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
University of Bayreuth
Faculty of Law, Business and Economics
henrik.guhling@uni-bayreuth.de

This version: May 5, 2025

Abstract: This paper introduces a novel attentional mechanism for protest mobilization, building on salience theory developed by Bordalo *et al.* (2013, 2020b). While existing models typically emphasize thresholds, networks, or informational factors, our model focuses on how high-profile events, i.e., fatal police encounters involving Black individuals, influence protest behavior. We examine the impact of such events on Black Lives Matter and anti-police violence mobilization, employing treatment evaluation and panel data techniques. Our results show that high-profile events significantly increase protest mobilization, with a substantial portion of this effect mediated by changes in salience. This framework enhances our understanding of protest dynamics, particularly in the immediate aftermath of major events.

Keywords: Protests; Salience Theory; Attention; Critical Events; Black Lives Matter.

JEL classification: D71; D74; D91.

Acknowledgments: The authors thank the participants of the following research conferences: Public Choice Conference 2025, Bayreuth Graduate Seminar 2024 and 2025, and 38th BGPE Research Workshop. The authors thankfully acknowledge helpful suggestions and feedback from Martin Leschke and Fabian Herweg. 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, sometimes referred to as critical moments, have the potential to trigger instant, far-reaching public uprisings in the form of protests. For instance, the murder of George Floyd provoked large-scale protests against police violence and racial discrimination. In the literature, such mobilization dynamics are commonly modeled relying on threshold models (e.g., Granovetter, 1978; Yin, 1998), informational cascades (e.g., Lohmann, 1994; Little, 2016), or models of strategic interaction (e.g., De Mesquita, 2010; Shadmehr, 2018). In this paper, we propose a novel attention mechanism to explain the effect of such high-profile events on protest mobilization. The proposed mechanism can explain protest timing, mobilization size, and contextual dependencies. We assess a set of theory-derived hypotheses by empirically investigating mobilization dynamics in the context of local police violence and Black Lives Matter (BLM) and anti-police violence protests in the US from 2017 to 2022.

The proposed mechanism starts with a context shock (critical event) that increases the salience of protesting, which in turn increases the number of protests. We claim the mechanism to have substantial explanatory power for rapid protest mobilization patterns after far-reaching events. Such events include, among others, the murder of George Floyd, which led to sustained anti-police violence protests, the Chinese extradition plans, which led to mass protests in Hong Kong, or the self-immolation of Mohamed Bouazizi, which started the Arab Spring[1]. While these cases powerfully illustrate the mobilization dynamics of interest, their empirical investigation faces a number of obstacles. First, there is no valid control group (no second Hong Kong) where no such event happened. Therefore, we can hardly make any comparisons between the entity with and without a critical event. Second, such critical events are defined ex-post, and thus all events we would assess have already survived a selection process. In this process, these events happened to be the critical ones because they sparked large protests in their aftermath, but ex-ante, it is far from obvious and hardly predictable which events turn out that way.

To overcome these econometric difficulties, we shift the analysis to the local level in the US and define fatal police encounters of black individuals as local context shocks.

[1] There are a number of excellent studies investigating protest dynamics in these context: among others, Moreno-Medina *et al.* (2025) investigate biased media attention of 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 reasons (Malik and Awadallah, 2013) and the role of social media (Khondker, 2011).

Especially in the context after George Floyd, such events received increased media attention and, thus, changed the context in which individual protest decisions are made. The investigation on 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 relying on comparisons between 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.

Based on Bordalo *et al.* (2013), we propose a binary choice model over the decision to protest versus not to protest. An individual decides based on two attributes of protesting: the benefits (e.g., expressive utility, network effects) and costs (e.g., opportunity cost, potential violence). To analyze the effect of critical events on the individual decision, we incorporate context by assigning weights to both attributes that depend on each attribute's salience level. Salience is modeled as the distortion between the actual and reference benefits. This reference benefit is formed based on memories following Bordalo *et al.* (2020a). A critical event is theorized to affect the reference benefit in a way that the salience of the benefit increases, indirectly leading to higher weights for the benefit. The rationale behind this mechanism is that protest decisions are not made in a vacuum, but are strongly affected by the personal attentional environment. These environments are shaped by someone's exposure to, among other things, social media content, news coverage, and personal conversations. These aspects are captured in the memory element of the model and eventually affect the decision weights. A crucial insight derived from the model is that even though neither the benefits nor the costs change qualitatively, a context shock due to a critical event can shift individuals' decisions to protest.

The proposed model yields multiple new predictions about protest mobilization dynamics. The first is the mechanism from crucial events to enhanced protest mobilization via salience. This mechanism highlights the importance of contextual and attentional features in individual decision-making. Furthermore, the response of salience to changes in the context is shown to have a concave curvature for increasing salience levels. This feature makes the model sensitive to pre-shock attention levels. Finally, the model predicts demobilization as context shocks fade over time.

We test the model's predictions on the US county level for 2022; for some specifications, we extend the analysis period to 2017 through 2022. We apply treatment evaluation methods (regression adjustment, propensity score methods, and matching estimators) to obtain the effects of context shocks on salience levels and protest outcomes. Conditional independence is established using LASSO double-selection, and robustness for spillovers is tested using sample trimming. We rely on Google Trends data to measure salience. To assess the dynamic predictions of the model, we furthermore apply a distributed lag model and staggered difference-in-differences estimator proposed by De Chaisemartin and d'Haultfoeuille (2024). Finally, we run an extension in which we explore the drivers of the size of context shocks.

The empirical evidence provides robust evidence for the major propositions derived from the model. Fatal police encounters of black individuals, interpreted as context shocks, are found to increase protest mobilization for BLM and anti-police violence demonstrations compared to a comparison group. The major part 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 became much larger, indicating that the nationwide context shock interferes 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 seems to have substantial explanatory power for short-run mobilization processes, while the persistence of protest activity over time likely hinges on alternative mechanisms beyond attention.

The paper's structure is the following. After discussing the paper's contribution within 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. A final discussion about what determines the size of context shocks is provided in Section 5.

Contribution and Related Literature The paper contributes to different strands of the literature. First, we add a novel attention mechanism to a literature that attempts to explain protests through economic theory, typically focusing on thresholds, networks, or information. Threshold models typically posit that an individual's decision to protest depends on the participation of others, with benefits exceeding costs in high-participation equilibria and costs exceeding benefits in low-participation equilibria (Granovetter, 1978). In Yin (1998), individuals refrain from protesting until a critical point is reached, from which protesting becomes the rational choice. Such tipping points can potentially be reached by a small minority of "certain actors," which lead to widespread mobilization in the aftermath as modeled by Wiedermann *et al.* (2020). Our proposed attention mechanism fits into such models as one way of achieving thresholds via context shocks.

Another, related strand of literature primarily focuses on strategic interactions among individuals. Shadmehr (2018) models the strategic nature of protest decisions, finding that individuals act as strategic substitutes when goals are modest and as strategic complements when goals are ambitious. In games of strategic complementarity, De Mesquita (2010) argues that signals can be used to manipulate individual beliefs about the anti-government mood to solve the coordination problem. Building on our proposed channel, we argue that such signals need not necessarily be sent by other individuals, but can emerge from increased attention due to a critical event.

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. In her seminal work, 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 and specifically models 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 as they affect an individual's context under which a decision is made. However, we believe the attention mechanism to be distinct from the information mechanism, as our definition of context shocks does not affect individuals' conscious beliefs but subconscious perception.

Furthermore, this paper contributes to the empirical literature of 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) derive a similar effect of the VK platform in Russia on protests and their size. However, 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. We argue that, beyond its informational role, social media content significantly shapes the perceptual context in which individuals decide to protest, a dynamic that is closely captured in our attention mechanism model.

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) examines 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 towards the stereotype of that group. Our application of salience theory to protesting is novel in precisely measuring the context and salience levels. This allows us to examine both context and salience through field data by testing hypotheses derived from the salience model.

2. Salience Model and Focal Points

In our model, we view protests as goods consisting of two attributes, i.e.,

$$P = (b \geq 0, c \geq 0) \tag{1}$$

where b is the potential benefit of the protest, and c is the participation costs. Both benefits and costs are very generally defined terms. Benefits encompass all aspects that make participation more attractive, including, among others, the number of participants, the perceived probability of success, or just the benefit of being in a group with the same interest and fighting for an overriding goal. Costs, on the other hand, capture disadvantageous aspects, e.g., travel or opportunity costs, potential risks of violent encounters with the police, as well as being in unfamiliar surroundings with unfamiliar people.

We consider the decision problem of an individual as a binary choice. Therefore, the choice set is given by

$$\mathbb{C} = \{(0,0); (b,c)\} \tag{2}$$

consisting in the protest P and an outside option $\underline{P} = (0,0)$.

The standard model evaluates the protest under consideration according to

$$U(b, c) = b - c \tag{3}$$

that is, an individual joins a movement if the benefits exceed the costs. For constant b and c, people either protest ($b \geq c$) or not ($b < c$). In order to incorporate the occurrence of high media-coverage events, we rely on a model in which the different attributes of a good can be weighted differently, i.e.,

$$U^w(b, c) = w_1 \cdot b - w_2 \cdot c \tag{4}$$

By redefining the utility function in this manner, it is not necessarily the costs and benefits that change the protest outcome but a shift in attribute weights, which may be influenced by changes in the political context. We argue that high media coverage events impact the attribute weights and, consequently, affect protest outcomes without altering the attributes themselves.

2.1. Salience and Memory Weights

The idea of salience as described by Bordalo *et al.* (2012, 2013) is that individuals focus on the options, attributes, or outcomes that attract most of their attention, i.e., the most salient ones. Thus, the weight of an attribute depends on the salience modeled by a salience function σ (Bordalo *et al.*, 2013).

Definition 1 (Salience function)**.** *For an attribute a, 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 $\epsilon + \epsilon' > 0$ we have*

$$\sigma(a + \mu\epsilon, \overline{a} - \mu\epsilon') > \sigma(a, \overline{a}) \tag{5}$$

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

$$\sigma(a + \epsilon, \overline{a} + \epsilon) < \sigma(a, \overline{a}) \tag{6}$$

The ordering property implies that the larger the difference between the salience functions' arguments, the more salient the attribute becomes. For example, $\sigma(10, 20) < \sigma(10, 30)$ holds as the difference between 10 and 30 is larger than between 10 and 20. The diminishing sensitivity property includes Weber's law, which states that the higher the level of attributes, the less salient the difference is. For example, $\sigma(10, 20) > \sigma(110, 120)$ holds since the percentage deviation on the left-hand side is larger than on the right-hand side, even though the absolute difference is the same.

The second component in the salience function takes the attribute of the reference good, which is given by

$$\overline{P} = (\bar{b}, \bar{c}) \tag{7}$$

Since salience theory considers decision problems as context-dependent choices, the reference good establishes this context. Consequently, attributes are typically assessed as deviations from this reference.

We propose that high-attentional events act as shocks that alter the decision-making context, thereby increasing the salience of the benefit attribute of protests. Consequently, some individuals may protest even if the perceived benefit is smaller than the cost, i.e., $b < c$. This implies that the attention spent on the protest benefit increases while the real good remains unchanged. Therefore, it is essential to define the reference good in a context-dependent manner.

We argue that the perception of a protest is influenced by previous experiences, encompassing all information received about the protest movement through social media, traditional news, conversations with friends, or other sources. For instance, reading a headline in the local newspaper constitutes an experience. If it is the first time we encounter a protest topic, we have no prior reference; this initial experience forms our reference. As we encounter similar headlines over time, the experiences do not significantly change, and the reference good remains stable. However, when a high-profile event occurs, two things happen simultaneously. The amount of information one receives about the topic increases drastically, and the problems and political demands are more explicitly formulated, such that its urgency is highlighted. Consequently, previous moderate experiences lose their influence in forming the reference good as they become too dissimilar from the current situation. This results in a change in the reference good, altering the salience of the attributes.

Our mathematical definition of experiences is based on Bordalo *et al.* (2020a). They also shape a norm regarding one particular good based on previous experiences. An experience at time t is defined as

$$e_t = (b_t, c_t, k_t) \in \mathbb{R}_0^+ \times \mathbb{R}_0^+ \times \mathbb{R}_0^+ \tag{8}$$

where b_t and c_t are the perceived benefits and costs of the protest, respectively. These perceived attributes do not necessarily have to equal the real attributes b and c. If $b_t = b$, individuals perceive the benefit as the true benefit. k_t reflects the context at time t.[2]

At first glance, context is a very general and abstract term. In our model, the context describes the situation in which we perceive an experience. For example, it can make a difference to read the same headline on the front page of a newspaper or a later page. In the same way, the context can also reflect the media coverage of the protest-related topic. The situation in which a certain event is reported can vary, i.e. if it is reported very frequently, the topic may appear more important to us. In addition, the reporting media outlet can provide a different context. Furthermore, the context can also represent whether we are personally affected or emotionally charged by an event at a certain time.

As mentioned earlier, the reference good is a weighted average of past experiences; the weight of a particular experience is given by the experience's similarity to the current experience at time t. Thus, more similar experiences are weighted more heavily. The reference good could also be seen as the expectation regarding the protest at time t. When the expectation for a particular attribute significantly exceeds its actual value, it receives much more attention, thereby influencing the decision more strongly.

We define similarity like Bordalo *et al.* (2020a) regarding the Euclidean distance.

Definition 2 (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:*

$$S(e_s, e_t) = S_1(|b_t - b_s|)S_2(|c_t - c_s|)S_3(|k_t - k_s|) \tag{9}$$

where $S_k : \mathbb{R}^0_+ \to \mathbb{R}_+$ 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.,

$$S(e_s, e_t) = exp(-[|b_t - b_s| + |c_t - c_s| + |k_t - k_s|]) \tag{10}$$

such that it takes values between 0 and 1. Therefore, it is 1 if e_s and e_t have the same parameters and go to 0 if they get very dissimilar in one or more dimensions.

[2]In general, we assume that different situations can be compared with each other so that an order relation exists. Therefore, the context k_t describes this comparability by assigning different numbers to different situations. Initially, it does not matter whether a number is high or low. However, the smaller the distance between two numbers, the more similar the situations should be.

Using the concept of similarity introduced above, we can determine the weight given to a particular past experience when molding the reference good.

Definition 3 (Memory weights and reference good)**.** *The memory weight of experience* $e_s = (b_s, c_s, k_s)$, $s \leq t$, *at time* t *is given by*

$$w(e_s, e_t) = \frac{S(e_s, e_t)}{\int S(e_u, e_t) dF}, \quad u \leq t \tag{11}$$

where F *is the distribution function of all experiences until time* t*. Then, for the reference good* $\overline{P} = (\bar{b}, \bar{c})$ *it holds*

$$\bar{b}(e_t) = \int b_u \cdot w(e_u, e_t) dF, \quad \bar{c}(e_t) = \int c_u \cdot w(e_u, e_t) dF, \quad u \leq t \tag{12}$$

Thus, the attributes of the reference good are determined by similarity-weighted averages of already experienced benefits or costs, respectively. Using this definition, we can define the utility resulting from the salience model based on Bordalo *et al.* (2013).

Definition 4 (Utility in the salience model)**.** *In the salience model, a protest* $P = (b, c)$ *is evaluated according to*

$$U^s(b, c) = \frac{2 \cdot \delta^{-\sigma(b, \bar{b}(e_t))}}{\delta^{-\sigma(b, \bar{b}(e_t))} + \delta^{-\sigma(c, \bar{c}(e_t))}} \cdot b - \frac{2 \cdot \delta^{-\sigma(c, \bar{c}(e_t))}}{\delta^{-\sigma(b, \bar{b}(e_t))} + \delta^{-\sigma(c, \bar{c}(e_t))}} \cdot c \tag{13}$$

where $\delta \in (0, 1]$ *measures the degree of salient thinking.*

[3]This approach is also closely linked to the proposal of Bordalo *et al.* (2012) for continuous weights. We retain δ because it describes an individual in terms of how salient they think, i.e., how strongly they focus on the more salient attribute. Using the δ-parameter, we can state which types go protesting and whether the number of protesters changes for different salience levels.

Considering the model, we can see that for $\delta = 1$, both attributes receive the weight 1 such that we get the standard model. This is also the case if both attributes are equally salient. That is why an individual of type $\delta = 1$ has no salience bias in preferences and is, thus, no salient thinker. For $\delta \to 0$, an individual only focuses on the most salient attribute.

[3]Please note that we have adapted the rank-based salience approach of Bordalo *et al.* (2013) so that the attribute weights are continuous to be able to make statements not only about which attribute is more salient but also about whether one is clearly more salient than the other or only a little. If the salience of one attribute is set to 0 and for the other to 1, we get the rank-based model.

2.2. Context Shocks and Reference Good

With the terms introduced above, we can define context shocks and how they impact attribute weights. From a general point of view, all high-profile events shock the context with regard to some entity. For our purposes, we are interested in events that shock the context vis-à-vis protest decisions. A large context shock was the murder of George Floyd with respect to attitudes towards anti-police and BLM protests. Other cases of police violence since then also received considerable media attention. We, thus, refer to such incidents as context shocks.

Mathematically, we say that a context shock occurs when the context reaches a level not previously achieved in a certain period of time.

Definition 5 (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:*

$$\max_{u,w \in [s,t)} |k_u - k_w| < \max_{u,w \in [s,t]} |k_u - k_w| \tag{14}$$

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 to the interval. In the simple case of 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.

As mentioned above, we say that the context influences how individuals perceive the benefit of a protest. Thus, we assume that the experienced benefit b_t at time t depends on the context k_t.

Assumption 1. *The perceived benefit at time t, b_t, depends on the context k_t, such that $b_t = f(k_t)$ with continuous f. Furthermore, we assume $b_t \geq b$. 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.*

We could also make the costs context-dependent while fixing the perception bias of the benefit. Since this adjustment would not alter the qualitative outcomes, we focus on changes in the benefit for the sake of tractability. Furthermore, we could also assume $b_t \leq b$. However, we want to exclude the cases in which the benefit is first salient by being strongly underestimated and then becomes salient due to a strong overestimation, and vice versa. It seems more realistic to assume either that a benefit is strongly underestimated and then weighted more heavily due to a positive surprise or, conversely, that the benefit is assumed to be higher and is therefore overweighed.

Since by Assumption 1, the reference benefit at time t only depends on k_t, we will write $\bar{b}(k_t)$ in the following. Furthermore, we write $\bar{c}$ as the reference costs are fixed.

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 reference benefits.

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., $\bar{b}(k_t) > (<) \bar{b}(k_{t-1})$, if and only if the following inequality is fulfilled:*

$$\sum_{s=0}^{t} S(e_s, e_t) \cdot [b_s - \bar{b}(k_{t-1})] > (<) 0 \tag{15}$$

For injective $f(k_s) = b_s$ and $\alpha_s = \frac{S(e_s,e_t)}{S(e_s,e_{t-1})}$, 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

$$\alpha_i \geq \alpha_j \Leftrightarrow b_i \geq (\leq) b_j, \quad \forall i, j \in \{0, ..., t-1\} \tag{16}$$

For $\bar{b}(k_{t-1}) \leq (\geq) \frac{1}{t} \sum_{s=0}^{t} b_s$, inequality 15 is always fulfilled by Chebyshev's sum inequality. Furthermore, $\Delta_{\bar{b}} = \bar{b}(k_t) - \bar{b}(k_{t-1})$ strictly decreases in $\bar{b}(k_{t-1})$ for fixed $S(e_s, e_t)$, $s = \{0, ..., t\}$.

Inequality 15 is always fulfilled for a constant context before time t, so a global context shock here always results in a higher (lower) reference benefit. However, the reference benefit and thus also the salience are not changed if inequality 15 is fulfilled with equality.[4]In the following, we define salience-changing context shocks as salience shifts.

Definition 6 (Salience shift)**.** *A context shock on $[s, t)$ is called*

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

$$\sigma(b, \bar{b}(k_t)) > \max_{u \in [s,t)} \sigma(b, \bar{b}(k_u)) \tag{17}$$

[4]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 10 always fulfills condition 16 such that a higher b_t leads to an increased $\bar{b}$.

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

$$\sigma(b, \bar{b}(k_t)) < \min_{u \in [s,t)} \sigma(b, \bar{b}(k_u)) \tag{18}$$

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 emphasizing the disadvantages or costs of protests, this is termed a negative salience shift.

Figure 1 shows how a context shock can affect the reference good and, thus, the perception of the benefit. It becomes visible that with a constant context and perceived benefit, every change leads to a context shock. By Lemma 1, for an injective f, the reference benefit reaches a new higher level. As the perceived benefits before t_{shock} are constant and equal to the true benefit b, this context shock increases the benefits' salience such that we observe a positive salience shift. Furthermore, in this example, the context returns to its original level, i.e., it normalizes again so that the reference benefit also returns to its previous level.

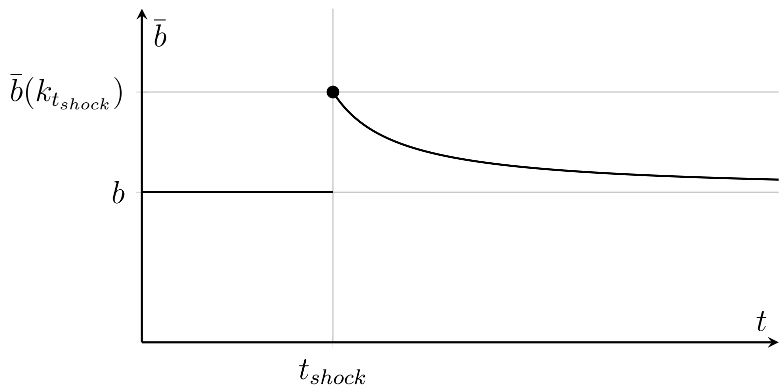

Figure 1: Shift in reference benefit $\bar{b}$ due to context shock in t_{shock}.

Finally, it is also worth taking a look at the extent to which a context shock has different effects at different previous levels of salience. The following proposition provides a statement about this.

Proposition 1 (Diminishing salience shifts)**.** *Consider discrete time with $s \in \{0, ..., t\}$. If $f(k_s) = b_s$ is injective, $x \mapsto \sigma(b, x)$ is a concave function and given b, for a global context shock k_t inducing a positive salience shift holds true that $\sigma(b, \bar{b}(k_t)) - \sigma(b, \bar{b}(k_{t-1}))$ gets smaller the higher $\bar{b}(k_{t-1})$.*

[5]The proposition, therefore, states that the higher the previous salience level, the smaller the increase in salience due to the same context shock. This statement is only true if the salience function σ is concave in its second argument. For a convex function, the statement can be reversed. In a later specification, we test this curvature behavior empirically.

2.3. Protest Participation

Applying the model to the decision problem with the choice set $\mathbb{C} = \{(0,0);(b,c)\}$, an individual protests if and only if

$$U^s(b,c) \geq U^s(0,0) = 0 \tag{19}$$

Rearranging that leads to

$$\delta^{\sigma(c,\bar{c})-\sigma(b,\bar{b}(k_t))} \geq \frac{c}{b} \tag{20}$$

As mentioned above, the model collapses to the standard model if the attributes are equally salient. Then, by considering equation 20, we see that for $c > b$, nobody protests irrespective of type δ, whereas everybody protests for $b \geq c$. These are the predictions of the standard model.

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 individuals participate. The next proposition gives an overview of protesting individuals based on the attributes' salience.

Proposition 2. *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 $\delta \leq \delta^*(k_t)$ where*

$$\delta^*(k_t) = \begin{cases} 0 & \textit{if } \sigma(b,\bar{b}(k_t)) \leq \sigma(c,\bar{c}), \\ \left(\frac{b}{c}\right)^{\frac{1}{\sigma(b,\bar{b}(k_t))-\sigma(c,\bar{c})}} \in (0,1) & \textit{if } \sigma(b,\bar{b}(k_t)) > \sigma(c,\bar{c}). \end{cases} \tag{21}$$

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 $\delta \geq \delta^*(k_t)$ where

$$\delta^*(k_t) = \begin{cases} 0 & \textit{if } \sigma(b,\bar{b}(k_t)) \geq \sigma(c,\bar{c}), \\ \left(\frac{c}{b}\right)^{\frac{1}{\sigma(c,\bar{c})-\sigma(b,\bar{b}(k_t))}} \in (0,1) & \textit{if } \sigma(b,\bar{b}(k_t)) < \sigma(c,\bar{c}). \end{cases} \tag{22}$$

[5]Even if the statement initially seems trivial due to the properties of salience functions, this is not the case. For example, convex salience functions may be constructed such as $\sigma(b,\bar{b}) = \frac{(\bar{b}-b)^2}{|b|+|\bar{b}|}$ for which the statement does not apply.

The proposition provides a threshold value δ^*, which depends on the attributes' salience and, thus, on the context at time t, k_t. Two trivial extreme cases arise by considering δ^*. First, no one will protest if the costs are greater than the benefits and more salient. This makes sense, as the focus is then on the higher costs. Secondly, if the benefits are greater than the costs and more salient, all individuals will protest regardless of the type δ. This also makes sense, as the focus is on the higher benefit.

In the other two cases, a proportion of individuals always protests, depending on the type δ. Suppose the costs outweigh the benefits, but the benefits are more salient. In that case, all individuals with a certain minimum level of salient thinking, i.e., those who focus enough on the more salient attribute, will protest. The more salient the benefit becomes, the greater the threshold value δ^* and the more individuals participate. If the benefit is greater than the costs, but the costs are more salient, all individuals who do not think too saliently, i.e., who do not focus too much on the more salient attribute, will protest. The smaller the salience difference between benefits and costs, the smaller the threshold value δ^* and the more individuals participate.

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* $s \in \{0, ..., t-1\}$.

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

(vi) *If* $\sigma(b, \bar{b}(k_{t-1})) - \sigma(c, \bar{c}) < 0$*, then every* ***positive salience shift*** *on* $[s, t)$ *leads to increased protest participation.*

The results show that not every context shock that changes salience necessarily affects protest participation. If the costs are greater than the benefits and are also salient, it takes a large increase in the benefits' salience for individuals to participate in the protest. The George Floyd incident, for example, can be seen as such a context shock. In our empirical analysis of Black Lives Matter protests, we will see that it needs a big shock like the George Floyd incident, such that smaller shocks like police shootings on the county level can impact protest participation.

Conversely, if the benefit is greater than the costs and is also salient, a strong decrease in the salience of the benefit is also required so that the costs become more salient again. Strong repressions in autocratic regimes like the Tiananmen massacre in China exemplify such powerful context shocks that lead to a sufficiently large negative salience shift. However, several smaller negative salience shifts can also lead to a decline in participation from a certain point onwards. For example, decreasing attention to the protest-related topic leads to the movements becoming smaller and eventually disappear if attention is not recaptured by a new context shock that leads to a positive salience shift.

Discussion of Protest Movements over Time Finally, we will discuss how our model can explain the life cycle of protest movements. Before protests arise, protest mobilization is generally very low hinting at the fact that the costs exceed the benefits, i.e., $c > b$ (and the costs are more salient). In our model, such a situation may arise even though $b > c$ applies, if the benefit's salience is sufficiently low, thus shifting all attention to the costs. In both cases, positive salience shifts can lead to increased participation. In the first case, the salience shift must be sufficiently large. Particularly large salience shifts, such as the George Floyd incident, are generally cited as trigger events for protest movements. However, the sum of several small shifts can also lead to the emergence of a protest movement, even if the movement then only builds up slowly. In many cases, however, a major trigger can be identified.[6] After such a large shock, several smaller shocks can still lead to an increased or constant protest participation over time. Over time, however, the attention will dissipate or shift to different topics leading to a decline in the salience level. With a falling salience level, the participation falls and returns to its original size

[6] Besides the George Floyd incident, there is, for example, the self-immolation of Mohamed Bouazizi, which heralded the Arab Spring, the shooting of Neda Agha-Soltan or the death of Mahsa Amini, which triggered the Green Movement or the Women's Rights Movement in Iran, respectively.

again. The protest "wave" finds its end, and only those individuals protest that have been motivated all along. Stronger negative salience shifts, such as repressions can accelerate this decreasing trend.

In principle, it is also possible for the ratio of b to c to change from $c > b$ to $b > c$ and back due to increasing network effects and a growing protest movement, for example. But here, too, the decrease in the benefits' salience leads to a drop in participation from a certain point onwards so that, in the end, only a few participants remain.

Figure 2 shows an example of the critical types' progression over time, $\delta^*(t)$, when $c > b$ and both attributes start with the same salience. In this case, a positive salience shift leads to a sudden increase in participation. Over time, however, the negative salience effects predominate, so participation decreases.

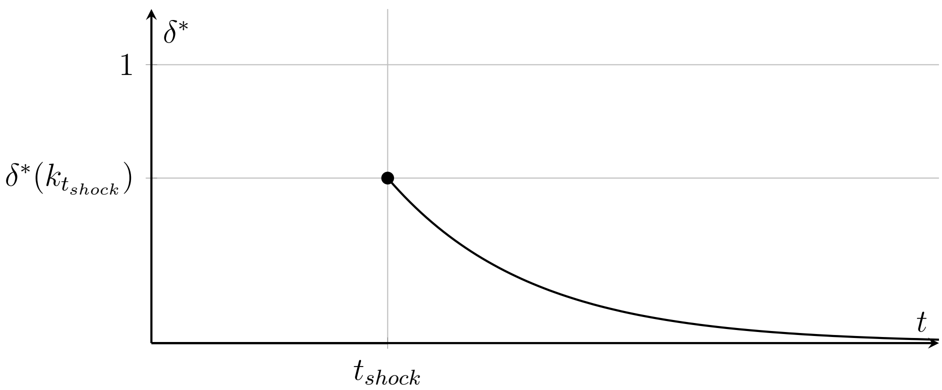

Figure 2: δ^* for $b = 0.7$, $c = 1$, and $\sigma(b, \bar{b}) = \frac{|b-\bar{b}|}{|b|+|\bar{b}|}$.

2.4. Hypotheses

Based on our theoretical model, we formulate empirically testable hypotheses. For the empirical specification, we select the mobilization patterns of BLM and anti-police protests between 2017 and 2022. This case study has some advantageous properties that make an assessment of the proposed model feasible. First, even before, but certainly after the George Floyd incident, occurrences of police violence received a lot of media attention. Thus, we have a setting in which comparable context shocks – fatal police encounters of black individuals – potentially increased the salience of protesting, subsequently increasing protest mobilization, too. This setting allows for the exploitation of different sources of variation. We assess the salience mechanism relying on (1) cross-sectional variation, (2) variation in prior salience levels, and (3) time variation. Due to the different timing of fatal police encounters, we can rely on proper comparisons with a clean control group throughout the paper.

As mentioned above, the kind of context shock under investigation is fatal police encounters of black individuals. The size of the resulting context shocks is proxied by the Google results count about these incidents, with a higher number of results indicating a larger shock. Additionally, by analyzing the publication dates of these search results, we can investigate the temporal dynamics of media coverage and assess the normalization of the context over time. In contrast, the salience level is measured using Google searches, which is an activity by individuals that signals engagement with a topic. The context shock is, thereby, a supply-side measure of how much the topic is covered in the media, while the salience level is the demand-side attention to this topic.

The outcome variable in question is the mobilization for BLM and anti-police protests on the US county level from 2017 to 2022. We aim to test whether the salience mechanism increases the likelihood of collective coordination for protests. The primary motivation for conducting this analysis at an aggregated level is that protests, by definition, involve multiple individuals. Therefore, we seek to capture whether context shocks enable groups to overcome the free-rider problem and engage in protest.

The first hypothesis relates to our basic statement that some specific acts or events can work as context shocks and impact protest outcomes.

Hypothesis 1 (Context shock effect)**.** *We expect to observe significantly more protest events after the occurrence of context shocks than before, i.e., the number of anti-police or Black Lives Matter protest events is significantly increased by police shootings.*

The second hypothesis aims to test the effect of salience on protests. Based on our model, a context shock is a positive salience shift if it increases salience. This salience shift, in turn, leads to an increase in 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).*

In the third hypothesis, we want to test how context shocks affect protests differently. In general, larger shocks should lead to more salience, so the effect on the number of protests should also be greater. As we will classify context shocks by the number of Google results regarding one shooting, more media coverage should lead to more salience, such that we should observe more protest events.

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.*

Furthermore, we test if there are diminishing salience effects for an increasing pre-shock salience level. In Proposition 2, we showed that this is the case if the salience function $\sigma(b, \bar{b})$ is concave in its second argument.

Hypothesis 4 (Diminishing salience effects)**.** $\sigma(b, \cdot)$ *is a concave function: The higher the pre-shock salience level, the smaller the increase in salience.*

In addition to the salience functions' concave curvature, the hypothesis also requires $f(k_s)$ to be injective and b to be constant. These two assumptions are not empirically testable. However, assuming individuals' ability to distinguish between different contexts does not seem controversial, which justifies the injectivity assumption. In addition, we may assume that, if there was any change, the benefit of protesting was higher in 2020 than in 2018. Such a change may be driven by higher moral loadings or stronger network ties among protesters during the high mobilization period in 2020. Furthermore, it can be shown that the diminishing salience effect described in Proposition 1 should be even stronger if there is a higher salience level in 2020 despite a higher b. Thus, if we observe (no) diminishing salience effects in 2020, a concave salience function cannot (can) be rejected.

In the last Hypothesis, we test the dynamic implications of our model. If the context normalizes over time, i.e., returns to its pre-shock level, a context shock's impact should diminish as salience also normalizes.

Hypothesis 5 (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, four empirical estimation strategies are presented: (1) binary treatment evaluation, (2) continuous treatment evaluation, (3) distributed lag model, and (4) difference-in-differences. These estimation strategies are applied in different parts of the results section to test the hypotheses. Finally, the last subsection of this section discusses potential challenges of the empirical estimation strategy. It describes sensitivity tests that are employed to demonstrate the robustness of the empirically identified 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 increases the salience for protesting, which in turn is hypothesized to increase 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 select 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 2017 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[7]. These events constitute (micro-level) context shocks, which are supposed to facilitate protest coordination. Having the database in place, the racial disparities between exposures 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. 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 attention about the fatal police encounters serves

[7]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 classify all counties with at least one fatal police encounter likewise as treated counties. Later, we specify media attention and aggregate a new variable on the county level, which accounts for heterogeneities between treatment intensities.

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

Salience, as a measure of how much attention a certain topic receives in the media and in political discussions, is measured by Google Trends data on the relative frequencies of Google searches for "police shooting." The intuition of 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. 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. Its 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 relations" 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.

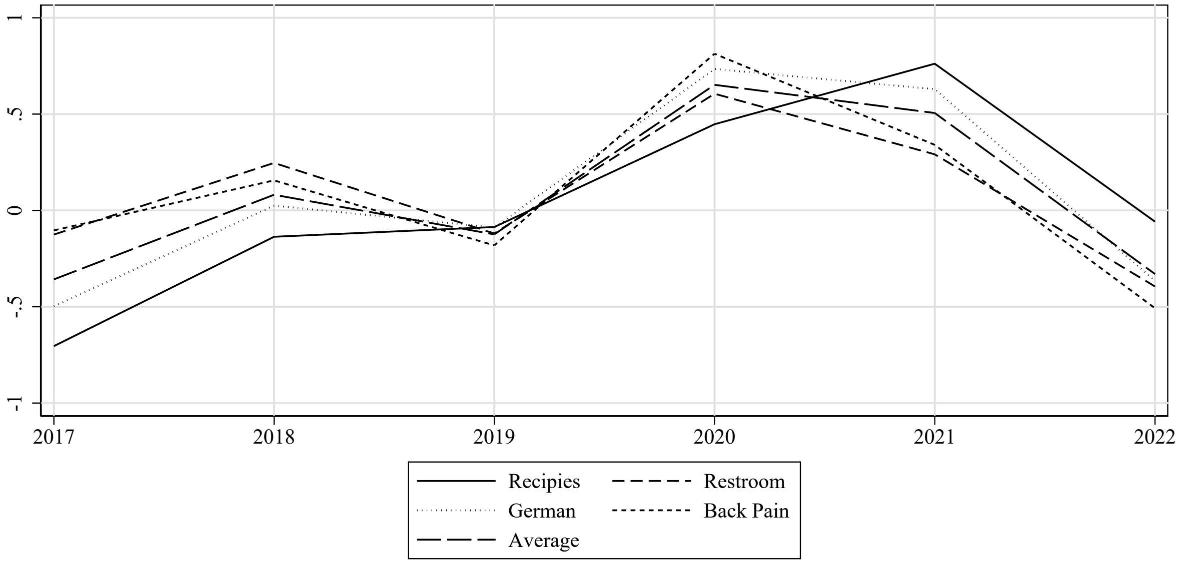

Figure 3: Salience of “Police Shooting” over time across all counties. Note that “Recipies” is misspelled on purpose.

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. Cross-Sectional Estimation Strategy

We employ a set of cross-sectional estimators to test the implications of Hypotheses 1-4. 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, we test to what extent we can observe a diminishing effect of context shocks on salience when the pre-shock salience level increases (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 only 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[8]. Equation 23 presents the target parameter where Y denotes the dependent variable, X denotes the covariates, and W denotes the binary treatment variable.

$$\tau_{\text{ATET}} = E[E[Y \mid X, W = 1] - E[Y \mid X, W = 0] \mid W = 1] \tag{23}$$

The identifying assumptions underlying the estimators are conditional unconfoundedness and common support (Abadie and Cattaneo, 2018). To alleviate potential concerns for a violation of conditional unconfoundedness, the LASSO-selection procedure is applied as suggested by, for example, Belloni *et al.* (2014). From a set of 76 potential control variables, only those that have predictive power for either the treatment or outcome enter as control variables. All relevant details about the LASSO selection procedure are presented in the Appendix[9].

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 8 in Appendix C.1. Much of the density of the

[8]This fact becomes visible in Figure 8. There, a large number of untreated counties has a treatment probability of close to zero.

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

control group is centered at 0, with very low propensity scores. As the target parameter is the ATET, this is no reason to worry because we want to assess whether there are 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 0.0025 or higher than 0.9975 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 with bias-adjust, (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 C.2.

II. Continuous Treatment

Investigating the effect of salience on protest outcomes requires an adaptation of 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 allows for a continuous treatment variable.

The generalized conditional unconfoundedness assumption implies that the potential outcome of all possible treatment dosages needs to 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 remove all bias stemming 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. Using this model, we obtain predicted values for each observation and the standard deviation of the residuals. The GPS for each observation is obtained by assuming a normal distribution of the disturbances. Therefore, we use the conditional Gaussian density given and estimate the GPS ($\hat{r}(w, X_i)$) for each observation.

Having the GPS at our disposal, we follow Hirano and Imbens (2004) in modeling the treatment effect in a quadratic conditional mean model. If conditional unconfoundedness holds, the model fully accounts for all potential confounding variations in the covariate set. The model includes a quadratic term and an interaction term between the treatment dosage and the GPS in order to impose weaker assumptions about the functional form of the treatment effect. The model is estimated as an OLS regression as illustrated in Equation 24. Due to the uncertainty of the first stage estimation of the GPS ($\hat{r}(w, X_i)$), standard errors are obtained using the Delta method.

$$Y = \alpha_0 + \alpha_1 w + \alpha_2 w^2 + \alpha_3 \hat{r}(w, X_i) + \alpha_4 \hat{r}(w, X_i)^2 + \alpha_5 w \hat{r}(w, X_i) \tag{24}$$

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

3.3. Dynamic Estimation Strategy

To test the final hypothesis, we rely on panel data and apply a distributed lag model (DLM) and a difference-in-differences (DID) specification. Hypothesis 5 states that after a normalization of the context over time, protest mobilization should also return to its prior level. The normalization of the context is assessed by considering the publication dates of articles about the context shocks. Then, we construct a finite DLM to have a first glance at the time dynamics of the effects that are so far only assessed in the cross-section. Second, we use a DID model based on De Chaisemartin and d'Haultfoeuille (2024) to more effectively control for potential confounders and assess the dynamics over time more precisely.

I. Distributed Lag Model

Equation 25 presents the model to be estimated. We are interested in the lagged effects of context shocks on the (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 25 is then estimated using an OLS regression with heteroskedasticity-robust standard errors. Since salience is observed at a high 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.

$$Y_i = \alpha_0 + \alpha_1 W_{y,i} + \alpha_2 W_{y-1,i} + \alpha_3 W_{y-1,i} + ...\beta X_i + \epsilon_i \tag{25}$$

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).

II. Difference-in-Differences

As a second, more advanced estimation strategy for identifying the dynamic effects of context shocks, this study employs the flexible DID estimator introduced by De Chaisemartin and d'Haultfoeuille (2024). The DID framework is particularly advantageous in this context, as it accounts for underlying time trends – an essential consideration given the period between 2017 and 2022, which includes the nationwide contextual shock following the George Floyd incident. The estimator proposed by De Chaisemartin and d'Haultfoeuille (2024) is especially well-suited to this application, as it accommodates staggered treatment adoption and permits non-absorbing or non-binary treatment effects. This flexibility enables assessing the dynamic effects of context shocks on both salience levels and protest outcomes across all counties.

The two identifying assumptions of the DID estimator are no anticipation and parallel trends (De Chaisemartin and d'Haultfoeuille, 2024). No anticipation rules out any adjustments of the outcome variable due to anticipated changes in future treatment. As police incidents are of a stochastic nature and happen without any prior planning, this assumption seems to be unproblematic for the present case. As the DID estimator compares the outcome variable's trajectory of treated to non-treated individuals and estimates the difference between them, it is necessary to assume that treated units would have developed in parallel to non-treated units if they had been untreated, too. As the De Chaisemartin and d'Haultfoeuille (2024) estimator compares treatment groups, the parallel trends assumption must be fulfilled for all units that share the same period 1 treatment status. This assumption is partly assessed using placebo estimators.

A treatment design restriction must be fulfilled for the estimation. Therefore, the treatment variable is adjusted to two designs. *Design 1* allows for binary treatment that can switch on and off. In each year in which there has been a fatal encounter between a black individual and the police, the treatment is thus 1 for the given county in the given year and 0 otherwise. For *Design 2*, the treatment intensity is allowed to vary. Hence, the treatment intensity is a variable that counts the current treatment status plus past treatment values. For example, if a county experienced three fatal police encounters in the last 5 years, the treatment status in year 5 is 3. Data on fatal police encounters are available from 2013 onward. Hence, we define starting values for the estimation with *Design 2* to be the sum of all fatal police encounters in the period between 2013 and 2016.

As a note on treatment frequencies, 2,696 units in the sample are never treated, while 526 units receive treatment at least once (treatment frequencies: treated once: 313; twice: 93; three times: 41; four times: 21; five times: 33; six times: 25). The yearly treatment mean of all counties is relatively stable over the time period and always within the range of 0.050 (2019) to 0.059 (2022). We estimate the effects over three successive years. The number of observations for the first lag is 501 (this number differs from 526 due to missing data), decreasing to 364 by the third lag. Further lags are not displayed as the number of observed units continues to decrease quickly. We estimate placebo effects for the treatment preceding years: one year prior to treatment, for which 306 observations are available, and two years prior to treatment, for which 141 observations remain.

3.4. Sensitivity Analysis

The estimation results are complemented with a set of robustness checks. The alternative estimators are presented in the Appendix and referenced in the results section. We focus on four potential sources of bias in the analysis: (1) outliers, (2) different margins, and (3) spatial spillovers. The strategies to circumvent these problems are presented in this subsection. Discussions about confounding and overlap are omitted, as the LASSO-selection and sample trimming strategies to address these concerns have already been discussed.

(1) 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 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.

(2) 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 1 if there was at least one protest event in the given county and 0 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 assess whether the results are only driven by the intensive margin or whether the treatment also has an impact on the extensive margin.

(3) 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. This event is specifically investigated as a context shifter for individual protest decisions in one of the empirical specifications. 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 assumption of independent observational units, 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 75- and 100-mile radius to test the result's sensitivity to spatial spillovers.

4. Results and Discussion

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 issue salience. We then assess potential temporal heterogeneity for different pre-shock salience levels by comparing effect sizes across the years 2018, 2020, and 2022, in line with Hypothesis 4. Finally, we incorporate a DLM and DID specification to explore the dynamic effects of salience over time (Hypothesis 5). A comprehensive set of robustness checks, detailed in the Appendix, 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 lead to increased protest mobilization. As outlined in Subsection 3.2, we investigate whether a context shock caused BLM or anti-police protest events. Also, 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
Context Shock	10.463**	8.768*	8.511*	9.044**	9.728**
	(4.364)	(4.474)	(4.456)	(4.525)	(4.292)
Observations	3007	3007	3007	3007	3007

Heteroskedasticity robust SE in parentheses; + p <0.2, * 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. LASSO selection covariate Set 1 is applied.

The results demonstrate considerable robustness, as evidenced by the sensitivity analyses reported in Table 10 in Appendix D.1. Estimates using the transformed dependent variable suggest that a context shock increases the number of protest events by 40-73.4 percent and raises the probability of at least one protest occurring in a county by 10.5-24.2 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. 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.

In conclusion, the results speak in favor of Hypothesis 1 and provide confidence that fatal police encounters are indeed context shocks for BLM and anti-police protests. Hypothesis 1 states that context shocks increase protest mobilization generally. This effect is found throughout a number of specifications and robustness tests, supporting Hypothesis 1.

4.2. Salience Mechanism

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 increasing the salience for protesting. Hypothesis 2 is thus divided into two mechanisms. Mechanism 1 implies that the context shock increases the salience in a given location. The subsequent Mechanism 2 hypothesizes that increased salience causes 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[10]. The effect of context shocks on salience is robust to the transformation of the dependent variable, accounting for spillovers, accounting for overlap, and shifting the observational unit to the Google Trends regions (Table 11 in Appendix D.1).

Mechanism 2 is also identified relying on the estimation technique for continuous treatments presented in Subsection 3.2. Figure 9 shows the dose-response function of each

[10]The salience distribution is illustrated in Figure 10 in Appendix D.2.

Table 3: Static Treatment Evaluation: Context Shock on Salience

	RA	IPW	IPW RA	C. Match	P.S. Match
Context Shock	2.612***	3.388***	3.464***	2.704***	2.979***
	(0.830)	(0.868)	(0.909)	(0.858)	(0.757)
Observations	2924	2924	2924	2924	2924

Heteroskedasticity robust SE in parentheses; + p <0.2, * 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 1 is applied.

salience level for its effect on protest mobilization. At 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 in Figure 9 in Appendix D.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.

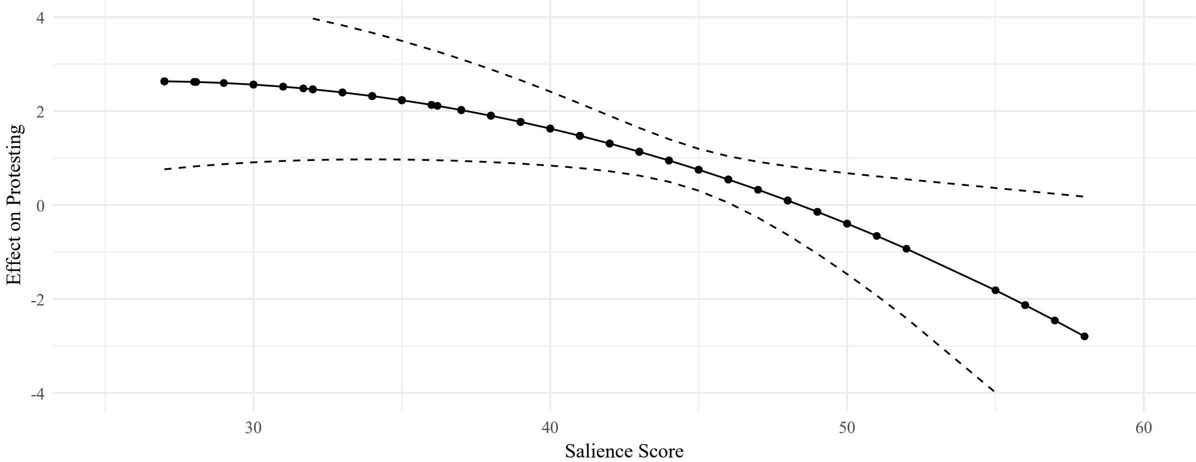

Figure 4: Estimation results continuous treatment of salience on protest events. The dashed lines indicate the 95% confidence interval. The salience score range from 20 to 60, which includes 92% of all salience values. Covariate Set 2 is applied.

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[11] 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

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

that the salience mechanism is the dominant effect for short-run protest mobilization. We return to the question of what drives the remaining gap between the estimated indirect effect via salience and the total effect in subsequent sections.

In conclusion, the results support Hypothesis 2 as they provide evidence for a salience channel from a context shock to protest outcomes. This effect holds for a variety of robustness tests. We continue this analysis in the next subsection, focusing on heterogeneity in the magnitude of context shocks and its implications for protest outcomes.

4.3. Differently Sized Context Shocks

To evaluate Hypothesis 3, which predicts that larger context shocks amplify both salience and protest outcomes to a larger extent relative to smaller shocks, 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 the article count. Under this framework, greater media attention is expected to produce stronger effects on both salience and protest mobilization.

Distinguishing between context shocks, those context shocks that received a lot of media attention seem to have led to more protest events. This tendency is visible 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, thus, seems to influence protest coordination substantially. It should be noted that the estimates for these results are imprecisely estimated, and many of them are identified with a statistical significance level of 20%. The salience shift seems to be highest for large context shocks, for which the salience level increases by 5.2-5.7 units. However, the coefficient size for small and medium context shocks is roughly the same, especially if we account for the standard errors.

The empirical findings are broadly consistent with the hypothesis, although 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 media shocks tend to correspond with larger estimated coefficients for both protest activity and salience level. While standard errors prevent definitive conclusions about statistical differences between these coefficients, the overall pattern suggests a monotonic relationship, particularly between high-salience events and mobilization.

In sum, the results provide suggestive evidence that context shocks in the top quartile

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	6.280	7.083+	3.883
	(4.858)	(5.139)	(5.082)	(5.125)	(5.913)
Observations	2905	2905	2905	2905	2905
Article C. ≥ 6	13.992*	10.201+	9.936+	12.414*	12.397*
	(7.171)	(7.531)	(7.518)	(7.194)	(6.974)
Observations	2875	2875	2875	2875	2875
Article C. ≥ 27	24.937*	21.277+	21.317+	21.214+	18.250
	(14.063)	(14.198)	(14.420)	(14.372)	(14.789)
Observations	2845	2845	2845	2845	2845
Dependent Variable: Salience Level					
Article C. ≥ 3	3.443**	4.256***	4.547***	3.481**	3.345**
	(1.338)	(1.409)	(1.439)	(1.486)	(1.579)
Observations	2824	2824	2824	2824	2824
Article C. ≥ 6	2.998*	3.973**	4.184**	3.046*	2.793+
	(1.639)	(1.679)	(1.682)	(1.769)	(2.165)
Observations	2794	2794	2794	2794	2794
Article C. ≥ 27	5.198**	5.659**	5.221**	5.155**	5.619***
	(2.279)	(2.422)	(2.564)	(2.567)	(1.555)
Observations	2764	2764	2764	2764	2764

Heteroskedasticity robust SE in parentheses; + p <0.2, * 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 2 for the latter 3 sets.

of media attention significantly increase protest activity, with a substantial portion of this effect mediated through increased salience. By contrast, the difference in response between small and medium-sized shocks is much more uncertain.

4.4. Diminishing Salience Effects

In this subsection, we exploit time variation in the salience levels for protesting to assess the functional form of the salience function. The salience level for protesting is much lower in 2018 and 2022 than in 2020 when the George Floyd incident sparked nationwide media coverage (see Figure 3). This allows us to assess the relative effect sizes at different salience levels throughout the different years. Specifically, we want to test Hypothesis 4, which implies that if we observe no diminishing salience effect for higher pre-shock levels, a concave curvature of the salience function can be rejected. This implication is governed by Proposition 1. Furthermore, we examine whether the impact of context shocks on protest outcomes varies across years characterized by differing baseline levels of salience.

Generally, the effect of context shocks on protest outcomes varies substantially, as the estimates in Table 5 show. A context shock is found to increase protest events in treated counties by 18.1-25 units in 2020. This effect is smaller for 2022 when the same context shock increased 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 is (obviously) related to the high level of attention paid to the topic since the murder of George Floyd in the US.

Table 5: 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*	0.530+	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***	18.358***	21.920***	18.146***
	(6.075)	(6.269)	(6.305)	(5.883)	(7.028)
Observations	3013	3013	3013	3013	3013
Context Shock 2022	10.463**	8.768*	8.511*	9.044**	9.728**
	(4.364)	(4.474)	(4.456)	(4.525)	(4.292)
Observations	3007	3007	3007	3007	3007
Dependent Variable: Salience Level					
Context Shock 2018	0.130*	0.225***	0.227***	0.246***	0.276***
	(0.076)	(0.079)	(0.083)	(0.076)	(0.086)
Observations	2926	2926	2926	2926	2926
Context Shock 2020	0.100+	0.148**	0.148*	0.114+	0.139+
	(0.069)	(0.072)	(0.077)	(0.070)	(0.091)
Observations	2934	2934	2934	2934	2934
Context Shock 2022	0.157***	0.172***	0.170**	0.176***	0.154***
	(0.059)	(0.064)	(0.066)	(0.058)	(0.050)
Observations	2907	2907	2907	2907	2907

Heteroskedasticity robust SE in parentheses; + p <0.2, * 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 2 for the latter 3 sets.

Interpreting the results in light of Hypothesis 4, we find evidence for a concave curvature of the salience function. The effect of a context shock on salience is the lowest (and barely statistically identified) for the year 2020, when the salience level was highest. For 2018 and 2022, the response of salience to a context shock is much stronger and, as a tendency, even higher in 2018. These results are consistent with a concave salience function, such that we cannot reject a curvature of this type.

A puzzle that remains is how the stark differences in the overall effects are explainable, given that the effect on salience even diminishes for high levels. The answer is based on Figure 5 where we investigate the response of protest outcomes salience increases as continuous treatments. It becomes visible that for a one-unit increase in salience, protest events increased by 2.5-5.5 units in 2020. This number is substantially lower for 2022 (as already assessed in the previous subsection) and even much lower for 2018. Thus, the differences in the effect size of context shocks on protest outcomes are explained by a higher sensitivity of protest mobilization to the salience level. This effect outweighs the decreasing effect of the context shock on salience, for which we can observe diminishing effects.

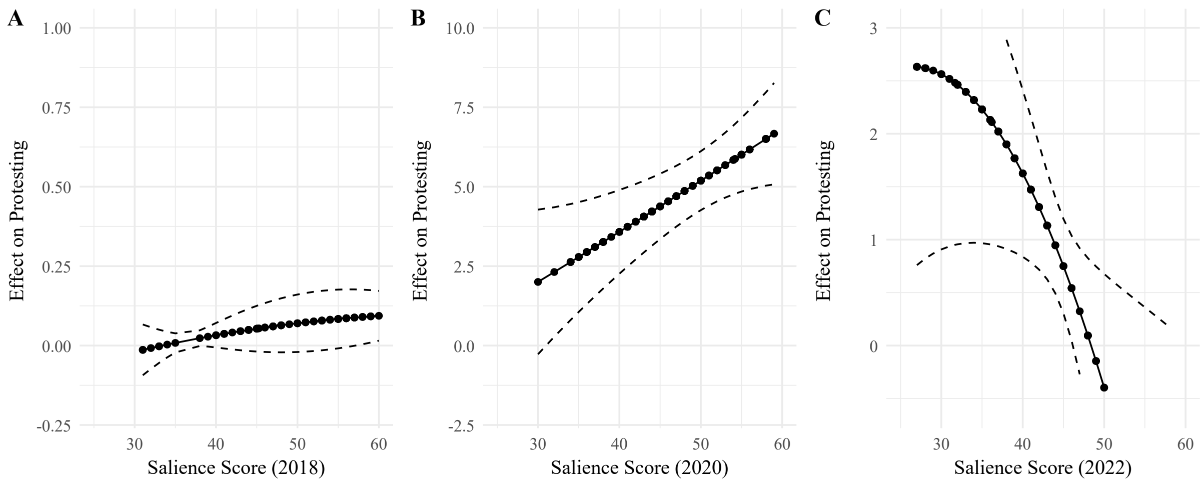

Figure 5: Estimation results fuzzy treatment of salience on protest event for the years 2018, 2020, and 2022. The dashed lines indicate the 95% confidence interval. LASSO selection Covariate Set 2 is applied.

Theoretically, these contrasting effects can be understood by considering differences in baseline salience. In 2020, salience levels were exceptionally high, such that even marginal increases were sufficient to trigger protest mobilization. By contrast, in 2018 and 2022, salience was substantially lower, and only sufficiently large shifts could generate a comparable response. Yet, given that salience levels in 2018 and 2022 were broadly similar, the differential effect across these years must stem from changes in other parameters. Most plausibly, the benefit of protesting, b, increased in the aftermath of 2020. Potentially, this may be due to the formation of protest networks or enhanced organizational capacities. We revisit this interpretation in our analysis of dynamic effects over time.

4.5. Dynamic Effects

The last collection of empirical results tests Hypothesis 5, which states that if a context shock fades out over time, the same fading out should be observable for protest mobilization and the salience level. The assessment begins by investigating the timing of media coverage about fatal police encounters. Next, we apply the DLM to assess the significance of lagged effects. Finally, the DID analysis is supposed to provide evidence in a setting where treated and untreated units are directly compared.

The first piece of evidence is presented in Figure 6, illustrating the media coverage timing of fatal police encounters. Within the first 30 days, more than one-third of articles are published about the fatal police encounters. This indicates that the context, as defined in the model, is shocked most heavily immediately after a fatal police encounter happens. 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 5, that the context normalizes over time, is therefore in line with the data.

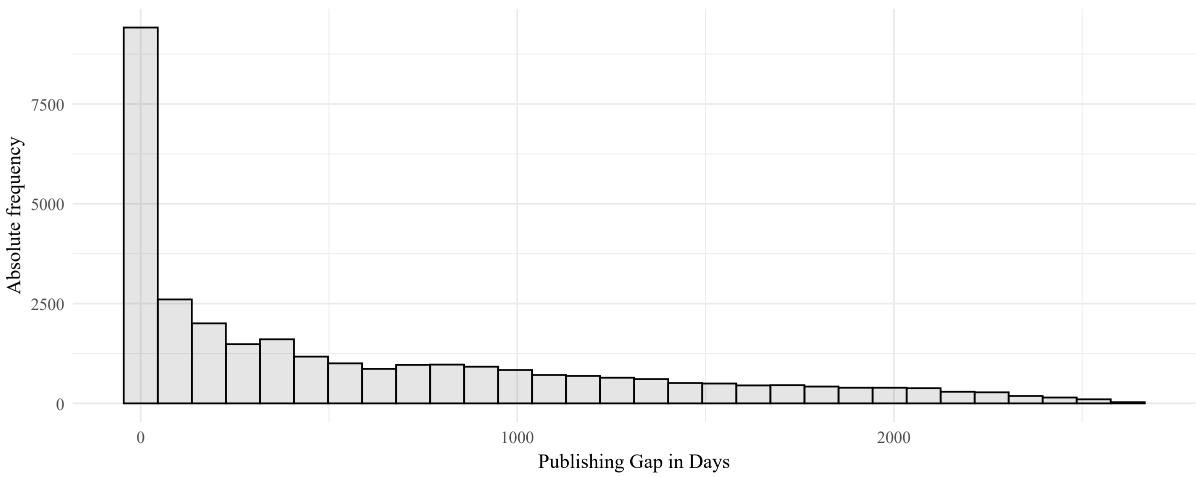

Figure 6: Publishing gap illustration as the difference in days between the fatal police encounter and article publication. The analysis includes all fatal police encounters occurring between 2017 and 2022. Given this window, the maximum observable publication lag for some of the incidents is three years. 35.8% of articles are published within the first month after the fatal police encounter.

After having argued that the context normalizes over time, the model predicts that protest mobilization and salience levels should follow a similar pattern. Interpreting the effect of the first two columns in Table 6, this seems to be the case for protest mobilization. In the current period, a context shock increases protests by 3.47 units or 40%. The level estimation is relatively imprecise, and therefore, differences in the estimates could be due to statistical imprecision. The estimates for the IHS transformation are much more precise, allowing for a cautious interpretation that the effect size decreases with increasing lag levels. But even there, most of the confidence bands intersect, which leaves justified doubt about this conclusion. However, evidence speaks strongly in favor of the claim that lagged context shocks affect protest mobilization in the current period based on both specifications.

Table 6: 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 <0.2, * 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 2 for salience as dependent variable.

In contrast to Hypothesis 5, this effect does not seem to be driven by the salience channel proposed in the model. The immediate increase in salience as a consequence of a context shock becomes visible in columns three and four in Table 6. There are statistically significant effects of contemporaneous context shocks, but no positive significant lagged effects, and for the year 2020, context shocks seem even to decrease the 2022 salience level. Thus, we may have arrived at a point where the explanatory power of the proposed salience mechanism is exhausted, and other approaches must be sought to explain sustained, slowly decreasing mobilization.

The DID results in Figure 7 tell a similar story as the DLM results. For the effect of context shocks on protests, we observe increasing effects over time. The lagged effect of a context shock seems to be much larger than for the contemporaneous treatment[12] (first row in Figure 7). This, again, does not seem to be driven by the salience mechanism. This becomes visible in the bottom two event-study plots, where the contemporaneous effect of a context shock on salience is positive – similar to the DLM result. However,

[12]These results are robust for a transformed dependent variable as visible in Appendix D.4 in Figure 11.

the lagged effects of a context shock on salience become insignificant, and for the non-absorbing treatment, even negative. This finding is plausible in light of the quick return of the context to its normal level as illustrated in Figure 6. The negative effect hint at the fact that there are “attention-hangovers”; that is, after a period of high salience on a certain topic, this topic is over after a year and then receives even less attention in the subsequent years after it.

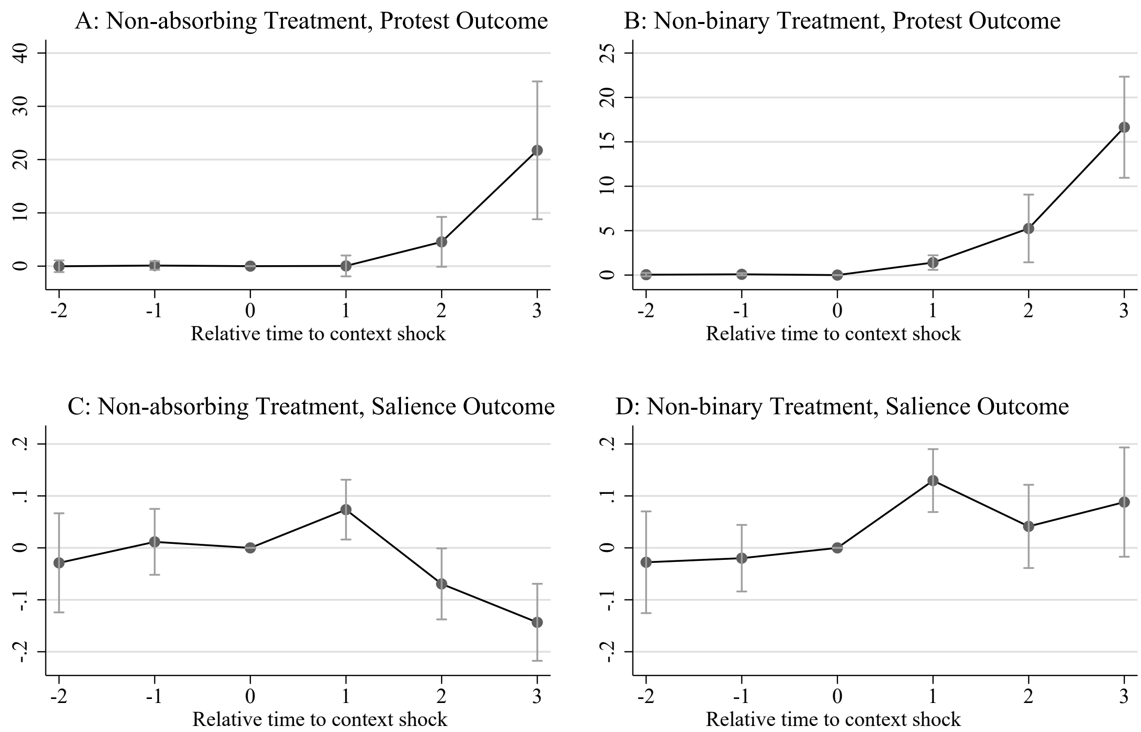

Figure 7: Results for the Difference-in-Differences specification. The graphs depict the effects of context shocks over the period from 2017 to 2022 on the protest event outcome. For the specification with protest outcomes (first row) and salience levels (second row), only units for which all effects are estimable are considered, as this makes the parallel trend more plausible to hold. Covariate Set 1 is applied.

So, if the lagged effects of context shocks are not driven by the salience mechanism, what else drives these effects? A number of candidates are discussed in the literature. For one, networks may form throughout the early mobilization period. Thus, as soon as a context shock happens, the salience mechanism mobilizes individuals to protests, who, during these protests, form protest networks. Such networks may be capable of sustaining protest mobilization within the smaller circle of individuals without having to rely on increased attention in the larger regional environment. A complementary mechanism may be that organizational capacities were built during the early protest stage, facilitating later mobilization, like networks could.

In summary, we presented a variety of empirical evidence supporting the existence of a salience mechanism. The mechanism is found to be of immediate effect and substantially large. Once it comes to sustained mobilization and demobilization, the explanatory power of the proposed mechanism is low, indicating that other approaches need to provide explanations. Considering current dynamics in attention markets, like social media, this finding should not be too surprising. The only thing faster than the emergence of a hype is its certain end. This seems to be true for attention-motivated protesting, too.

5. Extension

So far, we have proposed a new formal mechanism through which salience affects protest outcomes. Furthermore, we formulated five testable propositions for which we provided supporting evidence in the empirical section. As a last piece of evidence, we are interested in answering the question of what makes fatal police encounters larger or smaller context shocks. That is, what are the correlates of events that receive large media attention, and what leads to smaller coverage by the media?

To investigate the drivers of media attention, we regress the article count from Google SerpAPI on a number of incident, victim, and county characteristics and include fixed effects for years and US states. The observational units are all fatal police encounters (including all victim races) in the US from 2018 to 2022 taken from the *Mapping Police Violence* database. In this database, some interesting incident characteristics, such as victim race, victim gender, or whether the police officer wore a body camera, are already included. We merge this database with the county database from Section 3 and include Covariate Set 1 in the regression. Equation 26 states the model that we estimate using OLS.

$$\begin{aligned}\text{Article Count}_i = \beta_0 + \beta_1 \text{Incident Characteristics}_i + \beta_2 \text{County Characteristics}_i \\ + \delta_{\text{year}(i)} + \delta_{\text{state}(i)} + \epsilon_i\end{aligned} \tag{26}$$

Interpreting the results presented in Table 7, the paramount importance of the George Floyd incident becomes visible. All fatal police encounters before the George Floyd incident received 1.5-2.2 fewer article counts. Then, the variable *Since Floyd incident* implies that with each passed day after the incident, the article count decreases by 0.003-0.005. Thus, an incident one year after George Floyd already received 1.1-1.8 articles counts less.

The significance of these effects disappears once we control for time-fixed effects, which already account for such effects.

The media attention on fatal police encounters concentrates on black individuals, as for no other race, robust significant effects were identified. Context shocks, thus, seem to be larger if black individuals were the victims indeed. Furthermore, there are a couple of incident characteristics that drive the media coverage of these events: an armed victim decreases media attention, possibly because the police officer had more reason to apply violence; police body cameras increase media attention, perhaps because there is video material about the incident, the younger the victim, the more attention possibly because the involvement of kids makes such incidents even more tragic. Finally, the fact that a country has a democratic majority also increases media attention, perhaps because of the demand-side effect of media reporting.

6. Conclusion

In this paper, we develop a binary choice model in which protest decisions are driven by salient benefits and costs. Salience arises from distortions between actual and reference benefits, with the latter shaped by memory. Critical events, such as a fatal police encounter with a black individual, shift the reference good, increasing the salience of benefits and thus their decision weight. This mechanism captures how attentional context-mediated exposure to new experience can change protest decisions even when underlying costs and benefits remain unchanged.

We assess the model's predictions with 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 affect protest outcomes substantially and that a large part of this effect is mediated via salience. Furthermore, we highlight the importance of the pre-context shock salience level. Numerous robustness tests support the results.

The paper contributes to a rich literature of 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 be part of the explanation for an equilibrium change. The salience mechanism should be understood as complementary to previous studies, as it explicitly formalizes subconscious perceptual aspects of protest mobilization.

One promising path of future research is to assess more thoroughly the dynamic aspects of attention-based protesting. The results in this study indicate that attention-motivated protesting is a short-term phenomenon that can hardly explain sustained protesting. Other channels, like the formation of networks, organizational capacities, or information cascades, are thus more likely to be informative about intertemporal protest dynamics. A more global perspective combining explanatory elements of multiple theories may provide a very accurate and precise account of protest dynamics.

References

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

— and Imbens, G. W. (2011). Bias-corrected matching estimators for average treatment effects. *Journal of Business & Economic Statistics*, **29** (1), 1–11.

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.

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.

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.

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 Chaisemartin, C. and d'Haultfoeuille, X. (2024). Difference-in-differences estimators of intertemporal treatment effects. *Review of Economics and Statistics*, pp. 1–45.

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.

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.

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: First, for discrete-time $s \in \{0, ..., t\}$ the reference benefit simplifies to

$$\bar{b}(k_t) = \int b_s \cdot w(e_s, e_t) dF = \sum_{s=0}^{t} b_s \cdot w(e_s, e_t)$$

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

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

Please note that the similarities at time $t-1$, $S(e_s, e_{t-1})$, and the perceived benefits b_s until time $t-1$ are independent of the context at time k_t such that $\bar{b}(k_{t-1})$ can be seen as given in t. Thus, the change $\Delta_{\bar{b}}$ in the reference benefit strictly decreases as $\bar{b}(k_{t-1})$ increases.

If $f(k_s) = b_s$ is an injective function, then by the continuity of f 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

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

such that for $\bar{b}(k_{t-1}) \leq (\geq) \frac{1}{t} \cdot \sum_{s=0}^{t} b_s$, it always holds $\bar{b}(k_t) > (<) \bar{b}(k_{t-1})$. In general, inequality 15 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

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

The last equivalence transformation holds since we have

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

As $b_t > (<) \bar{b}(k_{t-1})$, inequality 15 is fulfilled.

q.e.d

Proof of Proposition 1: For given b and $\bar{b}_2 > \bar{b}_1 > b$, we get $\sigma(b, \bar{b}_1) < \sigma(b, \bar{b}_2)$ by ordering.

In the following, we will show the statement if $\bar{b}_1$ and $\bar{b}_2$ are raised by the same amount $k > 0$. As it follows from Lemma 1 that the increase in $\bar{b}_2$ is even smaller, this serves as an upper bound.

Since we have $\bar{b}_2 + k > \bar{b}_1 + k > \bar{b}_1$, there exists a $\lambda \in (0,1)$ such that

$$\begin{aligned} & \bar{b}_1 + k = \lambda \cdot \bar{b}_1 + (1-\lambda) \cdot (\bar{b}_2 + k) \\ \Leftrightarrow \quad & \lambda = \frac{\bar{b}_2 - \bar{b}_1}{\bar{b}_2 - \bar{b}_1 + k} \in (0,1) \end{aligned}$$

As $x \mapsto \sigma(b, x)$ is concave, it has to hold

$$\begin{aligned} & \sigma(b, \bar{b}_1 + k) \geq \lambda \cdot \sigma(b, \bar{b}_1) + (1-\lambda) \cdot \sigma(b, \bar{b}_2 + k) \\ \Leftrightarrow \quad & \sigma(b, \bar{b}_1 + k) - \sigma(b, \bar{b}_1) \geq \lambda \cdot \sigma(b, \bar{b}_1) - \sigma(b, \bar{b}_1) + (1-\lambda) \cdot \sigma(b, \bar{b}_2 + k) \\ \Leftrightarrow \quad & \sigma(b, \bar{b}_1 + k) - \sigma(b, \bar{b}_1) \geq (1-\lambda) \cdot [\sigma(b, \bar{b}_2 + k) - \sigma(b, \bar{b}_1)] \\ \Leftrightarrow \quad & \sigma(b, \bar{b}_1 + k) - \sigma(b, \bar{b}_1) \geq \frac{k}{\bar{b}_2 - \bar{b}_1 + k} \cdot [\sigma(b, \bar{b}_2 + k) - \sigma(b, \bar{b}_1)] \\ \Leftrightarrow \quad & \frac{\sigma(b, \bar{b}_1 + k) - \sigma(b, \bar{b}_1)}{k} \geq \frac{\sigma(b, \bar{b}_2 + k) - \sigma(b, \bar{b}_1)}{\bar{b}_2 - \bar{b}_1 + k} \end{aligned} \tag{A.1}$$

Furthermore, we consider $\bar{b}_2 + k > \bar{b}_2 > \bar{b}_1$. Then, there exists a $\lambda \in (0,1)$ such that

$$\begin{aligned} & \bar{b}_2 = \lambda \cdot \bar{b}_1 + (1-\lambda) \cdot (\bar{b}_2 + k) \\ \Leftrightarrow \quad & \lambda = \frac{k}{\bar{b}_2 - \bar{b}_1 + k} \in (0,1) \end{aligned}$$

As $x \mapsto \sigma(b, x)$ is concave, it has to hold

$$\begin{aligned} & \sigma(b, \bar{b}_2) \geq \lambda \cdot \sigma(b, \bar{b}_1) + (1-\lambda) \cdot \sigma(b, \bar{b}_2 + k) \\ \Leftrightarrow \quad & \lambda \cdot [\sigma(b, \bar{b}_2 + k) - \sigma(b, \bar{b}_1)] \geq \sigma(b, \bar{b}_2 + k) - \sigma(b, \bar{b}_2) \\ \Leftrightarrow \quad & \frac{k}{\bar{b}_2 - \bar{b}_1 + k} \cdot [\sigma(b, \bar{b}_2 + k) - \sigma(b, \bar{b}_1)] \geq \sigma(b, \bar{b}_2 + k) - \sigma(b, \bar{b}_2) \\ \Leftrightarrow \quad & \frac{\sigma(b, \bar{b}_2 + k) - \sigma(b, \bar{b}_1)}{\bar{b}_2 - \bar{b}_1 + k} \geq \frac{\sigma(b, \bar{b}_2 + k) - \sigma(b, \bar{b}_2)}{k} \end{aligned} \tag{A.2}$$

Combining inequality A.1 and A.2, we get

$$\begin{aligned} & \frac{\sigma(b, \bar{b}_1 + k) - \sigma(b, \bar{b}_1)}{k} \geq \frac{\sigma(b, \bar{b}_2 + k) - \sigma(b, \bar{b}_2)}{k} \\ \Leftrightarrow \quad & \sigma(b, \bar{b}_1 + k) - \sigma(b, \bar{b}_1) \geq \sigma(b, \bar{b}_2 + k) - \sigma(b, \bar{b}_2) \end{aligned}$$

such that the statement is shown.

q.e.d

B. LASSO Double Selection

B.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 subjective 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 **??** and **??**.

Following Belloni *et al.* (2014), the selection procedure and the post-LASSO estimations relies 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.3. 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).

$$\text{minimize}_{\boldsymbol{\beta}} \left\{ \sum_{i=1}^{n} \left(y_i - \beta_0 - \sum_{j=1}^{p} \beta_j x_{ij} \right)^2 + \lambda \sum_{j=1}^{p} |\beta_j| \right\} \tag{A.3}$$

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 **??** and **??**. 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).

B.2. Included Variables for LASSO Selection Procedure

B.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 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 age group 50-59, the share of white individuals, share with a university degree, the Gini coefficient, the share of immigrants (from after 2010), the share of immigrants (from after 1990), share of individuals without health insurance, binary variable whether democrats have majority support in the county, measure of political fractionalization, measure of institutional health, and a measure of collective efficacy.

Covariate Set 3: log population, the share of age group 30-39, the share of white individuals, share of married individuals, share with a high school degree, share with a university degree, the share of individuals with a yearly income over 200,000$, the Gini coefficient, the share of immigrants (immigrated after 2010), binary variable whether democrats have majority support in the county, a measure of collective efficacy, and the number of charitable organizations per capita.

Covariate Set 4: log population, the share of age group 30-39, the share of white individuals, share of married individuals, share with a high school degree, share with a university degree, the share of individuals with a yearly income over 200,000$, the Gini coefficient, the share of immigrants (immigrated after 2010), binary variable whether democrats have majority support in the county, a measure of collective efficacy, and the number of charitable organizations per capita, share of married individuals, share with a high school degree, the share of individuals with a yearly income over 200,000$, and the number of charitable organizations per capita.

Covariate Set 5: log population, the share of age group 30-39, the share of white individuals, share with a university degree, the Gini coefficient, the share of immigrants (immigrated after 2010), the share of immigrants (immigrated after 1990), binary variable whether democrats have majority support in the county, and a measure of collective efficacy.

C. Binary Treatment

C.1. Overlap Assumption

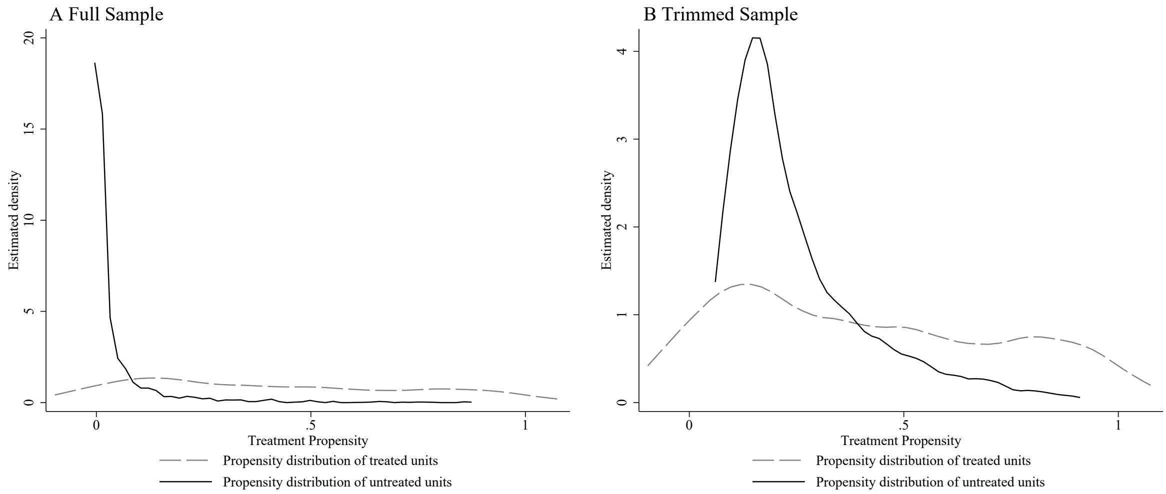

Figure 8: Smoothed density of the propensity score for the treatment and control group. B presents a trimmed sample in which all propensity scores below 0.1 were dropped for the control group to make the overlap of the remaining units visible. The smoothing is conducted relying on the Epanechnikov kernel.

C.2. Estimator Description

(1) A regression adjustment (RA) estimator is the first assessment of the target parameter from Equation 23 (Wooldridge, 2010). The estimator is depicted in Equation A.4 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.

$$\hat{\tau}_{\text{RA}} = \left(\sum_{i=1}^{N} W_i\right)^{-1} \left\{\sum_{i=1}^{N} W_i \left[\hat{\boldsymbol{m}}_1 \left(\mathbf{X}_i\right) - \hat{\boldsymbol{m}}_0 \left(\mathbf{X}_i\right)\right]\right\} \tag{A.4}$$

(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.5 (Wooldridge, 2010). Following Imai and Ratkovic (2014), 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.5.

$$\widehat{\tau}_{\text{IPW}} = N^{-1} \sum_{i=1}^{N} \frac{[W_i - P(X_i)]\, Y_i}{\frac{N_1}{N}\,[1 - P(X_i)]} \tag{A.5}$$

(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 Estimator (2). Given the inverse-probability weights, a weighted regression is fit for each treatment status as for Estimator (3). 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.4 and A.5.

(4) A fourth estimation strategy is propensity score matching (PSM) depicted in Equation A.6 and A.7. 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. The bias adjustment from Abadie and Imbens (2011) is applied to correct for potential bias due to imperfect matching.

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

where

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

(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 bias adjustment due to imperfect matching is also applied for this estimator (Abadie and Imbens, 2011). The ATET estimator is formally depicted in Equation A.8 and A.9 where $\|X_l - X_i\|$ denotes the Mahalanobis distance among observation l

and i.

$$\widehat{\tau}_{\mathrm{COM}} = \frac{1}{n_1} \sum_{i=1}^{n} W_i \left(Y_i - Y_{j(i)} \right) \tag{A.8}$$

where

$$Y_{j(i)} = \sum_{j:D_j=0} I \left(\|X_j - X_i\| = \min_{l:D_l=0} \|X_l - X_i\| \right) Y_j \tag{A.9}$$

D. Sensitivity Analysis

D.1. Robustness Static Model

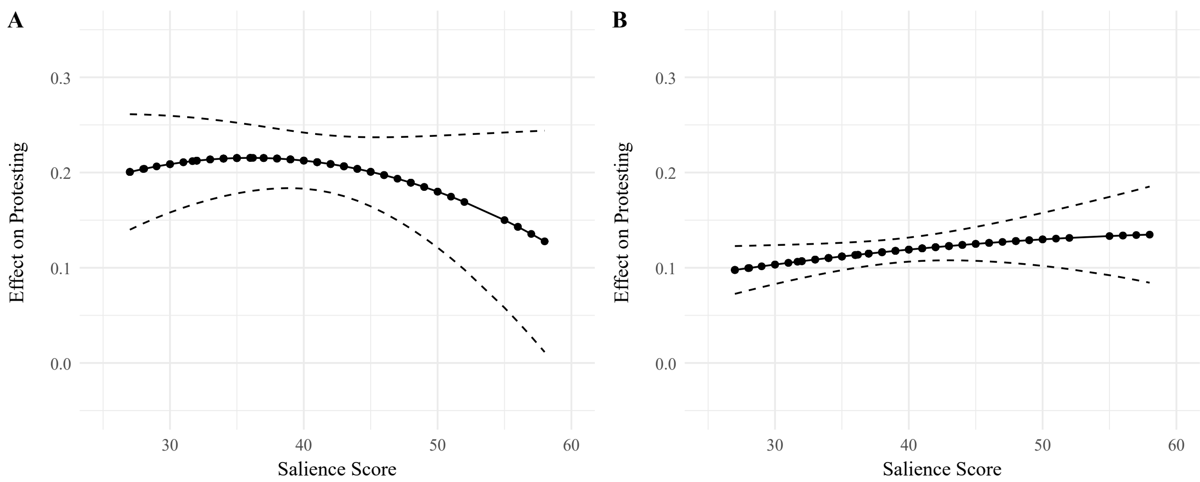

Figure 9: Robustness: Results continuous treatment of salience on protest outcome. The dashed lines indicate the 95% confidence interval. Graph A depicts the fuzzy treatment effect relying on the inverse hyperbolic sine transformed protest variable as the outcome, while Graph B illustrates the fuzzy effect for a binary protest variable as the outcome. LASSO selection Covariate Set 2 is applied.

D.2. Distribution Salience

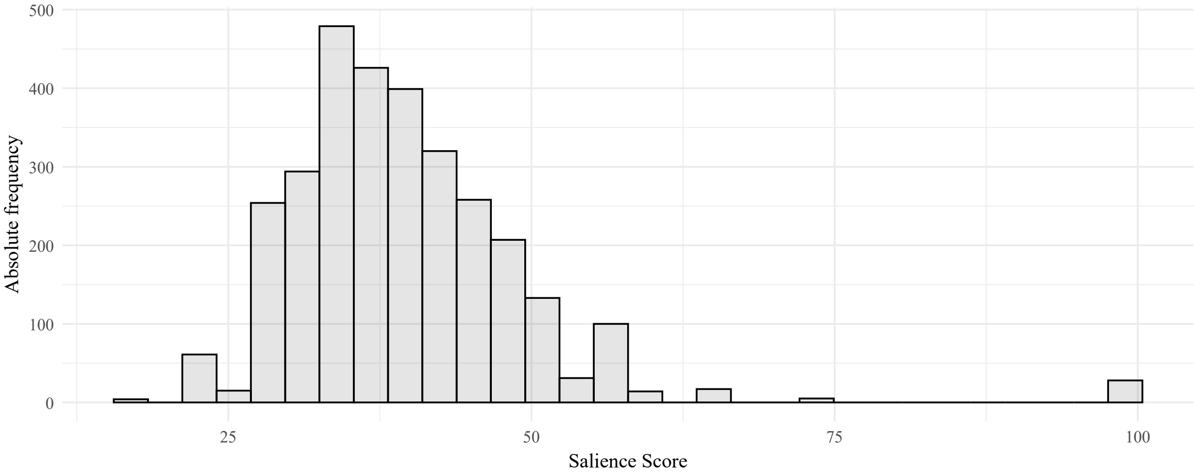

Figure 10: Distribution salience scores in 2022 based on Google Trends data.

D.3. Robustness Differently Sized Context Shocks

D.4. Robustness Difference-in-Differences

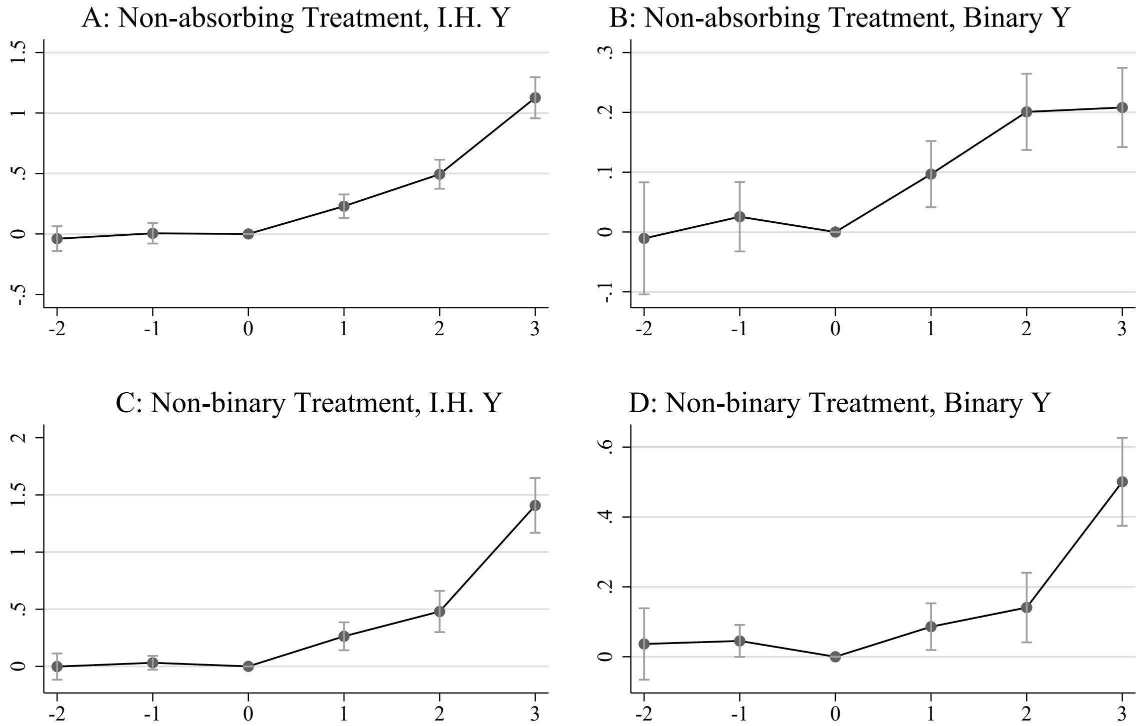

Figure 11: Robustness for the Difference-in-Differences results. The graphs depict the effects of context shocks over the period from 2017 to 2022 on the protest event outcome. Covariate Set 1 is applied.

Table 7: Extension: Drivers of Context Shocks.

	Dependent variable: Google SerpAPI Article Count				
	(1)	(2)	(3)	(4)	(5)
Before Floyd incident (dummy)	−2.192***	−1.516**	−1.696***	0.739	0.705
	(0.575)	(0.649)	(0.633)	(0.945)	(0.951)
Since Floyd incident (days)	−0.005***	−0.003**	−0.003**	−0.002	−0.003
	(0.001)	(0.001)	(0.001)	(0.003)	(0.003)
Race: Black	3.423***	2.846***	2.579***	2.591***	2.590***
	(0.549)	(0.556)	(0.524)	(0.523)	(0.537)
Race: Native	−0.698	−1.008	−0.774	−0.716	−1.543
	(0.816)	(0.823)	(1.466)	(1.462)	(1.569)
Race: Hispanic	1.210**	0.564	0.217	0.170	0.234
	(0.513)	(0.542)	(0.589)	(0.588)	(0.601)
Race: Asian	0.951	0.721	0.133	−0.074	−0.397
	(1.266)	(1.243)	(1.656)	(1.653)	(1.662)
Victim Armed	−2.330***	−1.787***	−1.830***	−1.871***	−1.883***
	(0.358)	(0.505)	(0.443)	(0.442)	(0.445)
Victim fleed	−0.224	−0.207	−0.156	−0.148	−0.041
	(0.403)	(0.440)	(0.438)	(0.437)	(0.440)
Police body camera	3.413***	3.486***	3.199***	3.123***	2.774***
	(0.622)	(0.637)	(0.509)	(0.509)	(0.523)
Victim male		−0.288	−0.272	−0.415	−0.513
		(0.984)	(0.822)	(0.820)	(0.824)
Victim age		−0.054***	−0.050***	−0.050***	−0.048***
		(0.016)	(0.016)	(0.016)	(0.016)
Call for service		−0.374	−0.388	−0.325	−0.147
		(0.434)	(0.420)	(0.419)	(0.422)
Democrat		2.710***	1.954***	2.025***	1.791***
		(0.505)	(0.515)	(0.515)	(0.543)
Log Population			0.486**	0.535***	0.422*
			(0.195)	(0.196)	(0.243)
County Controls	NO	NO	YES	YES	YES
Year FE	NO	NO	NO	YES	YES
State FE	NO	NO	NO	NO	YES
Observations	5,647	3,411	3,410	3,410	3,410
R^2	0.027	0.059	0.066	0.072	0.091
Adjusted R^2	0.026	0.056	0.059	0.064	0.070

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

Note: Results Extension. County controls include all variables of Covariate Set 1.

Table 8: 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 9: 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
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.

Table 10: Robustness Binary Treatment Evaluation: Protest Outcome

	RA	IPW	IPW RA	C. Match	P.S. Match
Inverse hyper. Y	0.734***	0.469***	0.401***	0.648***	0.244+
	(0.110)	(0.133)	(0.141)	(0.126)	(0.155)
Observations	2968	2968	2968	2968	2968
Binary Y	0.242***	0.129***	0.105**	0.193***	0.058+
	(0.034)	(0.041)	(0.043)	(0.046)	(0.036)
Observations	3007	3007	3007	3007	3007
Trimmed Sample	11.847***	10.905**	10.647**	12.249***	10.232**
(75 mi spillovers)	(4.302)	(4.394)	(4.437)	(4.280)	(4.411)
Observations	1345	1345	1345	1345	1345
Trimmed Sample	11.649***	11.177**	10.946**	12.586***	10.981**
(100 mi spillovers)	(4.338)	(4.425)	(4.471)	(4.286)	(4.346)
Observations	966	966	966	966	966
Trimmed Sample	9.465**	8.742*	8.557*	9.333**	7.887*
(overlap)	(4.499)	(4.507)	(4.482)	(4.589)	(4.637)
Observations	1941	1941	1941	1941	1941

Heteroskedasticity robust SE in parentheses; + p <0.2, * 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. For the inverse hyperbolic Y Covariate Set 4 and the binary Y Covariate Set 3 are applied. All other estimations are based on Covariate Set 1.

Table 11: Robustness Binary Treatment Evaluation: Salience Level

	RA	IPW	IPW RA	C. Match	P.S. Match
Inverse hyper. Y	0.043**	0.072***	0.073***	0.066***	0.086***
	(0.017)	(0.019)	(0.020)	(0.020)	(0.027)
Observations	2946	2946	2946	2946	2946
Trimmed Sample	3.819***	6.653***	7.144***	5.263***	8.938***
(75 mi spillovers)	(1.390)	(1.488)	(1.571)	(1.263)	(1.693)
Observations	1232	1232	1232	1232	1232
Trimmed Sample	3.632**	6.368***	6.574***	5.395***	8.094***
(100 mi spillovers)	(1.648)	(2.001)	(2.186)	(1.317)	(1.729)
Observations	859	859	859	859	859
Trimmed Sample	3.935***	3.353***	3.449***	2.538***	3.353***
	(0.876)	(0.882)	(0.912)	(0.902)	(1.043)
Observations	1787	1787	1787	1787	1787
Higher Level	5.363***	6.857***	6.835***	4.346***	5.886***
	(1.551)	(1.870)	(1.873)	(1.414)	(1.087)
Observations	201	201	201	201	201

Heteroskedasticity robust SE in parentheses; + p <0.2, * 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 2.

Table 12: Robustness Diminishing Sensitivity

	RA	IPW	IPW RA	C. Match	P.S. Match
Shooting 2019	0.792**	0.509+	0.357	0.760**	0.439
	(0.367)	(0.396)	(0.422)	(0.360)	(0.384)
Observations	3014	3014	3014	3014	3014
Shooting 2021	13.096**	3.361	-0.827	12.694**	0.525
	(5.474)	(9.093)	(11.118)	(5.089)	(8.786)
Observations	3013	3013	3013	3013	3013
Inverse hyper. Y 2018	0.182***	0.137**	0.103+	0.149**	0.093
	(0.057)	(0.065)	(0.077)	(0.061)	(0.097)
Observations	2974	2974	2974	2974	2974
Binary Y 2018	0.115***	0.087**	0.075*	0.082**	0.074*
	(0.031)	(0.038)	(0.044)	(0.035)	(0.041)
Observations	3014	3014	3014	3014	3014
Trimmed Sample 2018	0.656**	0.682**	0.683**	0.572*	0.470
(75 mi spillovers)	(0.329)	(0.330)	(0.330)	(0.340)	(0.372)
Observations	1351	1351	1351	1351	1351
Trimmed Sample 2018	0.688**	0.691**	0.693**	0.633*	0.711**
(100 mi spillovers)	(0.330)	(0.331)	(0.330)	(0.331)	(0.339)
Observations	987	987	987	987	987
Trimmed Sample 2018	0.595*	0.567*	0.538+	0.530+	0.482+
(overlap)	(0.332)	(0.335)	(0.339)	(0.334)	(0.341)
Observations	2015	2015	2015	2015	2015
Inverse hyper. Y 2020	0.518***	0.589***	0.467***	0.682***	0.301***
	(0.079)	(0.112)	(0.112)	(0.103)	(0.115)
Observations	2974	2974	2974	2974	2974
Binary Y 2020	-0.154***	-0.009	-0.020	0.019	-0.000
	(0.028)	(0.020)	(0.020)	(0.028)	(0.019)
Observations	3013	3013	3013	3013	3013
Trimmed Sample 2020	27.703***	23.465***	22.994***	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)	(6.259)	(6.756)	(6.754)	(6.453)	(5.977)
Observations	1000	1000	1000	1000	1000
Trimmed Sample 2020	21.967***	19.987***	18.361***	21.900***	17.740**
(overlap)	(6.129)	(6.362)	(6.392)	(6.032)	(7.490)
Observations	1999	1999	1999	1999	1999

Heteroskedasticity robust SE in parentheses; + p <0.2, * 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 4 and for the binary Y specifications, Covariate Set 3 are applied. All other estimations are based on Covariate Set 1.