observation. Importantly, all covariates are standardized. The nearest units are determined by the Mahalanobis distance between observations – this distance measure accounts for all differences between the standardized selected covariates while controlling for the correlation among them (Huber, 2023). The ATET estimator is formally depicted in Equation A.6 and A.7 where $\|X_l - X_i\|$ denotes the Mahalanobis distance between observation l and i.

$$\widehat{\tau}_{\mathrm{COM}} = \frac{1}{n_1} \sum_{i=1}^{n} W_i \left(Y_i - Y_{j(i)}\right) \tag{A.6}$$

where

$$Y_{j(i)} = \sum_{j:D_j=0} I\left(\|X_j - X_i\| = \min_{l:D_l=0} \|X_l - X_i\|\right) Y_j \tag{A.7}$$

D.3. Covariate Balance and Placebo Tests

Table 8: Covariate Balance and Placebo Tests

SMD	RA	IPW	IPW RA	C. Match	P.S. Match
Log Population		0.119	0.033	0.033	-0.045
Age Group 30-39		0.129	0.054	0.054	-0.055
Share White Ind.		-0.072	-0.026	-0.026	-0.047
Share Uni. Degree		0.057	0.001	0.001	-0.059
Gini Coefficient		0.166	0.138	0.138	0.184
Share Immig. (a. 2010)		0.012	-0.039	-0.039	-0.119
Democrat Majority		0.097	0.047	0.047	0.036
Collective Efficacy		-0.209	-0.201	-0.201	-0.225
Overidentification		0.8011	0.8943		
Placebo Outcome	14.123***	10.223*	8.501	12.518**	9.867*
(2021)	(5.125)	(5.356)	(5.243)	(5.180)	(5.137)
Placebo Outcome	21.433***	10.978	6.407	17.186***	6.893
(2020)	(5.402)	(6.974)	(6.660)	(5.377)	(7.015)
Placebo Outcome	0.654**	0.317	0.185	0.563*	0.102
(2019)	(0.325)	(0.356)	(0.374)	(0.327)	(0.327)
Observations	3007	3007	3007	3007	3007

Heteroskedasticity robust SE in parentheses; * p <0.1, ** p <0.05, *** p <0.01

Note: Covariate balance and placebo tests. The first rows present the standardized mean differences (SMD) of the control vs. treatment group after weighting and matching. Covariate balance is also tested using the overidentification test proposed by Imai and Ratkovic (2014) for weighting estimators. For the placebo tests, the protest outcome variable of the baseline specification is replaced with values for 2021, 2020, and 2019. We attribute the residual statistical significance to the strong persistence of protest activity over time. Specifically, the number of BLM and anti-police protests in 2022 is highly correlated with previous years: 0.88 with 2021, 0.69 with 2020, and 0.62 with 2019.

E. Sensitivity Analysis

E.1. Robustness Static Model

Table 9: Robustness Binary Treatment Evaluation: Protest Outcome

	RA	IPW	IPW RA	C. Match	P.S. Match
Trimmed Sample	11.847***	10.905**	9.042**	12.249***	10.232**
(75 mi spillovers)	(4.302)	(4.394)	(4.405)	(4.280)	(4.411)
Observations	1345	1345	1345	1345	1345
Trimmed Sample	11.649***	11.177**	9.347**	12.586***	10.981**
(100 mi spillovers)	(4.338)	(4.425)	(4.653)	(4.286)	(4.346)
Observations	966	966	966	966	966
Trimmed Sample	5.036	5.985	5.278	8.163**	6.060
(overlap)	(4.137)	(3.750)	(3.891)	(3.602)	(3.882)
Observations	714	714	714	714	714

Heteroskedasticity robust SE in parentheses; * p <0.1, ** p <0.05, *** p <0.01

Note: Robustness estimation results of the Average Treatment Effect on the Treated (ATET) from the context shock on protest events for the year 2022. All estimations are based on Covariate Set 1.

Table 10: Robustness Binary Treatment Evaluation: Salience Level

	RA	IPW	IPW RA	C. Match	P.S. Match
Inverse hyper. Y	0.043**	0.072***	0.074***	0.066***	0.086***
	(0.017)	(0.019)	(0.020)	(0.020)	(0.027)
Observations	2946	2946	2946	2946	2946
Trimmed Sample	3.480***	6.278***	6.497***	4.470***	8.129***
(75 mi spillovers)	(1.321)	(1.430)	(1.477)	(1.209)	(1.688)
Observations	1232	1232	1232	1232	1232
Trimmed Sample	3.281**	6.054***	5.335**	5.055***	7.806***
(100 mi spillovers)	(1.581)	(1.980)	(2.157)	(1.213)	(1.292)
Observations	859	859	859	859	859
Trimmed Sample	4.489***	4.032***	3.936***	2.888***	3.659***
(overlap)	(1.001)	(0.967)	(0.957)	(0.891)	(1.154)
Observations	683	683	683	683	683
Higher Level	5.363***	6.857***	5.922***	4.346***	5.886***
	(1.551)	(1.870)	(1.379)	(1.414)	(1.087)
Observations	201	201	201	201	201

Heteroskedasticity robust SE in parentheses; * p <0.1, ** p <0.05, *** p <0.01

Note: Robustness estimation results of the Average Treatment Effect on the Treated (ATET) from the context shock on salience for the year 2022. For the inverse hyperbolic Y Covariate Set 5 is applied. All other estimations are based on Covariate Set 4.

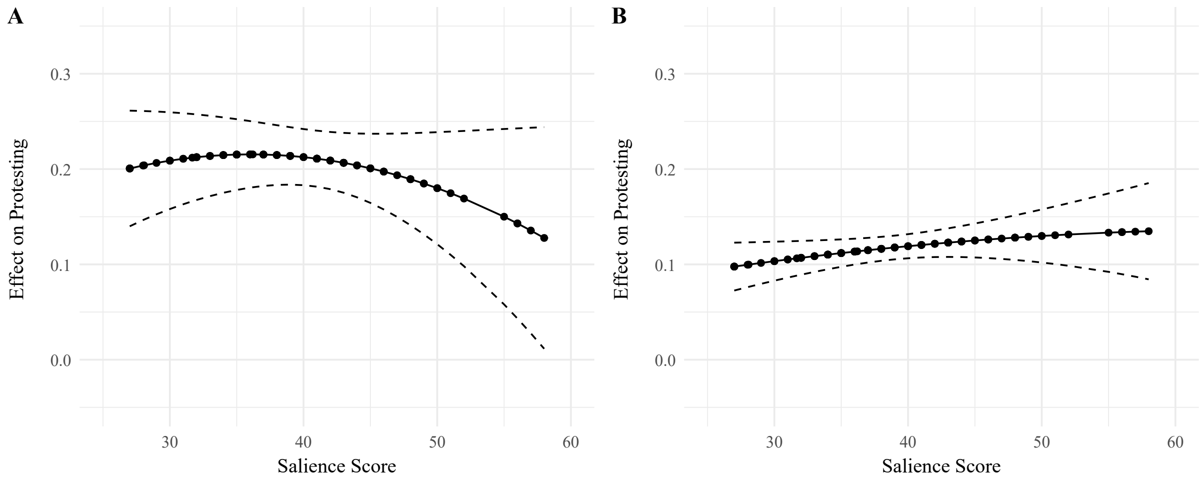

Figure 8: Robustness: Results of continuous treatment of salience on protest outcome. The dashed lines indicate the 95% confidence interval. Graph A depicts the fuzzy treatment effect relying on the inverse hyperbolic sine transformed protest variable as the outcome, while Graph B illustrates the fuzzy effect for a binary protest variable as the outcome. LASSO selection Covariate Set 2 is applied.

E.2. Distribution Salience

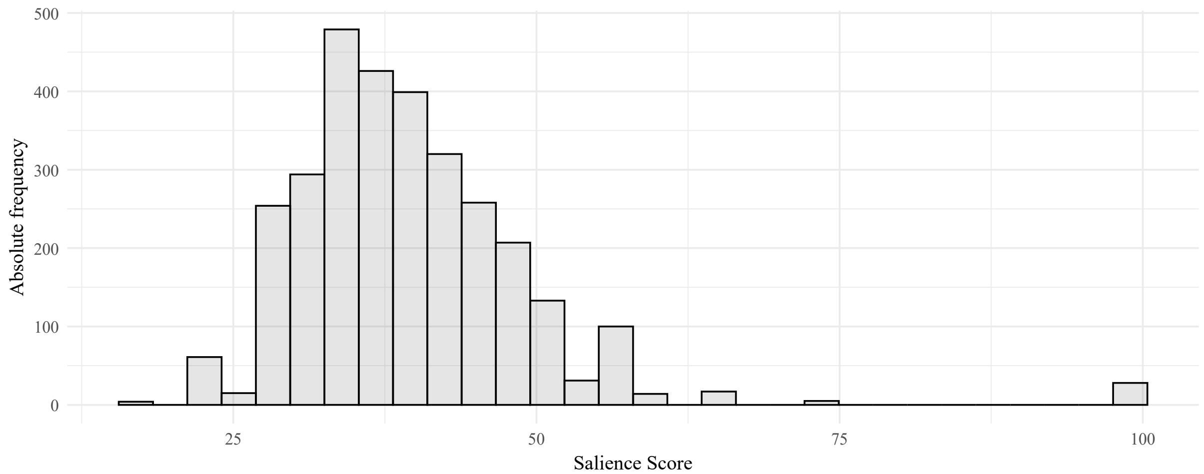

Figure 9: Distribution salience scores in 2022 based on Google Trends data.

E.3. Robustness Year Variation

Table 11: Context Dependence: Context Shocks on Protest Outcomes

	RA	IPW	IPW RA	C. Match	P.S. Match
Dependent Variable: Protest Outcome					
Context Shock 2018	0.620*	0.557*	0.517	0.526	0.496
	(0.327)	(0.331)	(0.334)	(0.328)	(0.337)
Observations	3014	3014	3014	3014	3014
Context Shock 2020	25.019***	20.213***	16.140***	21.920***	18.146***
	(6.075)	(6.269)	(5.890)	(5.883)	(7.028)
Observations	3013	3013	3013	3013	3013
Context Shock 2022	10.463**	8.768*	7.913*	9.044**	9.728**
	(4.364)	(4.474)	(4.560)	(4.525)	(4.292)
Observations	3007	3007	3007	3007	3007
Dependent Variable: Salience Level					
Context Shock 2018	0.130*	0.225***	0.240***	0.246***	0.276***
	(0.076)	(0.079)	(0.081)	(0.076)	(0.086)
Observations	2926	2926	2926	2926	2926
Context Shock 2020	0.100	0.148**	0.146*	0.114	0.139
	(0.069)	(0.072)	(0.081)	(0.070)	(0.091)
Observations	2934	2934	2934	2934	2934
Context Shock 2022	0.157***	0.172***	0.196***	0.176***	0.154***
	(0.059)	(0.064)	(0.061)	(0.058)	(0.050)
Observations	2907	2907	2907	2907	2907

Heteroskedasticity robust SE in parentheses; * p <0.1, ** p <0.05, *** p <0.01

Note: Estimation results of the Average Treatment Effect on the Treated (ATET) from the context shock on protest events and salience for the year 2022. LASSO selection Covariate Set 1 is applied for the first 3 sets of estimators and Covariate Set 4 for the latter 3 sets.

Table 12: Robustness Year Variation

	RA	IPW	IPW RA	C. Match	P.S. Match
Shooting 2019	0.792**	0.509	0.399	0.760**	0.439
	(0.367)	(0.396)	(0.367)	(0.360)	(0.384)
Observations	3014	3014	3014	3014	3014
Shooting 2021	13.096**	3.361	1.400	12.694**	0.525
	(5.474)	(9.093)	(7.871)	(5.089)	(8.786)
Observations	3013	3013	3013	3013	3013
Inverse hyper. Y 2018	0.182***	0.137**	0.073	0.149**	0.093
	(0.057)	(0.065)	(0.066)	(0.061)	(0.097)
Observations	2974	2974	2974	2974	2974
Binary Y 2018	0.115***	0.087**	0.066	0.082**	0.074*
	(0.031)	(0.038)	(0.040)	(0.035)	(0.041)
Observations	3014	3014	3014	3014	3014
Trimmed Sample 2018	0.656**	0.682**	0.543	0.572*	0.470
(75 mi spillovers)	(0.329)	(0.330)	(0.336)	(0.340)	(0.372)
Observations	1351	1351	1351	1351	1351
Trimmed Sample 2018	0.688**	0.691**	0.548	0.633*	0.711**
(100 mi spillovers)	(0.330)	(0.331)	(0.336)	(0.331)	(0.339)
Observations	987	987	987	987	987
Trimmed Sample 2018	0.259	0.211	0.218	0.228	0.165
(overlap)	(0.158)	(0.183)	(0.160)	(0.155)	(0.148)
Observations	683	683	683	683	683
Inverse hyper. Y 2020	0.518***	0.589***	0.187***	0.682***	0.301***
	(0.079)	(0.112)	(0.068)	(0.103)	(0.115)
Observations	2974	2974	2974	2974	2974
Binary Y 2020	-0.154***	-0.009	-0.056**	0.019	-0.000
	(0.028)	(0.020)	(0.023)	(0.028)	(0.019)
Observations	3013	3013	3013	3013	3013
Trimmed Sample 2020	27.703***	23.465***	16.622***	25.885***	25.053***
(75 mi spillovers)	(6.214)	(6.474)	(6.553)	(6.271)	(8.299)
Observations	1427	1427	1427	1427	1427
Trimmed Sample 2020	29.194***	24.224***	25.130***	28.286***	28.477***
(100 mi spillovers)	(6.214)	(6.474)	(6.137)	(6.271)	(8.299)
Observations	1427	1427	1427	1427	1427
Trimmed Sample 2020	15.269**	15.336**	13.739**	21.135***	13.056*
(overlap)	(6.177)	(6.443)	(5.757)	(5.814)	(7.571)
Observations	674	674	674	674	674

Heteroskedasticity robust SE in parentheses; * p <0.1, ** p <0.05, *** p <0.01

Note: Robustness estimation results of the Average Treatment Effect on the Treated (ATET) from the context shock on protest events for the years 2018, 2019, 2020, and 2021. For the inverse hyperbolic Y specifications, Covariate Set 2 and for the binary Y specifications, Covariate Set 3 are applied. All other estimations are based on Covariate Set 1.

E.4. Robustness Salience Effect Year Variation

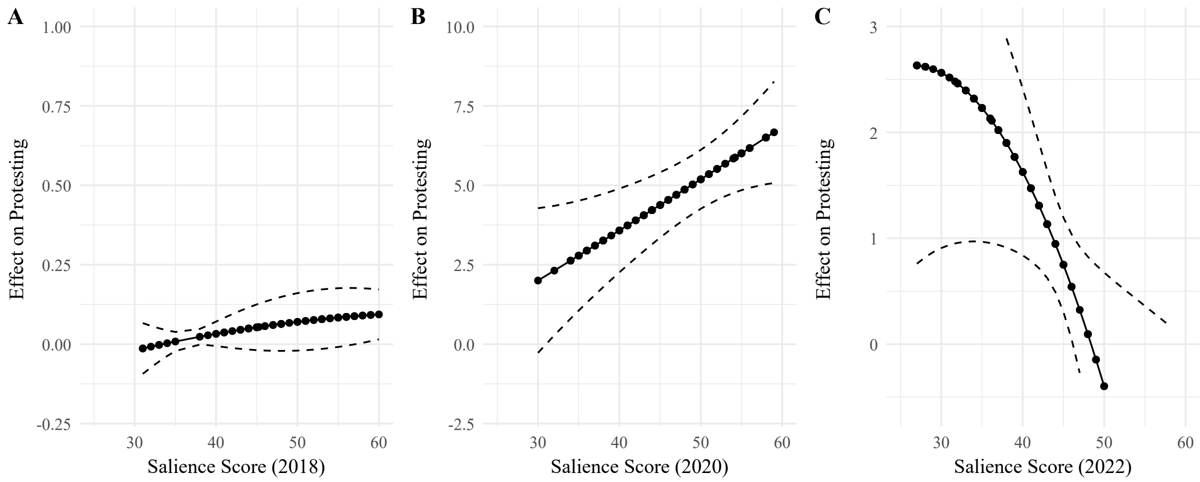

Figure 10: Estimation results fuzzy treatment of salience on protest event for the years 2018, 2020, and 2022. The dashed lines indicate the 95% confidence interval. LASSO selection Covariate Set 2 is applied.