Newcastle University | Business School

Discussion Paper Series

Economics

5

EFMD EQUIS ACCREDITED · ASSOCIATION OF MBAS ACCREDITED · AACSB ACCREDITED · Athena SWAN Bronze Award

The emotional effect of terrorism: Evidence from Twitter data*

Vincenzo Bove[1†], Georgios Efthyvoulou[2†], Armine Ghazaryan[2*†] and Harry Pickard[3†]

[1]Department of Politics and International Studies, University of Warwick, Coventry, CV4 7AL, United Kingdom.
[2]Department of Economics, University of Sheffield, Sheffield, S1 4DT, United Kingdom.
[3]Newcastle University Business School, Newcastle University, Newcastle, NE1 4SE, United Kingdom.

*Corresponding author(s). E-mail(s): a.ghazaryan@sheffield.ac.uk; Contributing authors: v.bove@warwick.ac.uk; g.efthyvoulou@sheffield.ac.uk; harry.pickard@newcastle.ac.uk;
†The authors contributed equally to this work.

Abstract

Terrorism causes emotional reactions among public audiences, with downstream consequences for their well-being, attitudes and policy preferences. We utilise a novel approach which harnesses a unique dataset of Twitter activity from 324K users to precisely capture emotional responses to terrorism. Our results demonstrate that terrorist attacks induce dramatic spikes in various discrete emotions of a negative valence, which vary based on the characteristics of the attacks. Furthermore, the effects on emotions are shown to engender changes in attitudes towards immigration.

Keywords: terrorism; sentiments; emotions; Twitter

*This research was funded by the Economic and Social Research Council (ESRC Grant number: ES/V002333/1).

1 Introduction

The impact of terrorism reaches far beyond its immediate victims. As a form of 'psychological warfare' [1], it aims to instil feelings of fear, anxiety, and insecurity. Thus, evaluating the impact of terrorism on public emotions is essential for assessing the effectiveness of terrorist tactics in achieving their political and social goals [2, 3]. Understanding the precise dynamics of these effects is also crucial to understanding the broad societal impacts of terrorism. Our emotional experiences play an important role in our long-term subjective well-being, flourishing, and life satisfaction [4, 5]. It is therefore not surprising that the emotional impact of terrorism can have detrimental effects on mental health [6, 7] and increase trauma- and stressor-related disorders, which can affect individuals even outside the target country [8, 9]. Our emotions also shape our cognition, processing, decision-making and normative judgments [10–14]. Exposure to terrorism can thus alter publics' cognitions and policy preferences, even where the immediate emotional reactions are short-lived [15–19]. Disaggregating the consequences of attacks based on discrete emotional states among audiences has broader implications for attitudinal responses, as divergent emotional reactions can provoke distinct policy preferences [17, 18, 20], often resulting in the inauguration of harsh policy responses [21–23]. The advent of social media has accelerated and intensified the dissemination of information, making the effects of terrorism even more far-reaching and problematic than before. Emotions expressed by others on social media can directly influence an individual's own emotional state through shared social networks [24].

Against this background, our study explores the emotional effect of eight major terrorist incidents that have occurred in the United Kingdom (UK) since 2016, using a large and unique dataset encompassing 7.6 million observations collected from 324K Twitter users. Though there has been some exploration of utilising Twitter to understand emotional reactions to terrorism [25, 26], no previous studies have employed a large dataset of this nature covering multiple and heterogenous attacks.

The use of Twitter data offers a unique opportunity to gain insight into public reactions to terrorist acts. The platform enables users to comment on news and events from their own accounts in *real-time* [26]. Thus, by analysing the content of tweets, we can measure emotions and attitudes based on individuals' own language and frames of references, and track changes in emotional reactions within short time intervals (every few minutes or hours), which provides an advantage over the use of survey responses. Notably, Twitter users' emotional state has previously been used to assess the population's well-being, to predict and measure mental disorders, and to detect emotional patterns [27–32]. Although we cannot infer to the whole UK population more broadly, focusing on Twitter users allows us to perform a comparatively hard test of the terrorism-induced effects on emotions, which should be more stable among social media users who are exposed to multiple news stories at the same time [33].

To isolate the causal effect of terrorist attacks on emotions, we focus on a short time window (from 3 days before up to and including 3 days after each attack) and exploit variation within individuals, net of potential temporal unobserved factors. In this way, we can estimate whether the tweets posted by a given individual in the short period after a specific attack convey more negative feelings than those posted by the same individual in the short period before the same attack. Our analysis reveals that terrorist attacks induce dramatic spikes in various discrete emotions of a negative valence, and that fear is the emotion that displays the largest and more persistent post-attack rise. The observed patterns persist when we focus on non-terror-related tweets, suggesting that people experience negative emotional reactions in all areas of online discourse following a terrorist attack. We also find that the emotional effect of terrorism is amplified by the motivation of the attacker and the number of victims, and can engender changes in attitudes towards immigrants.

2 Results

The emotional effect of terrorism

We start by comparing the emotional content of tweets posted 3 days after an attack to that of tweets posted 3 days before the attack. Panel A of Table 1 reports the post-attack change in the overall negative sentiment (column (1)) and in four negative emotions: fear, anger, sadness, and disgust (columns (2)-(5)). Overall, we can see that all emotions of negative valence are heightened in the aftermath of terrorist attacks: the treatment (post-attack) effect is positive and highly statistically significant throughout. Comparing the estimates in the last four columns, we can also see that fear is the emotion that displays that largest increase (by about 10% compared to the pre-attack mean). Importantly, these baseline results remain essentially the same when we add the fixed effects and control variables in a progressive manner (see *SI Appendix* Table A3.b).

The overall public mood in social media adjusts very quickly to new events and information [34]. This raises the possibility that the emotional reactions following a terrorist attack may be influenced by subsequent government activities, communications, or other unobserved factors, rather than the attack itself. Failing to isolate the impact of terrorism threat itself leads to the problem of "compound treatment" [35], which can undermine the quality and consistency of inferences. To address this issue, we estimate the post-attack effects on a narrower sample of tweets posted 24 hours before and after each attack. As shown in panel B of Table 1, all negative emotions are substantially more intense in the few hours after an attack (as captured by the interaction of the treatment variable with a 24-hour bandwidth), and then return to baseline levels in the following days (as captured by the treatment variable alone). Furthermore, fear and anger appear to be the dominant negative emotions, with estimates suggesting that tweets posted 24 hours after an attack contain 21% more fear and 14% more anger compared to those posted

24 hours prior to the attack.

[Table 1 about here]

To provide further insights on the temporal dynamics of the emotional effects, we aggregate the tweets at the hourly level and replace the treatment variable with time indicators representing each hour before and after the selected attacks. Figure 1 illustrates the 3-hour moving average estimates of these indicators, with the hour before the attack serving as the baseline. The post-attack emotional reactions can be classified into two groups: fear and anger, which exhibit an immediate surge and gradually decrease over time; and sadness and disgust, which slightly increase and persist at that level until about 14 hours after the attack. It is worth noting that the emotional content of tweets posted 1-24 hours before the attacks does not appear to differ significantly from that of tweets posted 1 hour before the attacks, indicating the absence of pre-existing trends. This is also corroborated when we examine an extended version of this figure based on the full time window (see *SI Appendix* Figure A.3).

[Figure 1 about here]

To what degree do the observed effects stem from tweets that mention terrorism? To answer this question, we run separate regressions for tweets that contain the word 'terror' and other related terms – as identified using a Word2Vec algorithm – and those that do not contain such terms. Table 2 presents the results for these two sub-samples, based on the specification with the 24-hour bandwidth. Examining the terror-related tweets in panel A, we observe that the immediate emotional reactions to a terrorist incident are quite large, in some cases up to three times larger than for the full sample of tweets, and persist for the next few days. This may be due to the fact that tweets about terrorism are more likely to capture feelings elicited by the terrorist event itself for several days after it occurs. Turning to the non-terror-related tweets in panel B, we find that the estimates are similar to those reported in Table 1, indicating that people experience negative emotional reactions in all areas of online discourse following a terrorist attack. In other words, exposure to terrorist violence affects individuals' overall emotional state and the language they use to express their thoughts, even when they are not explicitly discussing terrorism-related issues or the incident itself.

[Table 2 about here]

Heterogeneity analysis

The magnitude of emotional responses is not expected to be uniform across all eight events. In this section, we examine heterogeneity in the effects with respect to two attack characteristics that have been linked to heightened threat perceptions and increased negative emotions: the motivation of the attacker (i.e., whether the attack is motivated by Islamic extremism) and the number of victims. Additionally, we explore differences in emotional responses based on the amount of media coverage an attack receives, which can be used as a proxy for the event's relevance and national significance, as the media tend to give more attention to attacks that are perceived as more consequential and threatening to the general public [36].

Figure 2 presents the post-attack estimates for the 24-hour bandwidth (based on the specification in panel B of Table 1) when we run separate regressions for the following attack groups: (i) the six attacks with Islamist perpetrators versus the two attacks with far-right perpetrators; (ii) the four attacks with the highest number of victims (as indicated by the number of deaths and injuries) versus the remaining four attacks; and (iii) the four attacks with the highest media coverage (as measured by the number of LexisNexis hits in the week following the attack) versus the remaining four attacks. Two key conclusions emerge from this analysis. First, attacks motivated by a radical interpretation of Islam result in more fearful sentiments than far-right attacks. This can be attributed to the fact that the former attacks are generally perceived as posing a more systematic threat to national security and democratic values. Second, attacks with a high number of victims and extensive media attention elicit significantly more negative sentiment and emotional responses than those with relatively fewer victims and less media coverage. The difference in effects is substantial across all outcome variables, but it is particularly pronounced for the fear content in tweets, which is three to four times higher for high-victim / high-coverage attacks.

Due to high correlation between the three conditioning factors (see *SI Appendix* Table A3.c), one has to be very cautious in prioritising and uncovering links among them. Nevertheless, the analysis here clearly indicates that the emotional effect of terrorism is stronger for attacks that are deemed more threatening or consequential than others.

[Figure 2 about here]

Second-order effects

So far, we have focused on the 'first-order effects' of terrorism; that is, the emotions that are triggered by the attack itself. In this section, we examine one of the 'second-order echo effects' of terrorism, its impact on attitudes and emotions towards immigration. Extant research has shown how, after terrorist attacks, members of the broader audience tend to distance themselves from strangers and out-groups in general, and to develop negative attitudes towards immigrants [19, 37].

To accomplish this, we analyse the emotional content of tweets related to immigration; i.e., tweets containing the word 'immigration' and other related terms identified by a Word2Vec algorithm. We consider the period of three days before and after each attack (including the day of the attack). Figure 3 compares the pre- and post-attack average values of the negative sentiment and emotions about immigration, calculated using the share of words assigned to a given sentiment/emotion across all lexicon-identified words included in the immigration-related tweets. As shown in the upper panel of the figure, there is a notable increase in negative feelings about immigration following terrorist attacks, particularly in the overall valence, fear, and sadness.

These differences are expected to be influenced by the identity of the perpetrator. Previous research has shown that following Islamic attacks, the general public is more likely to view foreigners and out-groups as a threat to the homogeneity of the nation-state population [38–40]. Conversely, after far-right attacks, people may soften their views towards immigrants to distance themselves from the ideology of the perpetrator [41]. To test for this, we compare the pre- and post-attack emotional content of immigration-related tweets separately for Islamic attacks and far-right attacks. The results confirm the above expectations: while there is a substantial increase in negative sentiment and emotions about immigration following the six Islamic attacks (by 50% to 100%), the corresponding effects for the two far-right attacks are in the opposite direction, albeit small in magnitude (see the lower panels of Figure 3).

[Figure 3 about here]

In Table 3, we address the same question using regression analysis. In panels A and B, we present the results for all attacks and Islamic attacks, respectively, by including attack fixed effects, in addition to hour fixed effects and tweet-level controls (i.e., number of retweets, replies, likes and quotes). This approach enables us to compare emotional reactions about immigration around the same attack, while controlling for unobserved factors related to time and other tweet characteristics. Panels C and D present the results when we replace attack fixed effects with individual $\times$ attack fixed effects, which allows us to exploit within-individual variation around the same attack. The latter specification is useful as it enables more robust causal inferences, but it has the disadvantage of reducing the statistical power of our analysis, as only a small number of individuals post immigration-related tweets before and after the same attack. Regardless of the specification used, we observe a large and highly statistically significant increase in negative feelings about immigration after Islamic attacks, with fear (about immigration) being the emotion that displays the largest and most persistent post-attack rise, in line with the patterns observed in Figure 3.

[Table 3 about here]

Further analyses and robustness tests

We probe the robustness of the main results in a number of auxiliary analyses, which are all reported in detail in *SI Appendix*.

In Section B.1, we test the sensitivity of our results to using restricted samples of Twitter users: those who are present in our dataset before and after all sampled attacks; and those who posted the same number of tweets before and after a given attack. Overall, our inferences do not change: once again, we find that the tweets posted 24 hours after the attacks convey more negative feelings than those posted 24 hours before the attacks. This guards against the concern of selection bias due to terrorist attacks being correlated with Twitter users' engagement with the platform or the frequency of their posts.

In Section B.2, we perform a number of tests to strengthen our causal inference and rule out possibility of a spurious relationship. First, we focus on one of the most important attacks in our sample (the 2017 Westminster attack) and set the attack date to be 1 week prior to the actual date. Second, we benchmark our results against a failed and not immediately reported attack: the 2017 assassination attempt of PM Theresa May. Third, we examine the treatment effect on outcomes that should not be affected by terrorist events; namely, people's feelings about the weather. In all cases, we find no evidence of significant spikes in negative emotions in the aftermath of the incidents. As a further test, we perform Monte Carlo permutation tests that randomly shuffle the data 500 times and estimate a treatment effect for each random draw. The permuted data produce estimates which are lower than those in Table 1, suggesting that there is 0% probability that the observed effects are observed by chance.

In Sections B.3 and B.4, we examine the treatment effect on positive emotions – using the Emolex sentiment analysis tool – and check robustness to using alternative tools (VADER, Textblob, and composite indices). Overall, the patterns observed are in line with our previous findings and do not seem to be influenced by the method we use to measure sentiment.

In Section B.5, we explore the conditionality of effects upon geographic proximity, as captured by the proximity in kilometers between the user's geo-tagged location and the attack location. The analysis suggests that, while physical proximity can play a moderating role in how individuals respond to terrorism, this role is rather weak. This is likely due to the severity and emblematic nature of the attacks in our sample.

Finally, in Section B.6, we estimate our model separately for each of the eight individual attacks. In all cases, we find evidence of heightened negative emotions in the aftermath of the incidents, suggesting that our results are not driven a small subset of the sampled attacks. In line with our previous results, we also find that the effects are stronger and statistically more robust for attacks with a high number of victims, widespread media coverage, and Islamist perpetrators. This is also verified in Section B.7, where we consider the temporal dynamics of the emotional effects (using time-to-event analysis) for the different groups of attacks.

3 Discussion

Terrorist attacks trigger strong emotional responses that affect well-being, attitudes and decisions. In this study, we investigate the emotional impact of eight major terrorist attacks that occurred in the UK between 2016 and 2020 on the sentiments and emotions expressed in tweets. The results present evidence of dramatic spikes in negative emotions in the immediate wake of the incidents, particularly in fear and anger. Strikingly, these effects persist when we focus on tweets that are unrelated to terrorism, indicating the far-reaching and pervasive effects of terrorist acts on public sentiment.

Our results challenge the notion that European audiences have become desensitised to terrorism due to its frequency in recent years [42, 43]. Rather, we find that the emotional responses are conditioned by the characteristics of the attacks. Islamist attacks, which often result in a high number of casualties and receive extensive media coverage, elicit a stronger fear response from the public than far-right attacks. This may be due to the media and policymakers framing Islamist attacks as the work of organised terrorist cells, while portraying right-wing attacks as isolated, 'lone wolf' incidents [44, 45].

We also explore the attitudinal shifts that might accompany emotional reactions to terrorism. In particular, we investigate whether terrorist attacks affect attitudes towards immigration. Our findings reveal a significant increase in negative feelings about immigration following terrorist attacks, providing evidence that terrorism can fuel anti-immigrant sentiment, likely due to hostility towards the perceived out-groups [19, 37, 46]. Notably, these effects are driven by Islamic attacks.

A growing body of scientific literature interrogates the effects of societal challenges such as climate change and Covid-19 on emotional responses and well-being [47, 48]. The availability of vast amounts of Twitter data has allowed researchers to monitor large-scale emotional changes in real-time, offering valuable insights for public health campaigns [29]. Part of this body of research's aim is to provide policymakers and healthcare professionals the tools to develop and disseminate evidence-based strategies and interventions to promote psychological well-being. In a similar vein, understanding the emotional impact of terrorism is important for policymakers tasked with managing the public's 'collective trauma' [49].

Methods

Data and Variables

Our empirical analysis is based on data taken from Twitter. Twitter is the second highest ranking social media website in the UK (behind Facebook), and in April 2020, its monthly social network market share in the country was around 37 percent.

Users of Twitter can comment news and communicate their views on current events from their own accounts in *real time* [26]. As such, contrary to questionnaires, analysing Twitter data allows us to track changes in an individual's attitudes and emotions within very short time intervals – e.g., every few minutes or hours – which can help to produce valid causal estimates of theoretically relevant shocks. In addition, it allows us to study a person's emotions and attitudes based on their own language and frames of reference, rather than their responses to survey questions, which are inherently subject to some misinterpretation and bias. For instance, as stressed by the literature on response, survey questions do not only measure public opinion; they can also shape and channel it by the manner in which they frame issues, order the various alternatives, and set the context [50].

Although we cannot infer to the whole UK population more broadly, focusing on Twitter users allows us to perform a comparatively hard test of the terrorism-induced effects on emotions and attitudes, as these outcomes should be generally more stable among social media users who are exposed to multiple news stories at the same time [33]. There are also two important advantages in using Twitter rather than other social networking websites. First, in Twitter, a large percentage of the messages posted by users (tweets) are freely accessible contrary to other social networks; and second, Twitter allows also to geo-tagging the tweets, which can be used to analyse emotions and attitudes at some sub-national level [51].

We use Twitter's API V2 to obtain English language tweets with a geotag in the UK. We sample tweets that were posted around the timing of eight major terrorist incidents: the murder of MP Jo Cox in June 2016, the Westminster attack in March 2017, the Manchester Arena bombing in May 2017, the London Bridge attack in June 2017, the Finsbury Park attack in June 2017, the Parsons Green bombing in September 2017, the London Bridge stabbings in November 2019, and the Reading stabbings in June 2020. These are considered to be the most salient (domestic) attacks that occurred over the period 2016-2020: all eight attacks resulted in fatalities or a large number of injuries and received widespread national media coverage. This implies that, regardless of where each attack occurred, individuals from all over the UK were potentially exposed to them. *SI Appendix* Section A.1 offers background material on these attacks.

To minimise the possibility of other events driving the estimated effects, we employ a short-range time window around the attacks: 3 days before, the same day, and 3 days after each attack. Also, since we are interested in variation within individuals, we only keep Twitter users with both pre- and post-attack tweets for at least one of the sampled terrorist incidents. This procedure results in a large individual-level unbalanced panel consisting of around 7.6 million observations (24 observations, on average, per individual).

In order to measure the valence and emotional content of the text contained in each tweet, we use a dictionary-based method, the NRC Emotion Lexicon (EmoLex) [52, 53], developed by crowd-sourced manual annotations. The lexicon contains 14,182 words and 25,000 senses, and each one of these words/senses is linked to two sentiments (negative and positive) and eight emotions (anger, fear, sadness, disgust, anticipation, trust, surprise, and joy). The sentiments are assigned either a value 1 (associated) or a value 0 (not associated); whereas the emotions are assigned a value from 0 to 1, capturing the share of lexicon-identified words/senses in a tweet that are linked to a given emotion. We focus on the negative sentiment and emotions given their strong influence on judgement and choices [11, 54], and their high correlation with offline behaviour – see, e.g., evidence on the relationship between negative tweets about Islam and offline hate crimes [55]. Furthermore, analysing positive emotions in the aftermath of terrorist attacks can be troublesome since the text may also capture words of empathy towards the victims.

Besides the textual content and geotag data, we also retain some additional information about the tweets (number of retweets, replies, likes and quotes), which we introduce in our model to control for heterogeneity with respect to tweet-specific characteristics. More details about the Twitter data collection and coding are presented in *SI Appendix* Section A.2. Descriptive statistics of all variables used in our analysis are provided in *SI Appendix* Table A3.a.

Identification strategy

Our model specification allows us to exploit variation within individuals, net of potential temporal unobserved factors. More formally, it can be written as follows:

$$Y_{ijs} = \beta \mathit{Post\text{-}attack}_{ijs} + \delta \mathbf{X}_{ijs} + \vartheta_{ih} + \lambda_{js} + \varepsilon_{ijs} \tag{1}$$

where Y_{ijs} is the sentiment or emotion linked to tweet i posted by individual j around attack s; $\mathit{Post\text{-}attack}_{ijs}$ is a binary indicator that takes value 1 if the tweet was posted after the minute of the attack, and 0 otherwise; $\mathbf{X}_{ijs}$ is a vector of tweet-level controls, as described above; ϑ_{ih} represents hour fixed effects (capturing the hour of the day, h, that the tweet was posted); λ_{js} represents individual $\times$ attack fixed effects; and ε_{ijs} is an error term, clustered at the individual level. Our parameter of interest, β, measures the effect of terrorism on the outcome variable, with a positive (negative) value indicating that exposure to terrorism strengthens (weakens) the corresponding sentiment or emotion.

It must be stressed that the inclusion of individual $\times$ attack fixed effects eliminates any time-invariant sources of heterogeneity across individuals, and controls for the possibility that individuals posting tweets before the attacks are systematically different from those posting tweets after the attacks. Thus, rather than comparing the tweets of different individuals and around different attacks, we estimate whether the tweets posted by a given individual after a specific attack convey more negative feelings than those posted by the same individual before the same attack. Furthermore, adding hour fixed effects in Eq. (1) accounts for residual heterogeneity arising from the hour of the day that the tweet was posted (e.g., night hours vs day hours).

As noted above, by employing a short-range time window around the attacks, we can reduce the potential for bias due to other events. Given how quickly the overall public mood in social media changes and adjusts to new information [34], we also present results when we focus on a narrower time window. To do that, we augment Eq. (1) with a binary variable (*24hr-bandwidth*) capturing the 24 hours before and the 24 hours after each attack, together with its interaction with the *Post-attack* dummy. In this way, we can estimate how individual-specific feelings change in the few hours after the attacks compared to the few hours before the attacks.

Our identification strategy relies on the assumption that the timing of the event in question is exogenous and unexpected. This is clearly the case of violent events, such as the assassination of political leaders or terrorist attacks [35]. A remaining threat to our identification arises from the possibility of selection into tweeting around the attacks. An important reason why this threat is less acute in our context is that we exploit variation within individuals who have at least one tweet both before and after an attack. However, to further ensure that selection in not affecting our results – e.g., when Twitter users systematically change the topic and the frequency of their tweets in the wake of a terrorist incident – we adopt two complementary approaches. First, we run separate regressions for tweets that contain terror-related terms, and those that do not contain such terms (see Section 2). These are identified using the Word2Vec algorithm [56], which is trained on a Google News dataset containing about 100 billion words. Second, we test whether our results persist when we restrict the sample to include the Twitter users who are present in our dataset before and after all sampled attacks, and those who posted the same number of tweets before and after a given attack (see *SI Appendix* Section B.1).

Table 1 The emotional effect of terrorism: baseline results

	Negative (1)	Fear (2)	Anger (3)	Sadness (4)	Disgust (5)
Panel A: Simple specification					
Post-attack	0.003***	0.006***	0.003***	0.002***	0.002***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Mean of DV (pre-attack)	0.115	0.057	0.056	0.059	0.039
Observations	7,643,102	7,643,102	7,643,102	7,643,102	7,643,102
Number of users	323,992	323,992	323,992	323,992	323,992
Panel B: Interaction with a 24-hour bandwidth					
Post-attack	0.000	0.002***	0.001***	0.000	0.001***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
24-hour bandwidth	–0.000	–0.000	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Post-attack × 24-hour bandwidth	0.010***	0.012***	0.008***	0.005***	0.004***
	(0.001)	(0.000)	(0.000)	(0.001)	(0.000)
Mean of DV (pre-attack)	0.115	0.057	0.056	0.059	0.039
Observations	7,643,102	7,643,102	7,643,102	7,643,102	7,643,102
Number of users	323,992	323,992	323,992	323,992	323,992
Individual × attack FEs	✓	✓	✓	✓	✓
Hour FEs	✓	✓	✓	✓	✓
Tweet-level controls	✓	✓	✓	✓	✓

Notes: The dependent variable (DV) is the sentiment or emotion shown in the first row. Time window: 3 days before, the same day, and 3 days after each attack. The tweets are aggregated at the minute level. *Post-attack* is a binary variable that takes value 1 if the tweet was posted after the minute of the attack, and 0 otherwise. *24-hour bandwidth* is a binary variable capturing the 24 hours before and the 24 hours after each attack. Standard errors are clustered at the individual-level and reported in parentheses. * $p < .10$; ** $p < .05$; *** $p < .01$.

Fig. 1 The emotional effect of terrorism: time-to-event analysis

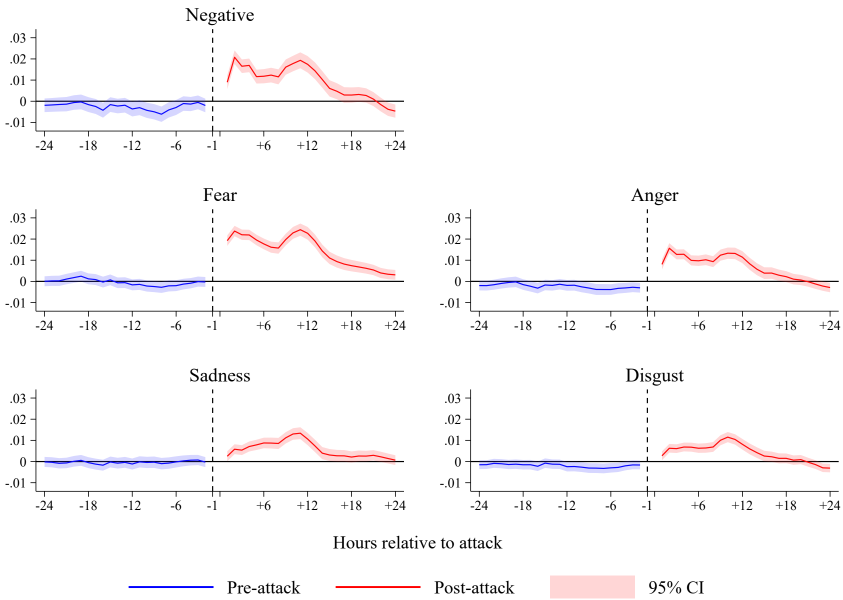

Notes: The figure shows the evolution of negative feelings 24 hours before and 24 hours after the sampled attacks. The tweets are aggregated at the hour level. The blue (red) solid line shows the 3-hour moving average estimates before (after) the attacks, taking the hour before the attack as the baseline. The tweets posted in the hour after the attack are dropped from the estimations. The shaded areas show the 95 percent confidence intervals.

Table 2 The emotional effect of terrorism: terror-related vs non-terror-related tweets

	Negative (1)	Fear (2)	Anger (3)	Sadness (4)	Disgust (5)
Panel A: Includes terror-related tweets only					
Post-attack	0.011	0.036***	0.020***	0.007	0.011**
	(0.008)	(0.006)	(0.006)	(0.005)	(0.005)
Post-attack × 24-hour bandwidth	0.020*	0.032***	0.026***	0.010	0.012
	(0.011)	(0.011)	(0.009)	(0.012)	(0.008)
Terror-related obs.	33,283	33,283	33,283	33,283	33,283
Panel B: Excludes terror-related tweets					
Post-attack	–0.001	0.001***	0.000***	–0.000	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Post-attack × 24-hour bandwidth	0.008***	0.010***	0.006***	0.003***	0.003***
	(0.001)	(0.000)	(0.000)	(0.001)	(0.000)
Non-terror-related obs.	7,584,596	7,584,596	7,584,596	7,584,596	7,584,596
Individual × attack FEs	✓	✓	✓	✓	✓
Hour FEs	✓	✓	✓	✓	✓
Tweet-level controls	✓	✓	✓	✓	✓

See notes for Table 1. The variable *24-hour bandwidth* is included in all estimations. Terror-related tweets are those that include the word ‘terror’ and/or any related terms, as identified using a Word2Vec algorithm; i.e., extremism; jihad; islamist; islamic; radical; militants; suicide/bomb; bombing; terrror; teror; isis; isil; far-right. Standard errors are clustered at the individual-level and reported in parentheses. * $p < .10$; ** $p < .05$; *** $p < .01$.

Fig. 2 The emotional effect of terrorism: heterogeneity analysis

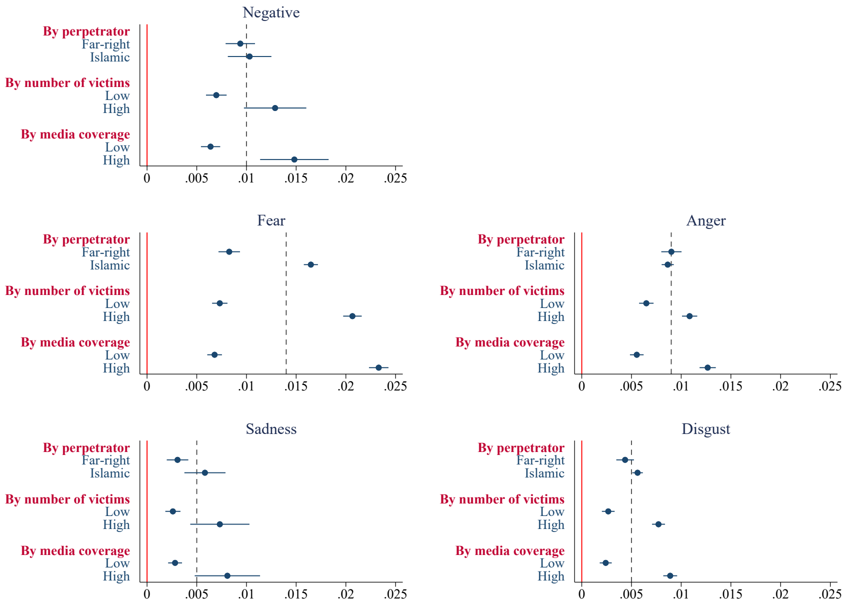

Notes: The figure shows the post-attack estimates for the 24-hour bandwidth (based on the specification in panel B of Table 1) when we run separate regressions for the attack groups displayed on the vertical axis. The horizontal lines signify the 95% confidence intervals of the corresponding estimates.

Fig. 3 The emotional content of immigration-related tweets

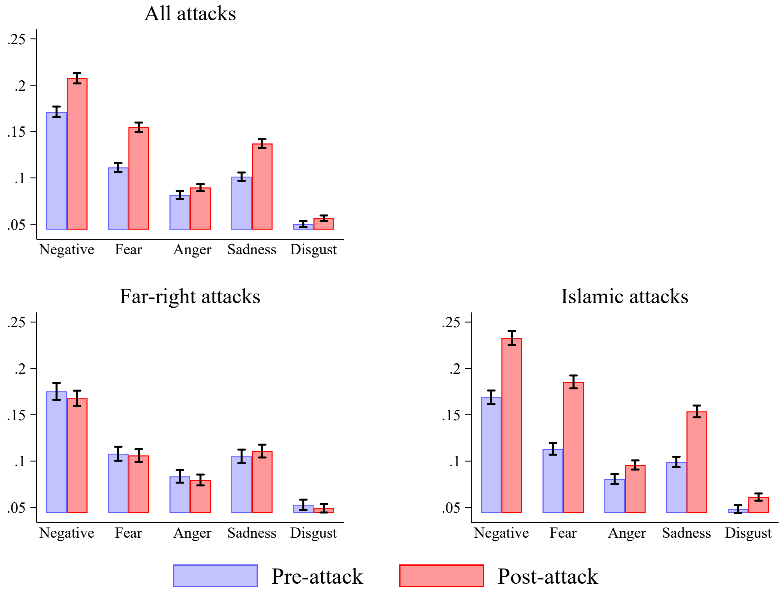

Notes: This figure shows the pre- and post-attack average values of the negative sentiment and emotions about immigration. The analysis is based on tweets that include the word 'immigration' and other related terms, as identified using a Word2Vec algorithm; i.e., migrant, deport, illegals, undocumented, refugee, citizenship, visa, illegal alien, expedited removal, asylum seeker, as well as typos of the word 'immigration'. Black bars denote the standard error of the mean.

Table 3 Terrorism and the emotional content of immigration-related tweets

	Negative (1)	Fear (2)	Anger (3)	Sadness (4)	Disgust (5)
Panel A: All attacks; includes attack FEs					
Post-attack	0.019***	0.027***	0.000	0.023***	0.001
	(0.005)	(0.005)	(0.004)	(0.004)	(0.003)
Post-attack × 24 hour bandwidth	0.069***	0.059***	0.029***	0.047***	0.019***
	(0.010)	(0.009)	(0.007)	(0.008)	(0.005)
Pre-attack dependent variable	0.172	0.112	0.082	0.102	0.050
Observations	15,979	15,979	15,979	15,979	15,979
Number of users	8,453	8,453	8,453	8,453	8,453
Panel B: Islamic attacks, includes attack FEs					
Post-attack	0.040***	0.051***	0.005	0.037***	0.006
	(0.007)	(0.007)	(0.005)	(0.006)	(0.004)
Post-attack × 24 hour bandwidth	0.089***	0.070***	0.038***	0.057***	0.022***
	(0.012)	(0.011)	(0.008)	(0.010)	(0.006)
Pre-attack dependent variable	0.169	0.113	0.081	0.099	0.048
Observations	9,849	9,849	9,849	9,849	9,849
Number of users	5,781	5,781	5,781	5,781	5,781
Panel C: All attacks; includes individual × attack FEs					
Post-attack	0.017	0.012	–0.007	0.009	–0.008
	(0.011)	(0.009)	(0.007)	(0.009)	(0.006)
Post-attack × 24 hour bandwidth	0.034*	0.038**	0.028**	0.016	0.029***
	(0.018)	(0.015)	(0.012)	(0.015)	(0.009)
Pre-attack dependent variable	0.177	0.113	0.084	0.102	0.049
Observations	8,085	8,085	8,085	8,085	8,085
Number of users	1,859	1,859	1,859	1,859	1,859
Panel D: Islamic attacks; includes individual × attack FEs					
Post-attack	0.031**	0.036***	–0.003	0.011	0.000
	(0.014)	(0.013)	(0.010)	(0.012)	(0.007)
Post-attack × 24 hour bandwidth	0.045**	0.037*	0.030*	0.028	0.024*
	(0.023)	(0.020)	(0.016)	(0.019)	(0.012)
Pre-attack dependent variable	0.179	0.117	0.082	0.104	0.044
Observations	4,545	4,545	4,545	4,545	4,545
Number of users	1,162	1,162	1,162	1,162	1,162
Hour FEs	✓	✓	✓	✓	✓
Tweet-level controls	✓	✓	✓	✓	✓

Notes: See notes for Table 1. The analysis is based on tweets that include the word 'immigration' and other related terms, as identified using a Word2Vec algorithm; i.e., migrant, deport, illegals, undocumented, refugee, citizenship, visa, illegal alien, expedited removal, asylum seeker, as well as typos of the word 'immigration'. Standard errors are clustered at the individual-level and reported in parentheses. * $p < .10$; ** $p < .05$; *** $p < .01$.

References

[1] Friedland, N., Merari, A.: The psychological impact of terrorism: A double-edged sword. Political Psychology, 591–604 (1985)

[2] Enders, W., Sandler, T., Gaibulloev, K.: Domestic versus transnational terrorism: Data, decomposition, and dynamics. Journal of Peace Research **48**(3), 319–337 (2011)

[3] Nussio, E., Böhmelt, T., Bove, V.: Do terrorists get the attention they want? Comparing effects of terrorism across Europe. Public Opinion Quarterly **85**(3), 900–912 (2021)

[4] Cohn, M.A., Fredrickson, B.L., Brown, S.L., Mikels, J.A., Conway, A.M.: Happiness unpacked: Positive emotions increase life satisfaction by building resilience. Emotion **9**(3), 361 (2009)

[5] Alexander, R., Aragón, O.R., Bookwala, J., Cherbuin, N., Gatt, J.M., Kahrilas, I.J., Kästner, N., Lawrence, A., Lowe, L., Morrison, R.G., *et al.*: The neuroscience of positive emotions and affect: Implications for cultivating happiness and wellbeing. Neuroscience & Biobehavioral Reviews **121**, 220–249 (2021)

[6] Whalley, M.G., Brewin, C.R.: Mental health following terrorist attacks. The British Journal of Psychiatry **190**(2), 94–96 (2007)

[7] Sønderskov, K.M., Dinesen, P.T., Hansen, B.T., Østergaard, S.D., Danckert, B.: Terrorism in the country of origin is linked to deterioration in the mental health of refugees. Nature Human Behaviour **5**(11), 1555–1561 (2021)

[8] Hansen, B.T., Østergaard, S.D., Sønderskov, K.M., Dinesen, P.T.: Increased incidence rate of trauma-and stressor-related disorders in Denmark after the September 11, 2001, terrorist attacks in the United States. American Journal of Epidemiology **184**(7), 494–500 (2016)

[9] Hansen, B.T., Dinesen, P.T., Østergaard, S.D.: Increased incidence rate of trauma-and stressor-related disorders in Denmark after the Breivik attacks in Norway. Epidemiology **28**(6), 906–909 (2017)

[10] Wenzlaff, R.M., Wegner, D.M., Roper, D.W.: Depression and mental control: The resurgence of unwanted negative thoughts. Journal of Personality and Social Psychology **55**(6), 882 (1988)

[11] Lerner, J.S., Keltner, D.: Beyond valence: Toward a model of emotion-specific influences on judgement and choice. Cognition & Emotion **14**(4), 473–493 (2000)

[12] Schwarz, N., Clore, G.L.: Feelings and phenomenal experiences (2007)

[13] Forgas, J.P.: Mood and judgment: The affect infusion model (AIM). Psychological Bulletin **117**(1), 39 (1995)

[14] Loewenstein, G., Lerner, J.S.: The role of affect in decision making (2003)

[15] Huddy, L., Feldman, S., Taber, C., Lahav, G.: Threat, anxiety, and support of antiterrorism policies. American Journal of Political Science **49**(3), 593–608 (2005)

[16] Skitka, L.J., Bauman, C.W., Aramovich, N.P., Morgan, G.S.: Confrontational and preventative policy responses to terrorism: Anger wants a fight and fear wants "them" to go away. Basic and Applied Social Psychology **28**(4), 375–384 (2006)

[17] Lambert, A.J., Scherer, L.D., Schott, J.P., Olson, K.R., Andrews, R.K., O'Brien, T.C., Zisser, A.R.: Rally effects, threat, and attitude change: An integrative approach to understanding the role of emotion. Journal of Personality and Social Psychology **98**(6), 886 (2010)

[18] Huddy, L., Feldman, S.: Americans respond politically to 9/11: Understanding the impact of the terrorist attacks and their aftermath. American Psychologist **66**(6), 455 (2011)

[19] Godefroidt, A.: How terrorism does (and does not) affect citizens' political attitudes: A meta-analysis. American Journal of Political Science (2022)

[20] Huddy, L., Feldman, S., Cassese, E.: Terrorism, anxiety, and war. Terrorism and Torture, 290 (2009)

[21] Pyszczynski, T., Abdollahi, A., Solomon, S., Greenberg, J., Cohen, F., Weise, D.: Mortality salience, martyrdom, and military might: The great Satan versus the axis of evil. Personality and Social Psychology Bulletin **32**(4), 525–537 (2006)

[22] Epifanio, M.: The politics of targeted and untargeted counterterrorist regulations. Terrorism and Political Violence **28**(4), 713–734 (2016)

[23] Epifanio, M., Giani, M., Ivandic, R.: Wait and see? Public opinion dynamics after terrorist attacks. The Journal of Politics, forthcoming (2022)

[24] Kramer, A.D., Guillory, J.E., Hancock, J.T.: Experimental evidence of massive-scale emotional contagion through social networks. Proceedings of the National Academy of Sciences **111**(24), 8788–8790 (2014)

[25] Fischer-Preßler, D., Schwemmer, C., Fischbach, K.: Collective sense-making in times of crisis: Connecting terror management theory with Twitter user reactions to the Berlin terrorist attack. Computers in Human Behavior **100**, 138–151 (2019)

[26] Giavazzi, F., Iglhaut, F., Lemoli, G., Rubera, G.: Terrorist attacks, cultural incidents and the vote for radical parties: Analyzing text from Twitter. American Journal of Political Science, DOI:10.1111/ajps.12764 (2023)

[27] De Choudhury, M., Gamon, M., Counts, S., Horvitz, E.: Predicting depression via social media. In: Proceedings of the International AAAI Conference on Web and Social Media, vol. 7, pp. 128–137 (2013)

[28] Ji, X., Chun, S.A., Wei, Z., Geller, J.: Twitter sentiment classification for measuring public health concerns. Social Network Analysis and Mining **5**, 1–25 (2015)

[29] Laranjo, L., Arguel, A., Neves, A.L., Gallagher, A.M., Kaplan, R., Mortimer, N., Mendes, G.A., Lau, A.Y.: The influence of social networking sites on health behavior change: A systematic review and meta-analysis. Journal of the American Medical Informatics Association **22**(1), 243–256 (2015)

[30] Guntuku, S.C., Yaden, D.B., Kern, M.L., Ungar, L.H., Eichstaedt, J.C.: Detecting depression and mental illness on social media: An integrative review. Current Opinion in Behavioral Sciences **18**, 43–49 (2017)

[31] Chen, X., Sykora, M.D., Jackson, T.W., Elayan, S.: What about mood swings: Identifying depression on Twitter with temporal measures of emotions. In: Companion Proceedings of the Web Conference 2018, pp. 1653–1660 (2018)

[32] Mayor, E., Bietti, L.M.: Twitter, time and emotions. Royal Society Open Science **8**(5), 201900 (2021)

[33] Eady, G., Hjorth, F., Thisted Dinesen, P.: Do violent protests affect expressions of party identity? Evidence from the Capitol insurrection. American Political Science Review, DOI:10.1017/S0003055422001058 (2022)

[34] Lansdall-Welfare, T., Dzogang, F., Cristianini, N.: Change-point analysis of the public mood in UK Twitter during the Brexit referendum. 2016 IEEE 16th International Conference on Data Mining Workshops (ICDMW), 434–439 (2016)

[35] Muñoz, J., Falcó-Gimeno, A., Hernández, E.: Unexpected event during

survey design: Promise and pitfalls for causal inference. Political Analysis **28**(2), 186–206 (2020)

[36] Bove, V., Efthyvoulou, G., Pickard, H.: Did terrorism affect voting in the Brexit referendum? British Journal of Political Science **52**(3), 1133–1150 (2022)

[37] Böhmelt, T., Bove, V., Nussio, E.: Can terrorism abroad influence migration attitudes at home? American Journal of Political Science **64**(3), 437–451 (2020)

[38] Breton, C., Eady, G.: Does international terrorism affect public attitudes toward Refugees? Evidence from a large-scale natural experiment. The Journal of Politics **84**(1), 554–559 (2022)

[39] Helbling, M., Meierrieks, D.: Terrorism and migration: An overview. British Journal of Political Science, 1–20 (2020)

[40] Mancosu, M., Ferrin Pereira, M.: Terrorist attacks, stereotyping, and attitudes toward immigrants: The case of the Manchester bombing. Social Science Quarterly **102**(1), 420–432 (2021)

[41] Pickard, H., Efthyvoulou, G., Bove, V.: What's left after right-wing extremism? the effects on political orientation. European Journal of Political Research **62**(1), 338–351 (2023)

[42] Doble, A., McGregor, N., de Gallier, T.: Are we becoming numb to terrorism from hearing about it too much. BBC Newsbeat (2016)

[43] Nussio, E.: Attitudinal and emotional consequences of Islamist terrorism. Evidence from the Berlin attack. Political Psychology **41**(6), 1151–1171 (2020)

[44] Powell, K.A.: Framing Islam: An analysis of US media coverage of terrorism since 9/11. Communication Studies **62**(1), 90–112 (2011)

[45] Meier, A.A.: The idea of terror: Institutional reproduction in government responses to political violence. International Studies Quarterly **64**(3), 499–509 (2020)

[46] Legewie, J.: Terrorist events and attitudes toward immigrants: A natural experiment. American Journal of Sociology **118**(5), 1199–1245 (2013)

[47] Hanlon, C., Jordans, M.: Guided self-help to reduce psychological distress: What's the potential for scaling up? The Lancet Global Health **8**(2), 165–166 (2020)

[48] Holmes, E.A., O'Connor, R.C., Perry, V.H., Tracey, I., Wessely, S., Arseneault, L., Ballard, C., Christensen, H., Silver, R.C., Everall, I., *et al.*: Multidisciplinary research priorities for the COVID-19 pandemic: A call for action for mental health science. The Lancet Psychiatry **7**(6), 547–560 (2020)

[49] Robben, A.C.: Massive violent death and contested national mourning in post-authoritarian Chile and Argentina: A sociocultural application of the dual process model. Death Studies **38**(5), 335–345 (2014)

[50] Zaller, J., Feldman, S.: A simple theory of the survey response: Answering questions versus revealing preferences. American Journal of Political Science **36**(3), 579–616 (1992)

[51] Curini, L., Iacus, S., Canova, L.: Measuring idiosyncratic happiness through the analysis of Twitter: An application to the Italian case. Social Indicators Research **121**, 525–542 (2015)

[52] Mohammad, S., Turney, P.: Emotions evoked by common words and phrases: Using Mechanical Turk to create an emotion lexicon. Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text (2010)

[53] Mohammad, S., Turney, P.: Computational intelligence. Crowdsourcing a Word-Emotion Association Lexicon **29**(3), 436–465 (2013)

[54] Skitka, L.J., Bauman, C.W., Mullen, E.: Political tolerance and coming to psychological closure following the September 11, 2001, terrorist attacks: An integrative approach. Personality and Social Psychology Bulletin **30**(6), 743–756 (2004)

[55] Müller, K., Schwarz, C.: From hashtag to hate crime: Twitter and anti-minority sentiment. American Economic Journal: Applied Economics, Forthcoming (2023)

[56] Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. arXiv preprint arXiv:1301.3781 (2013)

The emotional effect of terrorism: Evidence from Twitter data

Supplementary Information (SI) Appendix

For Online Publication

Contents

A Additional Information

A.1 Background material on the sampled attacks

We sample tweets that were posted around the timing of eight major terrorist incidents in the UK over the period 2016-2020: (i) the murder of MP Jo Cox in June 2016; (ii) the Westminster attack in March 2017; (iii) the Manchester Arena bombing in May 2017; (iv) the London Bridge attack in June 2017; (v) the Finsbury Park attack in June 2017; (vi) the Parsons Green bombing in September 2017; (vii) the London Bridge stabbings in November 2019; and (viii) the Reading stabbings in June 2020. Below we provide background material on these attacks.

(i) Murder of Jo Cox (June 2016)

Labour party MP Jo Cox was murdered in her constituency of Batley and Spen in Yorkshire, on June 16, 2016, a week prior to the Brexit referendum. The perpetrator, Thomas Mair, was a 53-years-old white supremacist, whose hatred extended to white people he deemed 'collaborators'. Mair's links to far-right movements and his obsession with Nazism, white supremacy and apartheid-era South Africa were well documented,[1] but the attack was not immediately identified as a terrorist incident. The media subsequently labelled Mair a 'far-right terrorist'. Though Mair was trialled for murder, the prosecutors argued that his crimes were "nothing less than acts of terrorism", while the judge noted in delivering Mair's life sentence that his "inspiration was not love of country but admiration for Nazism".[2]

The British public were deeply shocked: Union Flags on public buildings, including the Palace of Westminster, were flown at half mast, and the Brexit referendum campaign was suspended. Prominent UK leaders issued tributes and condolences, including the Conservative Prime Minister, David Cameron, the Labour party leader, Jeremy Corbyn, and MEP Nigel Farage, leader of the UK Independence Party (UKIP) and prominent Leave.EU campaigner.[3] Farage came under scrutiny as during the course of the attack, Mair shouted "Britain First", the name of a far-right organisation aligned with Farag's party and policies. Some commentators and scholars drew connections between the aggressive rhetoric oft-adopted by the Leave campaign, and Mair's motives (Jones, 2020; Bove et al., 2022a;

[1]Sources: voxpoliticalonline.com, splcenter.org, and theguardian.com.
[2]Source: bbc.co.uk.
[3]Source: theguardian.com.

Pickard et al., 2023). Farage rebutted these accusations, dismissing the murder as caused by 'one deranged, dangerous individual'.[4] Leaders of far-right groups – including Paul Golding, the leader of Britain First – rushed, at least initially, to distance themselves from the attack,[5] while extreme-right activists' reactions ranged from attributing the attack to Mair's mental-health issues to insinuating the 'truthfulness' of the attack.[6]

(ii) Westminster Attack (March 2017)

On 22 March 2017, 52-year-old Khalid Masood drove a car into pedestrians along Westminster Bridge in London. Masood abandoned the vehicle outside the Palace of Westminster and stabbed a police officer before being shot and killed by police forces.[7] The attack resulted in 5 deaths and at least 50 injuries. Masood was an Islamic extremist, claiming in a Whatsapp message uncovered by security forces that he was waging Jihad in response to Western military action in the Middle East.[8] On 23rd of March the Islamic State of Iraq and the Levant (ISIL), announced that the attacker was "a soldier of the Islamic State, executing the operation in response to calls to target citizens of coalition nations".[9] The Metropolitan Police, immediately designating the attack a terrorist incident, investigated these claims but found that Masood had acted alone, though he was inspired by ISIL rhetoric. Masood had previously been investigated as a peripheral figure in an MI5 investigation of a 2010 terror plot, but a risk assessment in that instance found that he posed no threat.[10]

The attack met with condemnation across the political spectrum. The prime-minister, Theresa May, attributed the attack to 'Islamic ideology', characterising it as an attempt to 'silence our democracy' while emphasising that Masood was not an active target of investigation by intelligence services despite his prior involvement in a terror plot.[11] Jeremy Corbyn, leader of the opposition, remarked in the Commons that the attack was an 'appalling atrocity'.[12]

[4]Source: mirror.co.uk.
[5]Source: huffingtonpost.co.uk.
[6]Source: globalcomment.com.
[7]Source: start.umd.edu/gtd.
[8]Source: telegraph.co.uk.
[9]Source: independent.co.uk.
[10]Source: telegraph.co.uk.
[11]Source: bbc.com.
[12]Ibid.

(iii) Manchester Arena Bombing (May 2017)

On 22 May 2017, Salman Ramadan Abedi, a 22-year-old Mancunian man of Libyan descent, detonated a home-made bomb in the foyer of Manchester Arena, as people were leaving an Ariana Grande concert. Twenty-three attendees – six of them children – died in the explosion,[13] and 1,017 were injured, 112 of whom required hospitalisation.[14] The attack proved the deadliest episode of terrorism in Britain since the London bombings of July 7, 2005.[15] The government immediately raised the terror threat level to 'critical', the highest level in a five-point scale, before reverting to the pre-existing level ('severe') five days later.[16]

The government pointed to the bombing as part of its motivation for updating its counter-terrorism strategy and introducing new policies such as broader data-sharing between counter-terrorism police and other agencies in 2018.[17] In 2022, testimonies in the Manchester Arena Inquiry reported that MI5 had sufficient intelligence to open an investigation against Abedi as a threat to national security a month prior to the attack. It had failed to do so because the agency was "struggling to cope" with increasing workload and could not carefully consider of the case.[18] British newspapers reported widely on the attack with graphic, emotional coverage. The Daily Mail shared 'horrifying videos' from inside the arena, as "terrified concert-goers flee for their lives",[19] while several newspapers emphasised the presence of children among the victims. PM Theresa May condemned how the attacker saw a "room packed with young children as an opportunity for carnage".[20] Data from LexisNexis confirms the relevance of the bombing in public debate: it is indeed the most widely covered attack in the 2016-2020 period, with over 11,900 results during the first month after the event.[21]

[13] Source: bbc.co.uk.
[14] Source: files.manchesterarenainquiry.org.uk.
[15] Source: kerslakearenareview.co.uk.
[16] Source: bbc.co.uk.
[17] Source: assets.publishing.service.gov.uk.
[18] Source: theguardian.com.
[19] Source: dailymail.co.uk.
[20] Source: gov.uk.
[21] Keywords ('terrorist' OR 'terrorism') AND 'Manchester' AND 'arena' were used to identify news reports about the attack.

(iv) London Bridge Attack (June 2017)

On the 3rd of June, 2017, a van was driven at pedestrians on London Bridge, then crashed on Borough High Street.[22] Three perpetrators, wearing fake suicide vests, exited the vehicle and began stabbing civilians around a cluster of restaurants and pubs along Stoney Street.[23] The assailants were shot dead by armed officers of the Metropolitan police Special Firearms Command. The attack resulted in 8 deaths, and 48 injuries.[24] In the days following the attack, numerous news stories emerged which vividly detailed how members of the public and a number of unarmed police officers had attempted to intervene.[25] On the 4th of June, ISIL claimed responsibility for the attack,[26] but no link between the assailants and the group could be confirmed by authorities.[27] The perpetrators were identified as Khuram Shazad Butt, a Pakistan-born British citizen, Rachid Redouane, a failed asylum seeker residing in Dagenham, and Youssef Zaghba, a Moroccan and Italian dual-national residing in East London.[28] Two of the assailants had been identified by various authorities as connected to Islamic extremism before the attacks.[29] Butt was a known member of the banned extremist group Al-Muhajiroun and was investigated in relation to his involvement with suspects involved in the July 2005 London bombing,[30] while Zaghba had been previously identified by the Italian authorities as a terror threat.[31]

National election campaigning was suspended by all political parties in the day following attack, with the controversial exception of UKIP whose leader, Paul Nuttall, claimed that suspending their campaign was "what the extremists would want".[32] Theresa May attributed the attack to 'evil Ideology of Islamic Extremism'.[33] The Mayor of London, Sadiq Khan praised London's 'defiant unity in the face of adversity' while also condemning the spike in hate crimes targeting Muslims in the wake of the attack.[34]

[22] Source: start.umd.edu/gtd.
[23] Ibid.
[24] Ibid.
[25] Source: news.sky.com.
[26] Source: gov.uk.
[27] Source: start.umd.edu/gtd.
[28] Source: theguardian.com.
[29] Source: nytimes.com.
[30] Ibid.
[31] Source: telegraph.co.uk.
[32] Source: theguardian.com.
[33] Source: gov.uk.
[34] Source: theguardian.com.

(v) Finsbury Park Attack (June 2017)

The second attack by a far-right perpetrator in our sample took place on June 19, 2017. A 48-year-old man, Darren Osborne, drove a van into a crowd of Muslims near the Finsbury Park Mosque, in north London, causing one death and injuring ten. Osborne was motivated by his anger over the Islamic attacks in London and Manchester, and a child grooming scandal in Rochdale involving men of Asian origin. The incident was immediately considered a terrorist attack by politicians, counter-terrorism police, and the media.

Most British newspapers' front pages on the following day focused on the perpetrator, rather than the victims, though there were some notable exceptions (e.g., the Guardian and the Independent). Importantly, Osborne's affiliation to far-right groups appeared secondary in reporting, in contrast with the pattern generally observed following Islamic attacks. Nonetheless, media coverage of the incident was rather high, with 1,152 results on LexisNexis within the first month.[35] The attack was generally condemned across the political spectrum: PM Theresa May praised London's multiculturalism and promised a stronger effort against Islamophobia, as did religious leaders from different creeds, while Prince Charles visited the Finsbury Park Mosque to meet community leaders.[36] Yet, some comments on the social media pages of far-right groups such as Britain First suggested that the attack was justified and painted the perpetrator as a 'hero'.[37]

(vi) Parsons Green Bombing (September 2017)

On September 15th, 2017, an explosion occurred on the District line train at Parsons Green Underground station in London, injuring 69.[38] The blast was caused by a homemade 'bucket bomb' packed with the explosive chemical triacetone triperoxide (TATP), which partially exploded.[39] The following day ISIL claimed responsibility for the attack, though the Metropolitan Police cast doubt on this claim.[40] An investigation led by the Metropolitan Police's Counter Terrorism Command immediately ensued,[41] culminating in the arrest of several individuals linked to Islamic extremism by the authorities.[42] Ahmed Hassan was eventually identified as the sole perpetrator, and was later sentenced to life imprison-

[35]Keywords used: ('terrorist' OR 'terrorism') AND 'Finsbury Park'.

[36]Sources: theguardian.com, independent.co.uk, and bbc.co.uk.

[37]Source: thetimes.co.uk.

[38]Source: start.umd.edu/gtd.

[39]Ibid.

[40]Source: bbc.co.uk.

[41]Ibid.

[42]Source: bbc.com.

ment on 23rd March 2018.[43] Hassan had entered the UK illegally via Calais, but successfully sought asylum, despite admitting in his asylum interview that he had spent three months in an ISIS training camp.[44] During the sentencing, it was concluded that Hassan was driven by "a mind-set of ISIS extremism, a deep-seated hatred of this country, a desire for revenge against Britain and America whom he blamed for his father's death in Iraq and anger at the continued bombing of Iraq by Western Coalition forces".[45]

Following the attack, PM Theresa May announced that the terror threat level would be raised to the highest level, 'critical'.[46] By the 17th of September, it was lowered to 'severe', the level previously designated during the months after the Manchester Arena attack.[47] The attack was widely condemned by actors across the political spectrum. Theresa May branded the incident a 'cowardly attack' while the Mayor of London, Sadiq Khan emphasised that London "will never be intimidated or defeated by terrorism".[48]

(vii) London Bridge Stabbings (November 2019)

On 29 November 2019, Usman Khan, a former British prisoner of Pakistani descent convicted of terror offences, stabbed five people inside and outside Fishmongers' Hall, adjacent to London Bridge. Two of the victims died from their stab wounds.[49] Khan, released on license just one year prior on the day of the attack,[50] was attending a conference on offender rehabilitation.[51] After initially threatening to detonate what turned out to be a fake suicide vest, he began stabbing people in the building.[52] Khan then ran outside and stabbed pedestrians on London Bridge, where a civilian eventually managed to restrain him, until the police arrived and shot him dead.[53] ISIL claimed responsibility for the attack, without evidence.[54] In 2021, Independent Reviewer of Terrorism Legislation Jonathan Hall QC, considering Khan's early release, recommended that those who participate in the planning or preparation of terrorist attacks are given automatic life sentences.[55]

---

[43]Source: judiciary.uk.
[44]Ibid.
[45]Ibid.
[46]Source: bbc.co.uk.
[47]Source: london.gov.uk.
[48]Ibid.
[49]Source: thetimes.co.uk.
[50]Ibid.
[51]Source: bbc.co.uk.
[52]Ibid.
[53]Source: bbc.co.uk.
[54]Source: washingtonpost.com.
[55]Source: richmondandtwickenhamtimes.co.uk.

While investigators concluded that police had lawfully killed Khan,[56] a separate inquiry found that the attacker had not been sufficiently monitored and that the security planning at the event had been sub-par. These factors, the jury concluded, contributed to the death of the two victims.[57]

(viii) Reading Stabbings (June 2020)

On June 20th, 2020, a single assailant with a kitchen knife attacked two groups of civilians at Forbury Park, a public park in the centre of Reading, England.[58] The perpetrator, 25-year-old Khairi Saadallah, was tackled by police called to the scene.[59] The attack resulted in three fatalities, and three serious injuries. Though initial police statements suggested that the motivation for the attacks was unknown,[60] Counter-Terrorism Policing South East, who took over the investigation of the incident in conjunction with MI5, confirmed that the attack was being treated as a 'terror incident'.[61] Investigations confirmed that Saadallah was an Islamic extremist, inspired by ISIL.[62] Saadallah claimed to police that the attack was 'jihad'.[63] It was later uncovered that Saadallah, who had successfully claimed asylum in the UK in after escaping the Libyan Civil War, was a known quantity to MI5, which had obtained evidence that he planned to travel for extremist reasons in 2019.[64] Before the attack, Saadallah was convicted on six occasions for 15 crimes, 8 of which were violent crimes. It was found by Westminster Magistrates court that in Libya, Saadallah had trained and fought for extremist group Ansar Al-Sharia in Libya.[65]

PM Boris Johnson was 'appalled and sickened' by the incident, and hinted at possible legislative action, stating that "if there are lessons that we need to learn about how we handle such case, we will not hesitate to take action where necessary".[66] One of the victim's families subsequently criticised the government for failing to deport Saadallah prior to the attack, despite his violent crime convictions.[67]

[56] Source: bbc.co.uk.
[57] Source: standard.co.uk.
[58] Source: start.umd.edu.
[59] Source: independent.co.uk.
[60] Source: thamesvalley.police.uk.
[61] Source: independent.co.uk.
[62] Source: telegraph.co.uk.
[63] Ibid.
[64] Ibid.
[65] Source: news.sky.com.
[66] Source: theguardian.com.
[67] Source: getreading.co.uk.

A.2 Description of Twitter data

We use Twitter API v2 to retrieve data on tweets posted three days before, the same day, and three days after each attack. Focusing on a short-range time window around the attacks allows us to minimise the possibility of other events driving the estimated effects and draw robust causal inferences (Muñoz et al., 2020). In addition, tweet extraction and data processing is very time-consuming – as well as subject to the cap of 10 million tweets per month as part of API v2 – which prevents us from considering a wider time frame, especially in the context of multiple attacks.

We retrieve tweets based on following criteria: (i) they are written in English; and (ii) they include a place or location based in the UK which is tagged by the user. Twitter geographical information comes from two sources: tweet-level geographic metadata and account-level geographic metadata. While account-level geographic data constitutes a substantial 30-40% of all tweets, it is prone to bias as the location is selected by the user; e.g., users can specify incorrect locations (for instance, the information might not be up to date) or fictional locations (for instance, Atlantis). Tweet-specific location information, on the other hand, is retrieved from the tweet (geo-tagged tweets), and can be either specific (based on GPS data), or a Twitter "Place"[68] – which includes the area from which the tweet is posted, or a specific place in this area, which is selected by the user. The most precise location data is determined from geo-tagged tweets, which constitute 1-2% of all tweets. The majority of tweets or profiles do not have any geographical metadata specified, albeit there are various methods and techniques for location prediction (see, e.g., Zheng et al. (2018) for a review). We use geo-tagged tweets that are tagged within the UK, which generates a very large number of tweets (more than 9 million) with precise locations. The corresponding information includes the name of the location, the type of place (country, government administrative unit, city, and point of interest), as well as the bounding box of the tweet. 85.3% of the tweets in our sample have a city as a location, 10.8% have a government administrative unit, 0.7% have a point of interest, 0.2% have a country, and 3.1% do not have a specific location. In addition to the tweet content, the time and location, we also extract data on retweet count, reply count, like count, quote count, and total number of tweets per user. We focus on Twitter users with both pre- and post-attack tweets for at least one of the sampled terrorist incidents, which brings the total number of tweets used in our analysis to 7,643,102.

We use the NRC Emotion Lexicon (EmoLex) (Mohammad and Turney, 2013, 2010) to

[68]These tweets constitute around 80% of all geo-tagged tweets.

measure the sentiment and emotions of tweets. The lexicon contains 14,182 words and 25,000 senses, and each one of these words/senses is linked to two sentiments (negative and positive) and eight emotions (anger, fear, sadness, disgust, anticipation, trust, surprise, and joy). The sentiments are assigned either a value 1 (associated) or a value 0 (not associated); whereas the emotions are assigned a value from 0 to 1, capturing the share of lexicon-identified words/senses in a tweet that are linked to a given emotion. Figure A.2 shows the number and visual proportions of words (in our dataset) associated with each sentiment and emotion.

Figure A.2: Number of words linked to each sentiment and emotion

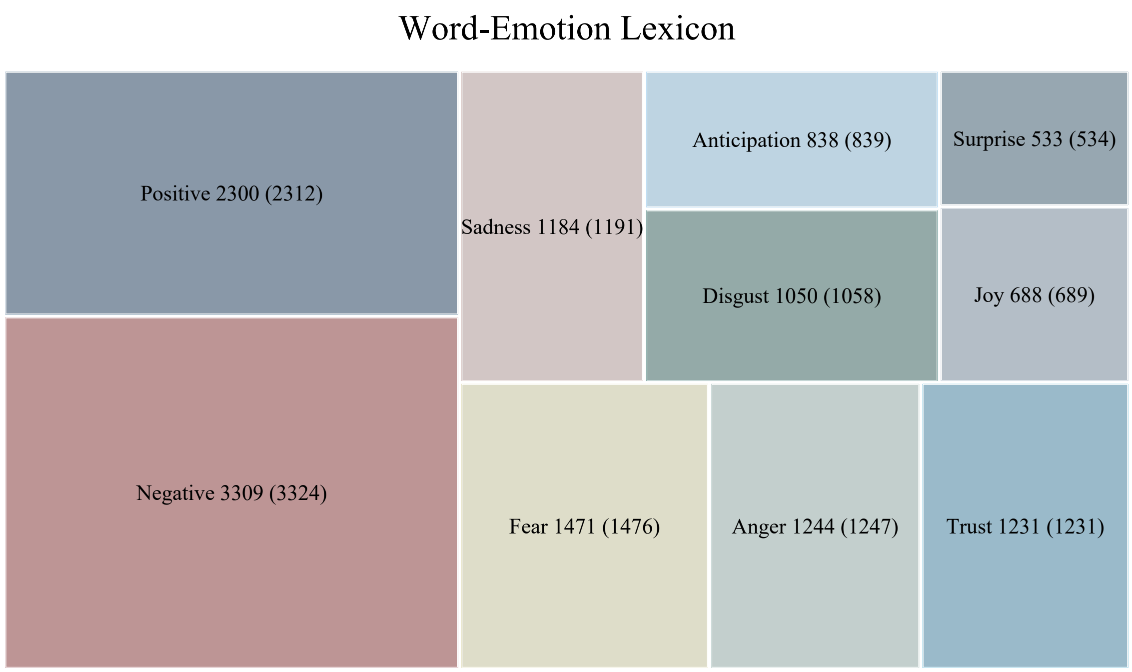

Notes: This figure shows the number and visual proportions of words (in our dataset) associated with each sentiment and emotion. The total number of words in the EmoLex lexicon associated with the corresponding sentiment or emotion is reported in parenthesis.

A.3 Additional tables and figures

- Table A3.a provides descriptive statistics of the variables used in our analysis for the full sample, the treated (post-attack) sample and the control (pre-attack) sample. It also presents the results of t-tests for differences in means across the pre- and post-attack groups.

- Table A3.b examines the sensitivity of the baseline results to alternative model specifications. We adopt an 'incremental strategy', where we start from a simple specification that includes the treatment (post-atack) variable and individual $\times$ attack fixed effects, and we then add hour fixed effects, the tweet-level controls and individual-level error clustering in a progressive manner, until we reach the full specification of panel A in Table 1. As shown in columns (1)-(4), the treatment effects for all outcome variables retain the size and statistical significance throughout these specifications. Finally, in column (5), we can see how the results change once we introduce an interaction term with a 24-hour bandwidth, as in panel B of Table 1.

- Figure A.3 presents an extended version of the time-to-event analysis in Figure 1 based on the full time window (from 3 days before to 3 days after the attacks). This rejects, once again, the presence of pre-existing patterns: the emotional content of tweets posted 1-93 hours before the attacks is very similar to that of tweets posted 1 hour before the attacks (the baseline hour). The figure also confirms that the heightening of negative feelings in the aftermath of the attacks lasts for about 24 hours. This can arguably capture the *direct* effect of terrorism violence on people's emotional state; i.e., before they are exposed to subsequent (related) activities and communication, or other unrelated events.

- Table A3.c shows the attack characteristics and the classification used for the heterogeneity analysis in Section 2.

Table A3.a: Descriptive statistics and balancing tests

	Full sample mean (sd)	Pre-attack mean (sd)	Post-attack mean (sd)	Difference (3)-(2) p-value
	(1)	(2)	(3)	(4)
Negative	0.12	0.11	0.12	
	(0.24)	(0.24)	(0.24)	0.00
Fear	0.06	0.06	0.06	
	(0.18)	(0.17)	(0.18)	0.00
Anger	0.06	0.06	0.06	
	(0.18)	(0.18)	(0.18)	0.00
Sadness	0.06	0.06	0.06	
	(0.17)	(0.17)	(0.18)	0.00
Disgust	0.04	0.04	0.04	
	(0.15)	(0.15)	(0.15)	0.00
Retweet count	0.74	0.72	0.76	
	(46.29)	(36.49)	(54.33)	0.16
Reply count	0.37	0.37	0.37	
	(3.70)	(2.88)	(4.36)	0.23
Like count	2.79	2.65	2.92	
	(146.31)	(93.95)	(184.25)	0.01
Quote count	0.07	0.06	0.07	
	(4.42)	(4.63)	(4.20)	0.22
Number of tweets	76.52	76.53	76.51	
	(379.21)	(378.21)	(380.21)	0.95
Observations	8,079,246	4,033,469	4,045,777	8,079,246

Notes: This table shows the mean values and standard deviations (in parentheses) of the variables used in our analysis, as well as the results of t-tests for differences in means across the pre- and post-attack groups.

Table A3.b: The emotional effect of terrorism: alternative model specifications

	(1)	(2)	(3)	(4)	(5)
Panel A: Negative					
Post-attack	0.003***	0.003***	0.003***	0.003***	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
24-hour bandwidth					–0.000
					(0.000)
Post attack × 24-hour bandwidth					0.010***
					(0.001)
Panel B: Fear					
Post-attack	0.006***	0.006***	0.006***	0.006***	0.002***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
24-hour bandwidth					–0.000
					(0.000)
Post-attack × 24-hour bandwidth					0.012***
					(0.000)
Panel C: Anger					
Post-attack	0.004***	0.003***	0.003***	0.003***	0.001***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
24-hour bandwidth					0.000
					(0.000)
Post-attack × 24-hour bandwidth					0.008***
					(0.000)
Panel D: Sadness					
Post-attack	0.002***	0.002***	0.002***	0.002***	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
24-hour bandwidth					0.000
					(0.000)
Post-attack × 24-hour bandwidth					0.005***
					(0.001)
Panel E: Disgust					
Post-attack	0.002***	0.002***	0.002***	0.002***	0.001***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
24-hour bandwidth					0.000
					(0.000)
Post-attack × 24-hour bandwidth					0.004***
					(0.000)
Individual × attack FEs	✓	✓	✓	✓	✓
Hour FEs		✓	✓	✓	✓
Tweet-level controls			✓	✓	✓
Error clustering at individual level				✓	✓

Notes: Time window: 3 days before, the same day, and 3 days after each attack. The tweets are aggregated at the minute level. *Post-attack* is a binary variable that takes value 1 if the tweet was posted after the minute of attack, and 0 otherwise. *24-hour bandwidth* is a binary variable capturing the 24 hours before and the 24 hours after each attack. Standard errors are clustered at the individual-level and reported in parentheses. * $p < .10$; ** $p < .05$; *** $p < .01$.

Figure A.3: The effect of the attack for 3 days post attack (extended)

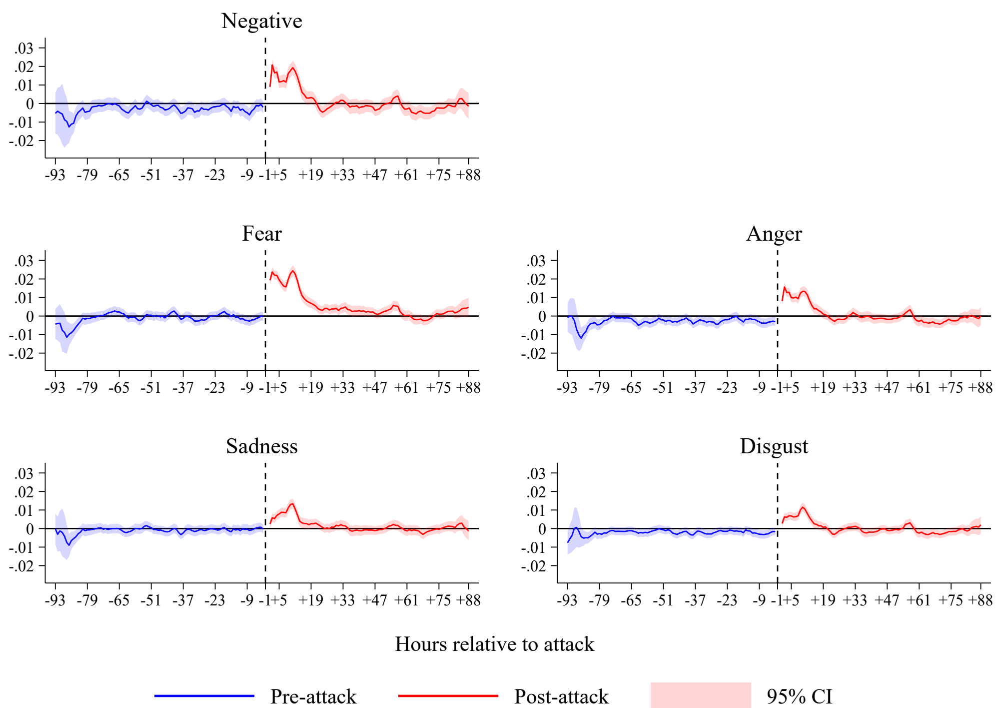

Notes: The figure shows the evolution of negative feelings from 3 days before to 3 days after the sampled attacks. The tweets are aggregated at the hour level. The blue (red) solid line shows the 3-hour moving average estimates before (after) the attacks, taking the hour before the attack as the baseline. The tweets posted in the hour after the attack are dropped from the estimations. The shaded areas show the 95 percent confidence intervals.

Table A3.c: Attack characteristics

	Perpetrator type (1)	Number of victims (2)	Number of newspaper articles (3)
Murder of MP Jo Cox in June 2016	Far-right	2	68
Westminster attack in March 2017	Islamic	55 †	4300 †
Manchester Arena bombing in May 2017	Islamic	141 †	11906 †
London Bridge attack in June 2017	Islamic	56 †	4551 †
Finsbury Park attack in June 2017	Far-right	12	3474 †
Parsons Green bombing in September 2017	Islamic	69 †	1282
London Bridge stabbings in November 2019	Islamic	5	2692
Reading stabbings in June 2020	Islamic	6	472

Notes: † indicates that the attack is classified as a high-victim / high-coverage attack.

B Robustness Tests and Further Insights

B.1 Using restricted samples of Twitter users

A possible threat to our identification arises from the possibility of selection into tweeting or changing the topic and the frequency of posts in the aftermath of a terrorist incident. An important reason why this threat is less acute in our context is that we exploit variation within individuals who have at least one tweet both before and after an attack. Furthermore, as shown in Section 2, our results hold when we run separate regressions for tweets that contain terror-related terms, and those that do not contain such terms. To further address this concern, we perform two additional checks. First, we restrict the sample to include the users who are present in our dataset before and after all sampled attacks. Second, we only keep the users who posted the same number of tweets before and after a given attack. Table B.1 reports the results when we run the same regression set-up as in panel A of Table 1 using these two 'restricted' samples. Overall, our inferences do not change: once again, we find that the tweets posted 24 hours after the attacks convey more negative feelings than those posted 24 hours before the attacks, and that fear is the emotion that displays the largest and more persistent post-attack rise.

Table B.1: The emotional effect of terrorism: using restricted samples of Twitter users

	Negative (1)	Fear (2)	Anger (3)	Sadness (4)	Disgust (5)
Panel A: Users present before and after all attacks					
Post-attack	0.001	0.003***	0.001	0.000	0.001
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Post-attack × 24-hour bandwidth	0.007***	0.011***	0.005***	0.003*	0.002*
	(0.002)	(0.002)	(0.002)	(0.002)	(0.001)
Pre-attack dependent variable	0.112	0.058	0.051	0.061	0.036
Observations	259,423	259,423	259,423	259,423	259,423
Number of users	980	980	980	980	980
Panel B: Users with the same number of tweets before and after a given attack					
Post-attack	0.001	0.001**	0.001	–0.000	0.001*
	(0.001)	(0.001)	(0.001)	(0.001)	(0.000)
Post-attack × 24-hour bandwidth	0.012***	0.014***	0.009***	0.008***	0.006***
	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)
Pre-attack dependent variable (mean)	0.100	0.052	0.047	0.054	0.036
Observations	546,568	546,568	546,568	546,568	546,568
Number of users	119,486	119,486	119,486	119,486	119,486
Individual × attack FEs	✓	✓	✓	✓	✓
Hour FEs	✓	✓	✓	✓	✓
Tweet-level controls	✓	✓	✓	✓	✓

Notes: See notes for Table 1 and Table A3.b. Standard errors are clustered at the individual-level and reported in parentheses. * $p < .10$; ** $p < .05$; *** $p < .01$.

B.2 Further identification validity tests

In this section, we perform a number of additional tests to strengthen our causal inference.

- First, we consider a placebo treatment during the pre-attack period. More precisely, we focus on one of the most important attacks in our sample (the 2017 Westminster attack) and set the placebo attack date to be 1 week prior to the actual date. Figure B.2a displays the evolution of negative feelings 24 hours before and 24 hours after the placebo attack, based on the same time-to-event analysis as in Figure 1. The patterns show no leaps after the simulated attack, while the occasional peaks before and after the 'event' are short-lived and are characterised by wide confidence intervals.

- Second, we benchmark our baseline results against a failed and not immediately reported attack. To do so, we perform the same analysis as before but we now compare the tweets posted around the 2017 assassination attempt of PM Theresa May, which was confirmed by the media one week later.[69] As can be seen in Figure B.2b, the treatment effects are close to zero and statistically insignificant throughout.

- Third, we examine the treatment effect on outcomes that should not be affected by terrorist incidents; namely, people's feelings about the weather. To do that, we create a sample of weather-related tweets (tweets containing the word 'weather') and compare their emotional content around the eight sampled attacks. Figure B.2c presents the evolution of negative feelings about the weather 24 hours before and 24 hours after the attacks. As expected, there is no evidence of significant spikes in the aftermath of the incidents.

- Finally, to ensure that the baseline estimates are unlikely to be observed by chance, we perform Monte Carlo permutation tests that randomly shuffle the data 500 times and estimate a treatment effect for each random draw. The resulting distributions are displayed in Figure B.2d. In all cases, the permuted data produce estimates which are lower than those reported in Table 1 (panel A), suggesting that there is 0% probability that the observed treatment effects are observed by chance.

[69]On the 28th November 2017, an Islamist extremist, Naa'imur Zakariyah Rahman, planned to bomb the gates of 10 Downing Street, kill guards and then attack Theresa May with a knife or gun. The suspect was arrested in London after collecting a suicide vest and a fake bomb from undercover operatives.

Figure B.2a: Placebo test based on an earlier cut-off point

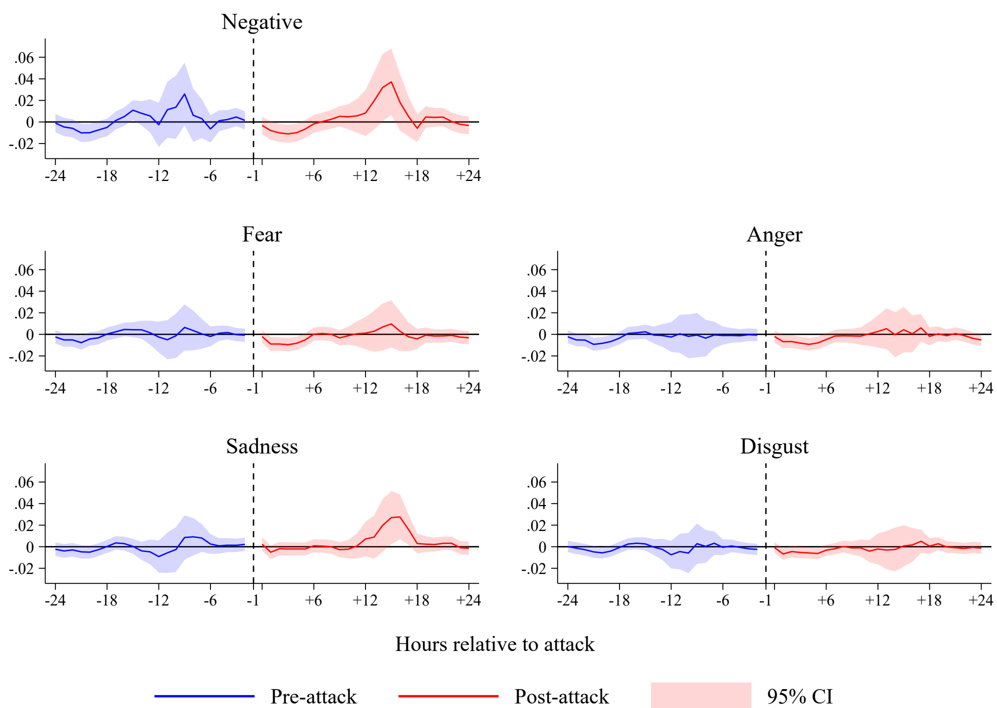

Notes: The figure shows the evolution of negative feelings 24 hours before and 24 hours after the placebo attack. To plot this figure, we rely on a sample of 1,095,903 tweets posted 3 days before and 3 days after the placebo attack date (i.e., 1 week before the 2017 Westminster attack). The tweets are aggregated at the hour level. The blue (red) solid line shows the 3-hour moving average estimates before (after) the attacks. The shaded areas show the 95 percent confidence intervals.

Figure B.2b: Comparison with a failed attack

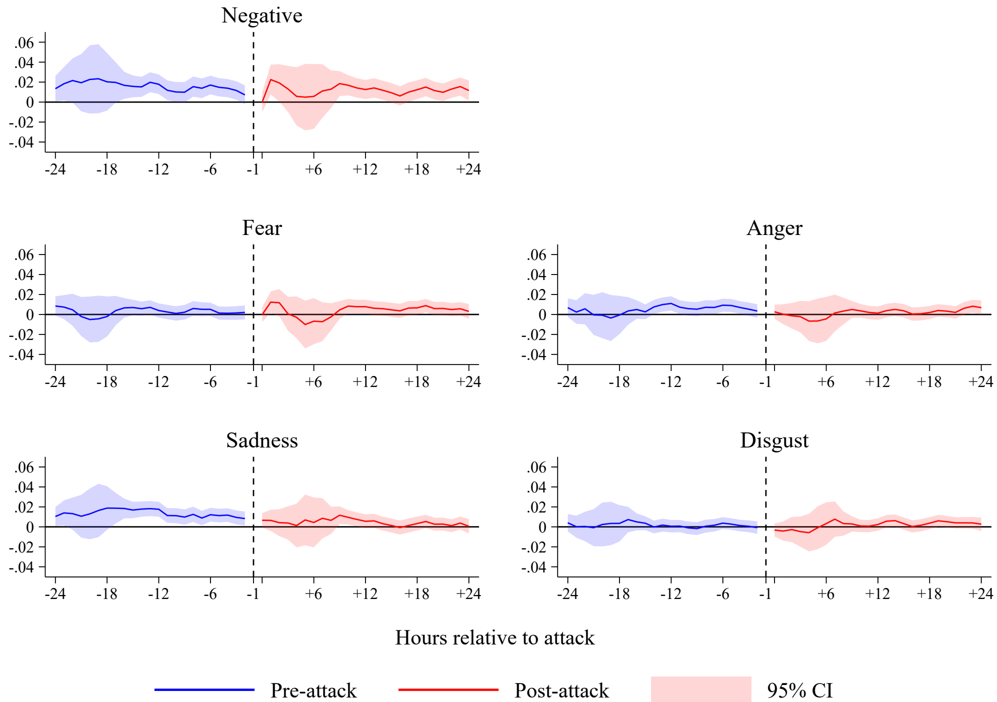

Notes: The figure shows the evolution of negative feelings 24 hours before and 24 hours after the 2017 assassination attempt of PM Theresa May. To plot this figure, we rely on 867,551 tweets posted three days before and three days after the failed attack. We omit the tweets posted on day of the attack, as the time of the attempted assassination and the time of the arrest were not reported. The tweets are aggregated at the hour level. The blue (red) solid line shows the 3-hour moving average estimates before (after) the attacks. The shaded areas show the 95 percent confidence intervals.

Figure B.2c: Placebo test based on unrelated outcomes: feelings about the weather

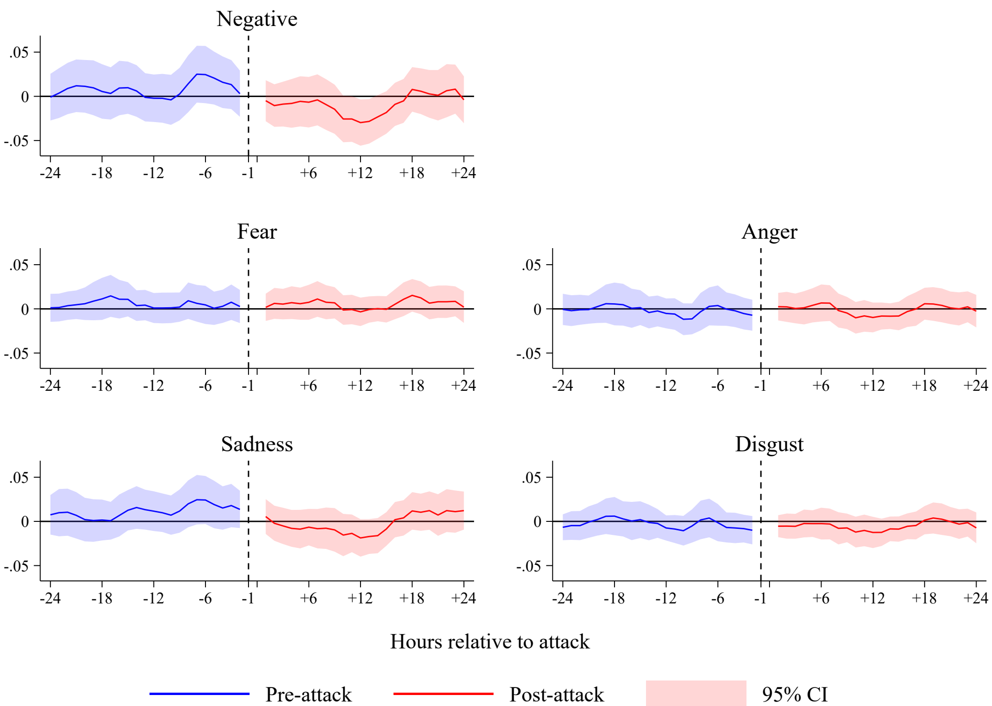

Notes: The figure shows the evolution of negative feelings about the weather 24 hours before and 24 hours after the sampled attacks. To plot this figure, we rely on a sample of 67,940 weather-related tweets posted from 3 days before up to and including 3 days after the attacks. The tweets are aggregated at the hour level. The blue (red) solid line shows the 3-hour moving average estimates before (after) the attacks, taking the hour before the attack as the baseline. The tweets posted in the hour after the attack are dropped from the estimations. The shaded areas show the 95 percent confidence intervals.

Figure B.2d: Permutation effect estimates

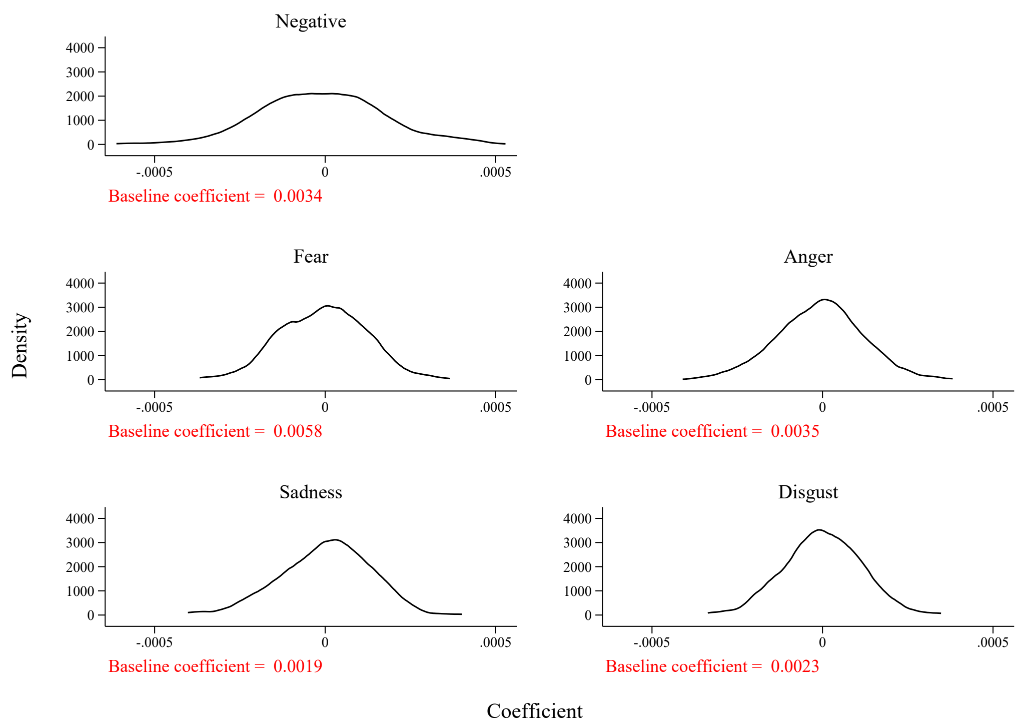

Notes: The figure shows the results of Monte Carlo permutation tests that randomly shuffle the data 500 times and estimate a treatment effect for each random draw. In all cases, the permuted data produce estimates which are lower than those reported in Table 1 (panel A).

B.3 The impact of terrorism on positive feelings

In this section, we examine how positive feelings respond to terrorist attacks. To do so, we carry out the same time-to-event analysis as in Figure 1, but we now focus on the overall positive sentiment. The patterns are displayed in Figure B.3. Generally speaking, we observe the opposite patterns to those of the overall negative sentiment – though the corresponding effects appear to be smaller in magnitude and shorter-lived. This is most likely due to the overall negative effect being counterbalanced by sympathy and compassion towards the attack victims.

Figure B.3: The emotional effect of terrorism: positive sentiment

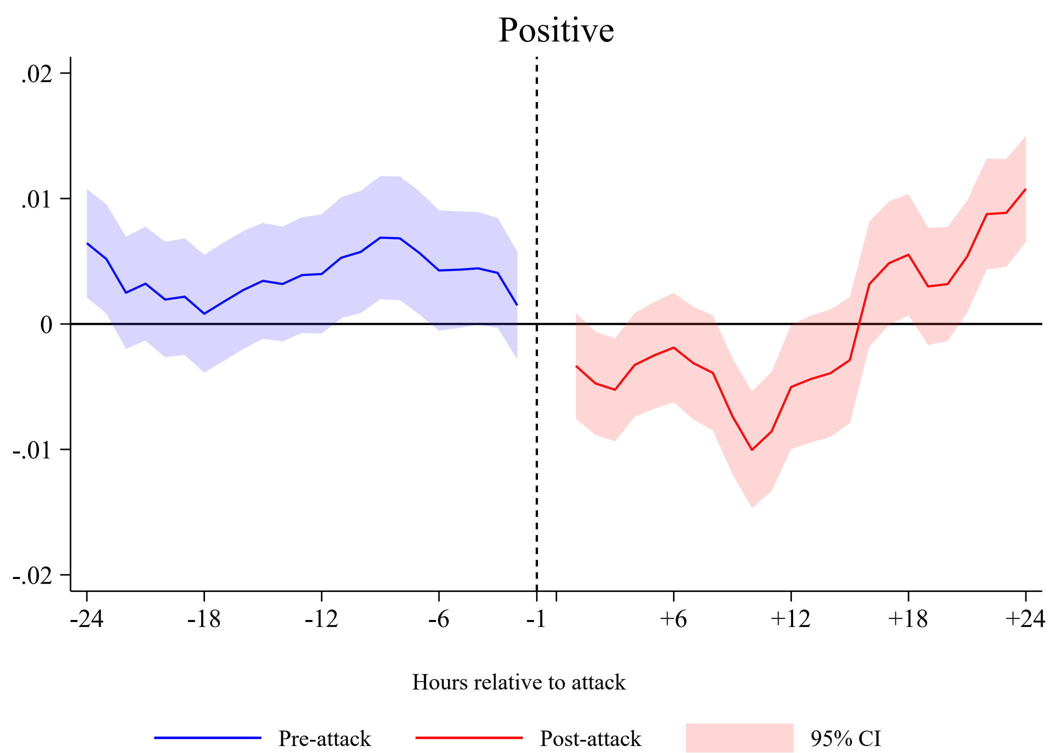

Notes: The figure shows the evolution of positive feelings 24 hours before and 24 hours after the sampled attacks. The tweets are aggregated at the hour level. The blue (red) solid line shows the 3-hour moving average estimates before (after) the attacks, taking the hour before the attack as the baseline. The tweets posted in the hour after the attack are dropped from the estimations. The shaded areas show the 95 percent confidence intervals.

B.4 Alternative approaches of measuring sentiment

So far, we have measured the valence and emotional content of the text contained in tweets using the NRC Emotion Lexicon (EmoLex). In this section, we test the sensitivity of our results to using alternative methods.

First, we replicate our analysis using VADER and Textblob. VADER (Hutto and Gilbert, 2014) is a lexicon- and rule-based sentiment analysis tool, and it is specifically attuned to sentiments expressed in social media. It returns negative, neutral and positive scores as the proportion of text that falls in each category, and a compound score that is computed by summing the corresponding scores of each word in the lexicon, adjusted according to the VADER grammatical and syntactical rules, and then normalised to be in the range between -1 (extreme negative) and +1 (extreme positive). Textblob (Loria et al., 2018) is another lexicon-based sentiment analysis tool, and its analyser returns the sentiment in the form of polarity and subjectivity scores. We consider the polarity score, which falls within the range [-1.0, +1.0], where -1 signifies negative sentiment and +1 positive sentiment. Second, we consider a composite index based on EmoLex. This is calculated by the difference between positive and negative sentiments, weighted by the ratio of the number of words in the tweet that are present in EmoLex to the total number of words in the tweet (LexRatio).

Figure B.4a shows the results of time-to-event analysis for the three VADER scores (positive, negative and neutral); whereas Figure B.4b shows the corresponding results for three alternative composite indices (EmoLex-based index, VADER compound score, and TextBlob polarity score), all capturing a net positive score. Across these figures, there is a sharp drop in neutral, positive, and net positive scores (or a sharp increase in the negative score) just right after the attacks, which persists for about 12 hours and is then followed by a gradual return to baseline levels within the next 12 hours. Overall, the patterns observed are in line with our previous findings and do not seem to be influenced by the method we use to measure sentiment.

Figure B.4a: The emotional effect of terrorism: VADER scores

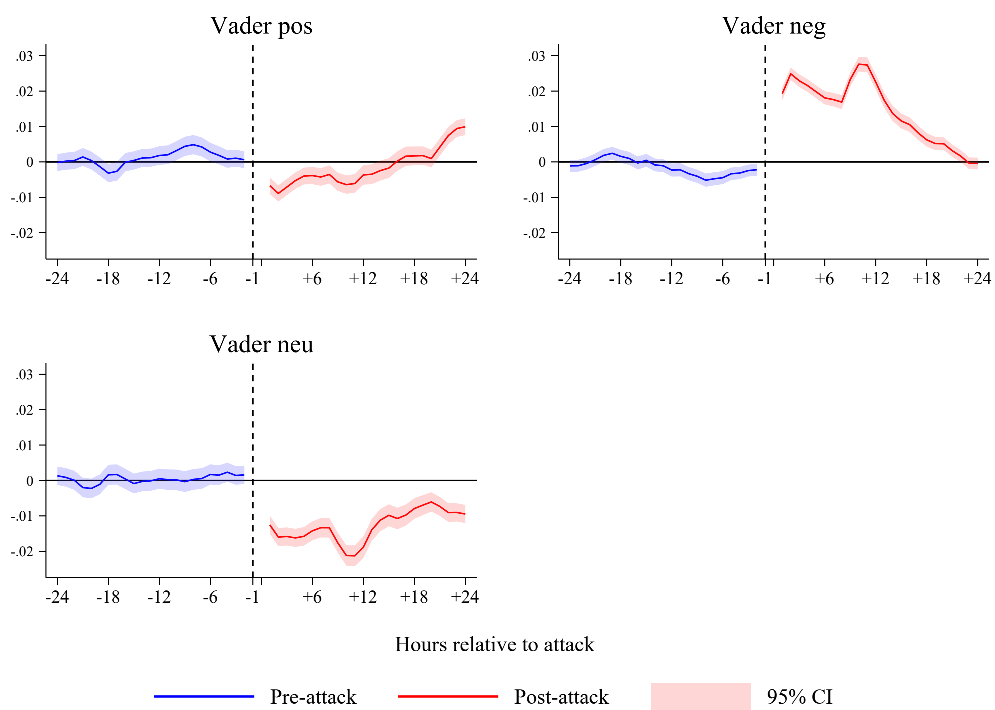

Notes: The figure shows the evolution of the VADER scores (positive, negative and neutral) 24 hours before and 24 hours after the sampled attacks. The tweets are aggregated at the hour level. The blue (red) solid line shows the 3-hour moving average estimates before (after) the attacks, taking the hour before the attack as the baseline. The tweets posted in the hour after the attack are dropped from the estimations. The shaded areas show the 95 percent confidence intervals.

Figure B.4b: The emotional effect of terrorism:
alternative composite (net positive) indices

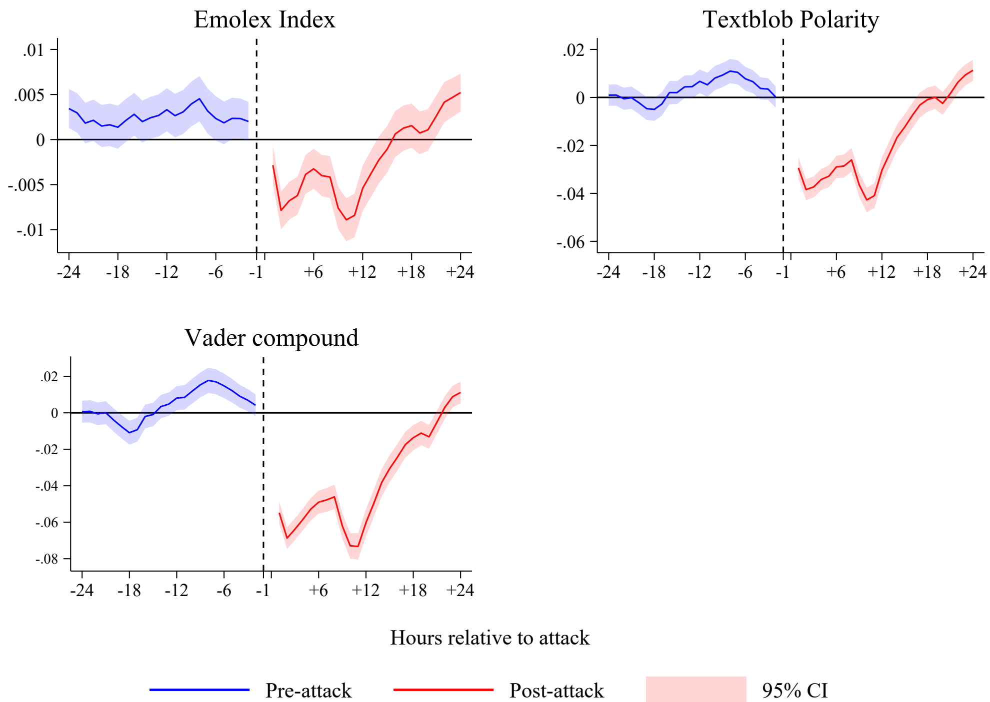

Notes: The figure shows the evolution of three alternative composite indices (all capturing a net positive score) 24 hours before and 24 hours after the sampled attacks. The tweets are aggregated at the hour level. The blue (red) solid line shows the 3-hour moving average estimates before (after) the attacks, taking the hour before the attack as the baseline. The tweets posted in the hour after the attack are dropped from the estimations. The shaded areas show the 95 percent confidence intervals.

B.5 The moderating role of geographic proximity

Geographic proximity to a terrorist incident can amplify the perception of threat and the personal sense of vulnerability, increase mortality salience as individuals feel more connected to the environment where the attack occurred, and affect the extent to which the event is covered by the local media (Nussio et al., 2021; Bove et al., 2022b). Following these arguments, one would expect that proximity to terrorism will act as a moderating factor whereby individuals that reside closer an attack are more likely to exhibit negative emotions. To test for this, we interact our treatment variable (*Post-attack*) with the physical distance between the user's geo-tagged location and the attack location. We normalise the distance measure by splitting it into decile groups, where individuals in group 10 are the most proximate to the attack and those in group 1 are the furthest away. Using the estimates from the model with the interaction term and relying on the 24-hour window before and after the attacks, we calculate the margins of the *Post-attack* variable and plot them over the respective decile values of proximity.

Figure B.5 reports the margins for the negative sentiment and the four negative emotions. The results indicate that the closer the individual's geo-tagged location is from the attack, the stronger the effect is on the outcome variables, which verifies the moderating role of geographic proximity in how individuals respond to terrorism. It should be acknowledged, however, that the estimated effect is positive across all values of proximity and only fails to reach statistical significance when we consider the lowest decile groups for sadness – which points to a rather weak dependence on proximity. This is likely due to the severity and emblematic nature of the attacks in our sample – see also Pickard et al. (2023) for a similar finding.

Figure B.5: The moderating role of geographic proximity

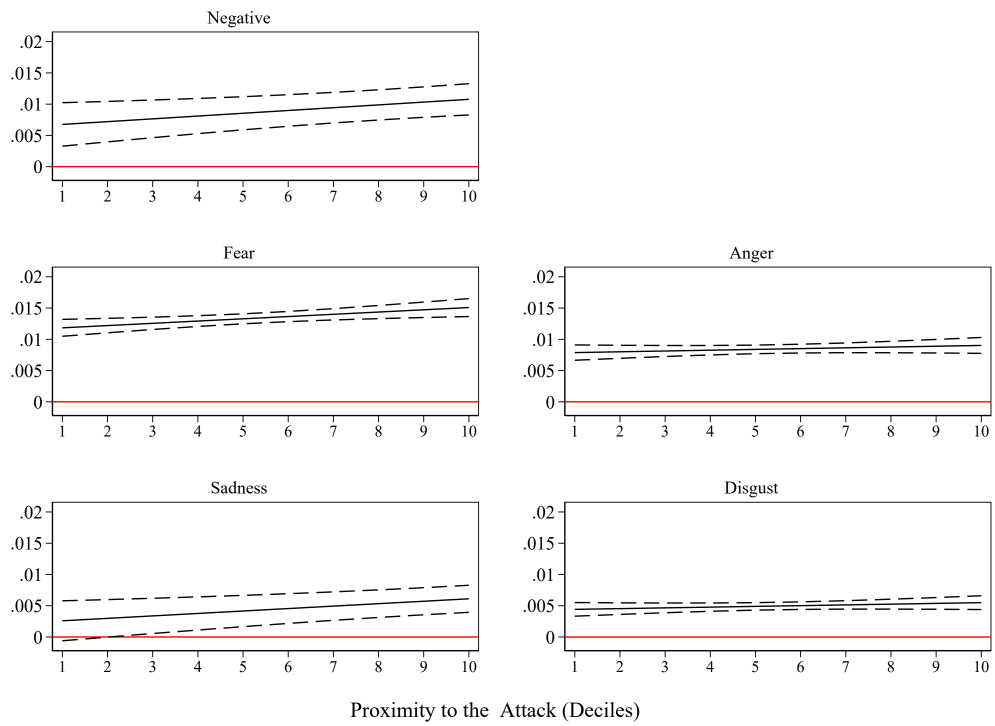

Notes: Proximity to the attack is the kilometer proximity (binned into deciles) between the user's geo-tagged location and the attack location. Dashed lines signify 95% confidence intervals.

B.6 Results for individual attacks

Given the large volume of tweets, we are able to estimate our model separately for each of the eight sampled attacks. Table B.6 presents the corresponding results based on the 24-hour bandwidth. Generally speaking, we find consistent effects across all attacks: the tweets posted 24 hours after the attacks convey more negative feelings than those posted 24 hours before the attacks. The differences in the magnitude of the estimates can be attributed to the context surrounding the attacks. For instance, the effects appear to be stronger and statistically more robust for the 2017 Manchester Arena bombing and the 2017 London Bridge attack, owing to the fact that these attacks had a high number of victims, widespread media coverage and Islamist perpetrators.

Table B.6: The emotional effect of terrorism: inidvidual attacks

	Negative	Fear	Anger	Sadness	Disgust
	(1)	(2)	(3)	(4)	(5)
Panel A: 2016 Jo Cox Murder					
24-hour bandwidth	0.010***	0.009***	0.010***	0.004***	0.005***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	1,464,544	1,464,544	1,464,544	1,464,544	1,464,544
Panel B: 2017 Westminster Attack					
24-hour bandwidth	0.009**	0.018***	0.011***	0.004	0.007***
	(0.003)	(0.001)	(0.001)	(0.003)	(0.001)
Observations	912,502	912,502	912,502	912,502	912,502
Panel C: 2017 Manchester Arena Bombing					
24-hour bandwidth	0.024***	0.033***	0.021***	0.014***	0.018***
	(0.003)	(0.001)	(0.001)	(0.003)	(0.001)
Observations	907,344	907,344	907,344	907,344	907,344
Panel D: 2017 London Bridge Attack					
24-hour bandwidth	0.013***	0.025***	0.010***	0.008***	0.005***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	1,009,967	1,009,967	1,009,967	1,009,967	1,009,967
Panel E: 2017 Finsbury Park Attack					
24-hour bandwidth	0.002	0.003**	0.003*	-0.000	0.000
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	828,258	828,258	828,258	828,258	828,258
Panel F: 2017 Parsons Green Bombing					
24-hour bandwidth	0.003**	0.004***	0.000	0.002**	0.001
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	904,900	904,900	904,900	904,900	904,900
Panel G: 2019 London Bridge Stabbings					
24-hour bandwidth	0.007***	0.011***	0.006***	0.001	0.001
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	809,494	809,494	809,494	809,494	809,494
Panel H: 2020 Reading Stabbings					
24-hour bandwidth	0.001	0.001	0.001	0.003**	-0.000
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	806,093	806,093	806,093	806,093	806,093

Notes: The table shows the short-term effect of individual terrorist attacks on the respective sentiment or emotion. The short term effect is measured using the 24-hour bandwidth. Standard errors are clustered at the individual-level and reported in parentheses. * $p < .10$; ** $p < .05$; *** $p < .01$.

B.7 Heterogeneity analysis: time-to-event figures

In this section, we present the time-to-event figures for the attack groups we considered in our 'heterogeneity analysis' (see Section 2): (i) the six attacks with Islamist perpetrators versus the two attacks with far-right perpetrators (Figure B.7a); (ii) the four attacks with the highest number of victims versus the remaining four attacks (Figure B.7b); and (iii) the four attacks with the highest media coverage versus the remaining four attacks (Figure B.7c). For brevity and comparability, we focus on the three outcomes with the most pronounced post-attack effects: the overall negative sentiment, and the emotions of fear and anger. Overall, the patterns displayed in these figures support our key conclusions. First, attacks motivated by a radical interpretation of Islam result in more fearful sentiments than far-right attacks. Second, attacks with a high number of victims and extensive media attention elicit more negative sentiment and emotional responses than those with relatively fewer victims and less media coverage, and the corresponding effects last longer.

Figure B.7a: Time-to-event analysis: far-right versus Islamic attacks

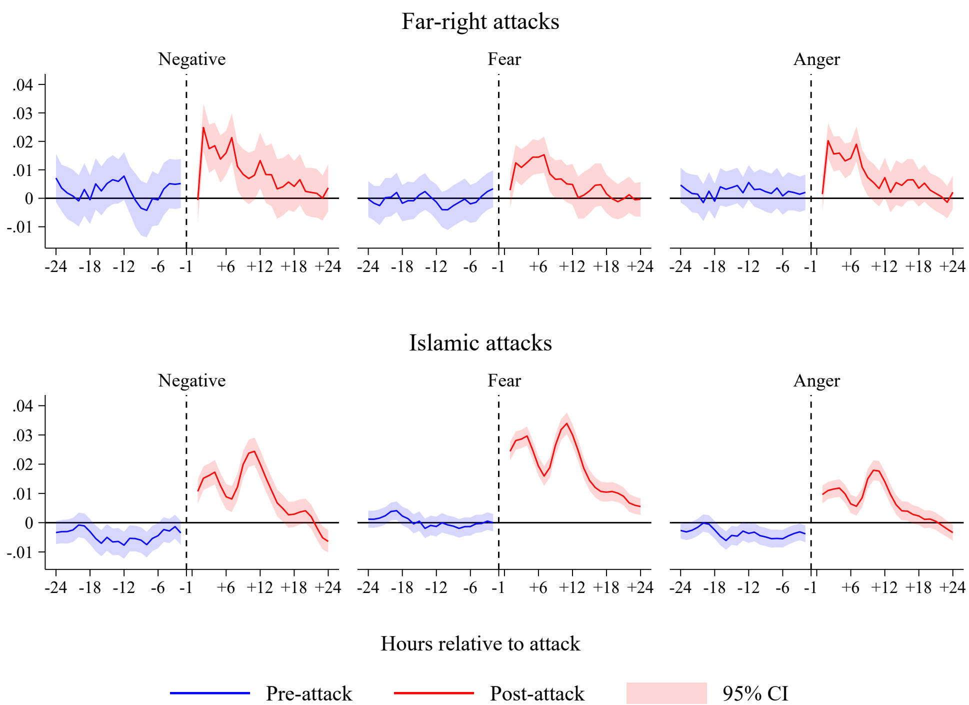

Notes: The figure shows the evolution of negative feelings 24 hours before and 24 hours after the sampled attacks. The tweets are aggregated at the hour level. The blue (red) solid line shows the 3-hour moving average estimates before (after) the attacks, taking the hour before the attack as the baseline. The tweets posted in the hour after the attack are dropped from the estimations. The shaded areas show the 95 percent confidence intervals.

Figure B.7b: Time-to-event analysis: high-victim versus low-victim attacks

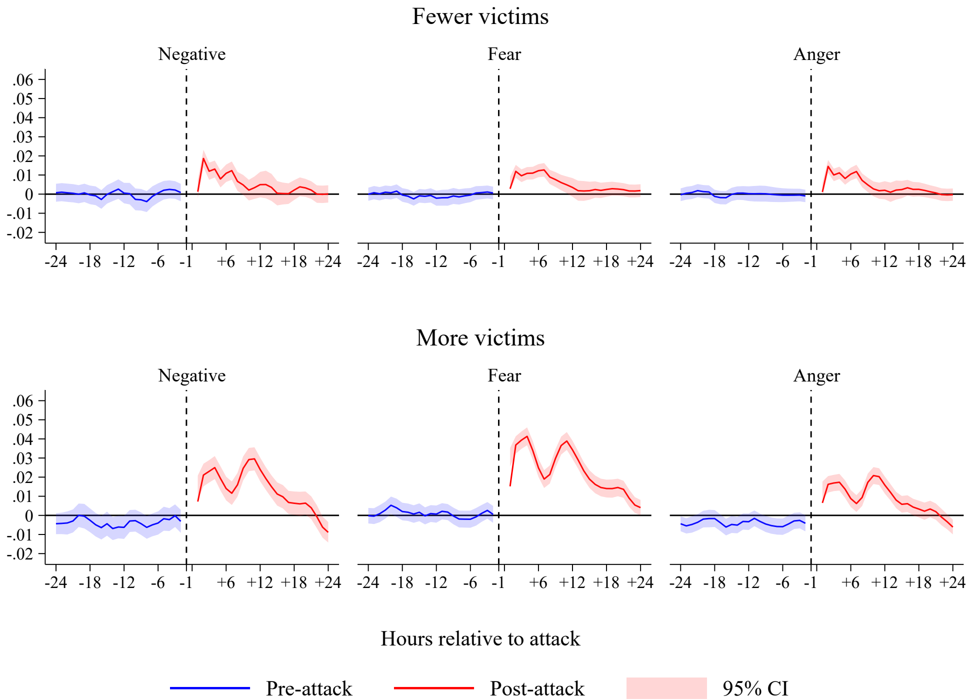

Notes: The figure shows the evolution of negative feelings 24 hours before and 24 hours after the sampled attacks. The tweets are aggregated at the hour level. The blue (red) solid line shows the 3-hour moving average estimates before (after) the attacks, taking the hour before the attack as the baseline. The tweets posted in the hour after the attack are dropped from the estimations. The shaded areas show the 95 percent confidence intervals.

Figure B.7c: Time-to-event analysis: high-coverage versus low-coverage attacks

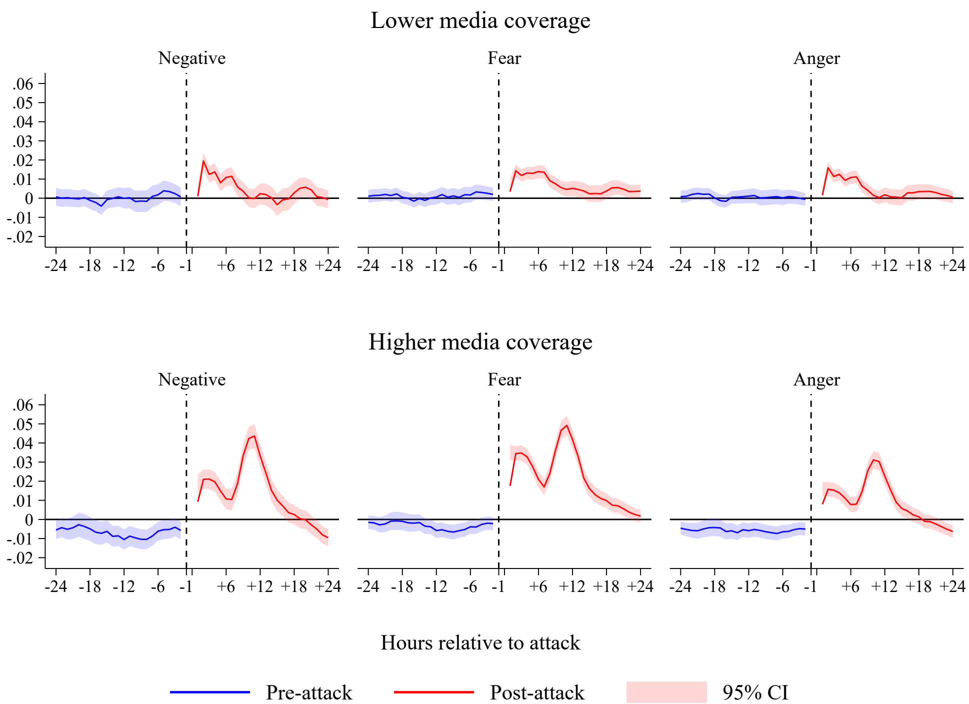

Notes: The figure shows the evolution of negative feelings 24 hours before and 24 hours after the sampled attacks. The tweets are aggregated at the hour level. The blue (red) solid line shows the 3-hour moving average estimates before (after) the attacks, taking the hour before the attack as the baseline. The tweets posted in the hour after the attack are dropped from the estimations. The shaded areas show the 95 percent confidence intervals.

C. References

Bove, V., Efthyvoulou, G., and Pickard, H. (2022a). Are the effects of terrorism short-lived? *WIDER Working Paper Series, No. 2022/66, Helsinki: UNU-WIDER*.

Bove, V., Efthyvoulou, G., and Pickard, H. (2022b). Did terrorism affect voting in the Brexit referendum? *British Journal of Political Science*, 52(3):1133–1150.

Hutto, C. and Gilbert, E. (2014). Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 8, pages 216–225.

Jones, H. (2020). More in common: the domestication of misogynist white supremacy and the assassination of Jo Cox. In *Racial Nationalisms*, pages 39–57. Routledge.

Loria, S. et al. (2018). textblob Documentation. *Release 0.15*, 2(8).

Mohammad, S. and Turney, P. (2010). Emotions evoked by common words and phrases: Using Mechanical Turk to create an emotion lexicon. In *Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text*, pages 26–34, Los Angeles, CA. Association for Computational Linguistics.

Mohammad, S. M. and Turney, P. D. (2013). Crowdsourcing a word-emotion association lexicon. *Computational Intelligence*, 29(3):436–465.

Muñoz, J., Falcó-Gimeno, A., and Hernández, E. (2020). Unexpected event during survey design: Promise and pitfalls for causal inference. *Political Analysis*, 28(2):186–206.

Nussio, E., Böhmelt, T., and Bove, V. (2021). Do terrorists get the attention they want? Comparing effects of terrorism across Europe. *Public Opinion Quarterly*, 85(3):900–912.

Pickard, H., Efthyvoulou, G., and Bove, V. (2023). What's left after right-wing extremism? the effects on political orientation. *European Journal of Political Research*, 62(1):338–351.

Zheng, X., Han, J., and Sun, A. (2018). A survey of location prediction on twitter. *IEEE Transactions on Knowledge and Data Engineering*, 30(9):1652–1671.