ONLINE APPENDIX

Are the effects of terrorism short-lived?

WIDER Working Paper 2022/66

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

June 2022

[*] Corresponding author: v.bove@warwick.ac.uk; [1] Department of Politics and International Studies, University of Warwick, Coventry, UK; [2] Department of Economics, University of Sheffield, Sheffield, UK; [3] Newcastle University Business School, Newcastle University, Newcastle upon Tyne, UK

Contents

A Further insights

A1 Background material on the three attacks

We focus on three major terrorist incidents that occurred over the period for which we have CMS data available: the 2005 London bombings, the 2007 Glasgow airport attack, and the 2013 Lee Rigby murder.

On 7 July 2005, Hassib Hussain, Mohammad Sidique Khan, Germaine Lindays, and Shezad Tanweer detonated four explosive devices in the London underground stations Aldgate, Edgware Road, and Russell Square, and a double-decker bus in Tavistock Square. A total of 52 people were killed and over 700 were injured—not including the four suicide bombers who were killed instantly upon detonating their explosive-filled rucksacks. Three of the four men left Leeds in a rented car in the early morning of that day and travelled to Luton where they met the fourth perpetrator. They then travelled by train to King's Cross Station where they split up and travelled to each of the aforementioned locations. The underground bombs were detonated at 08:50. The fourth bomber failed to do so because the Northern Line was closed and instead got on a bus and triggered the device at 09:47. This was the largest terrorist incident that had occurred in Great Britain since the Second World War.[1] Poignantly, this attack marked the day in which Al-Qaeda-linked terrorism came to the shores of Britain. It was the first attack of its kind in the UK after 9/11 in the USA and the 2004 Madrid train bombings.

The second attack occurred at the Glasgow airport on 30 June 2007. At 15:11, two men drove at the glass doors of the Glasgow airport terminal in a car filled with propane canisters. The vehicle was set ablaze, and upon leaving the vehicle, the driver poured petrol around and on himself, suffering severe burns. Five members of the public were injured in their attempts to help the police detain the perpetrators, but none sustained serious injuries. The attackers were identified as Bilal Abdullah, a British Muslim doctor of Iraqi ancestry, and Kafeel or Khalid Ahmed, an Indian engineer. Ahmed was the severely injured driver, who died as a result of his burns on 2 August. Immediately after the attack, the police evacuated the airport and all remaining flights for the day were suspended. The attack is historically significant for Scotland, as it was the first terrorist incident to have occurred in the devolved nation since the Lockerbie bombing in 1988.

The third attack happened on 22 May 2013 at 14:20. Off-duty Fusilier Lee Rigby of the Royal Regiment of Fusiliers was ran down with a car and subsequently stabbed and hacked to death with knives and a cleave in Woolwich, Southeast London. The perpetrators were Michael Adebolajo and Michael Adebowale. The men did not flee the scene and remained next to the victim's body until the police arrived nine minutes after a witness called the emergency services. The attackers were filmed telling passers-by that they had killed a soldier as revenge for the killing of Muslims by the British Army abroad. The assailants charged at the police when these arrived and, as a result, were shot. Both survived their injuries and were later found guilty of murder. Both attackers were British-born citizens of Nigerian descent who had converted to Islam. During the sentencing, Mr Justice Sweeney stated that their extremist views constituted a 'betrayal of Islam'. In response to this Adebowale shouted that '[t]hat [was] a lie' and Adebolajo shouted 'Allahu Akbar' (Allah is the greatest).

A2 Terrorism and emotions: evidence from tweets

To lend further empirical support to our main findings, we use Twitter data and analyse the sentiment and emotion content of terrorism-related tweets. We use Twitter's Academic Research product track that provides full access to Twitter's archive through their API. We then search for tweets that contain

[1] https://tinyurl.com/2p9hdpr7

the keyword 'terror' to build a corpus of tweets. Stop words (commonly used functional words such as 'the' and 'is') are removed and tweets are separated into individuals words. We then apply the NRCLex method that assigns each word an emotion (anger, fear, anticipation, trust, surprise, sadness, joy, and disgust) or a sentiment (positive or negative).

Out of the three attacks in our sample, we use tweets for the 2013 murder of Lee Rigby only, since Twitter was not available during the first attack in 2005 and had a very low user count during the second attack in 2007. We obtain a total of 251,965 tweets that occurred in the seven days after the attack and that contained the keyword 'terror'. Figure A1 presents the total counts of words that are registered as either positive or negative, together with the range of emotions. The figure shows that the overriding sentiment (in the very short run) is indeed negative and that the most common emotions are anger and fear.

A3 Media attention for the three attacks

In Section 2.3, we highlighted the importance of media coverage of a given attack in explaining the differences in temporal dynamics of risk perceptions and negative feelings. In this section, we provide some descriptive evidence that shows how media coverage evolved in the weeks after the three sampled attacks.

To begin with, we provide some images of newspapers published on the day after each attack (see Figure A2). We then analyse data on UK newspaper coverage over the first nine weeks after the attacks, using LexisNexis. To locate relevant articles, we include the keywords 'terrorist' or 'terrorism' and a set of attack-specific keywords including the location. We calculate the number of *new* newspaper articles written in each week after each attack and plot them in Figure A3. The figure shows three marked differences between the 2005 London bombings and the subsequent two attacks. First, over the first nine weeks considered, the total number of articles written is persistently higher for the 2005 London bombings.[2] Second, the initial spike in coverage for this attack is more than double that for the other two attacks, which may reflect the severity of the incident. Third, for the 2005 London bombings, the descent from the initial peak is much slower, and coverage persists at a non-zero number of new articles even nine weeks after the attack. In contrast, for the other two attacks, the coverage quickly dissipates to zero at one month after the attacks.

These findings reinforce, and potentially explain, the results from our main analysis: the effects for the 2005 London bombings are stronger and temporally more persistent, whereas the effects for the other two attacks start at a lower point and display a more linear decline.

[2] It is worth noting that online news coverage was not especially widespread in 2005 and thus most articles come from the newspaper daily publication. In later years, following the growth of online media, the coverage for the two last attacks is still below that for the 2005 London bombings.

Figure A1: Sentiment analysis of tweets

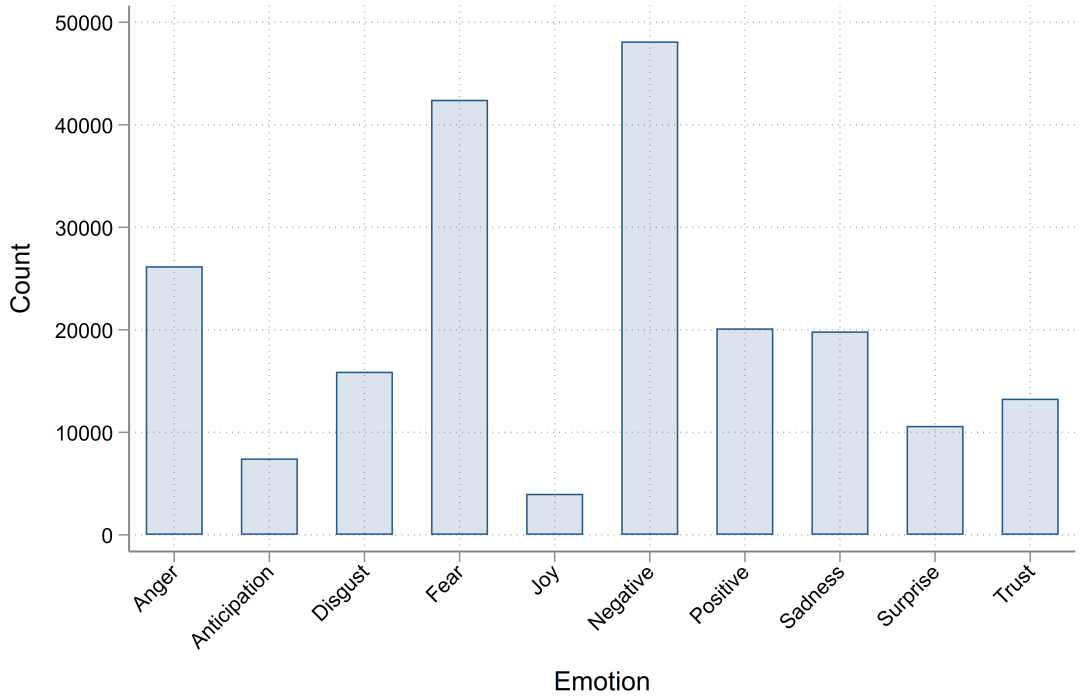

Note: the figure shows the count of emotions, obtained via the NRCLex package, contained in 251,965 tweets from the 7 days after the Lee Rigby murder on 22 May 2013.

Source: authors' illustration based on the NRCLex package.

Figure A2: Newspaper front pages

Note: selected front pages of newspapers published the day after each attack occurred. Row 1 relates to the 2005 London bombings; row 2 to the 2007 Glasgow airport attack; and row 3 to the 2013 Lee Rigby murder.

Source: authors' elaboration.

Figure A3: Timeline of newspaper coverage by attack

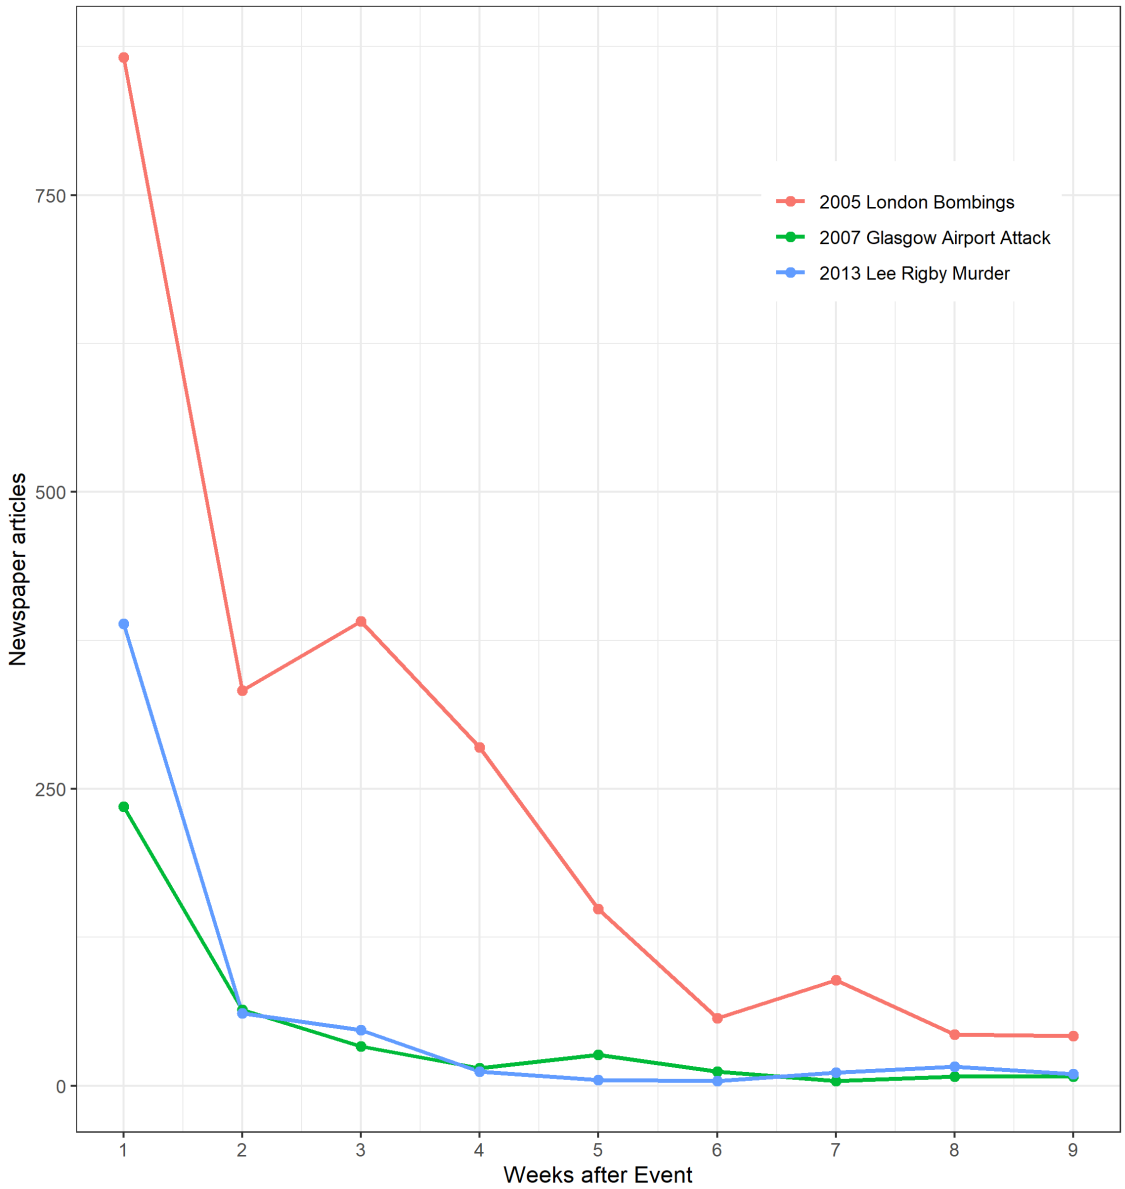

Source: authors' illustration based on LexisNexis.

B Additional empirical analyses

B1 Covariates and imbalances

A possible threat to our identification strategy is that individuals with specific characteristics may respond to the survey at different points in time, and these characteristics may be predictive of the outcome. To ensure that our results are not affected by such imbalances, we report estimates both before and after augmenting the baseline model with the following individual-level controls: gender (dummy: female vs male), age, age squared, ethnicity (dummy: White vs non-White), family status (dummy: has children vs does not have children), education (dummy variables capturing six education groups), and income (dummy variables capturing nine income groups). As shown in Figure B1, controlling for all these variables has no impact on our estimates, despite the fact that the sample sizes are now much smaller—see also Tables C1 and C2 for the full regression results.

As a further step, we perform balancing tests in the aforementioned characteristics across treatment and control units. Tables B1, B2, and B3 report the corresponding results for each time frame (very short run, short run, and medium run, respectively). We can see that, when we exploit information from the short and medium runs, there is a strong balance across treated and control units for nearly all attributes. On the other hand, when we exploit information from the very short run, we can observe some significant differences in a number of attributes (age, age squared, gender, and the last education group), which is not surprising given the smaller number of treated units in this case.

To correct for the imbalances reported above, we re-weight the sample through entropy balancing (Hainmueller 2012) such that the distribution of covariates among control units matches the moment conditions (until skewness) of the treated units. As shown in Figure B2, this exercise produces similar results as in Figure B1 and does not change our inferences. As an alternative approach, we rely on coarsened exact matching (CEM) (Blackwell et al. 2009) to pre-process the data and produce covariate balance between the treatment and control groups. In other words, instead of using the full sample of treated and control units, we now match treated units with a carefully selected group of matched control units before comparing their responses to the survey questions of interest. Figure B3 shows the results when we perform CEM on the full set of characteristics (mentioned above) and restrict the matched control units to come from the same attack-by-region group as the treated units. The evidence obtained is in line with our previous findings.

Figure B1: Main results—with and without control variables

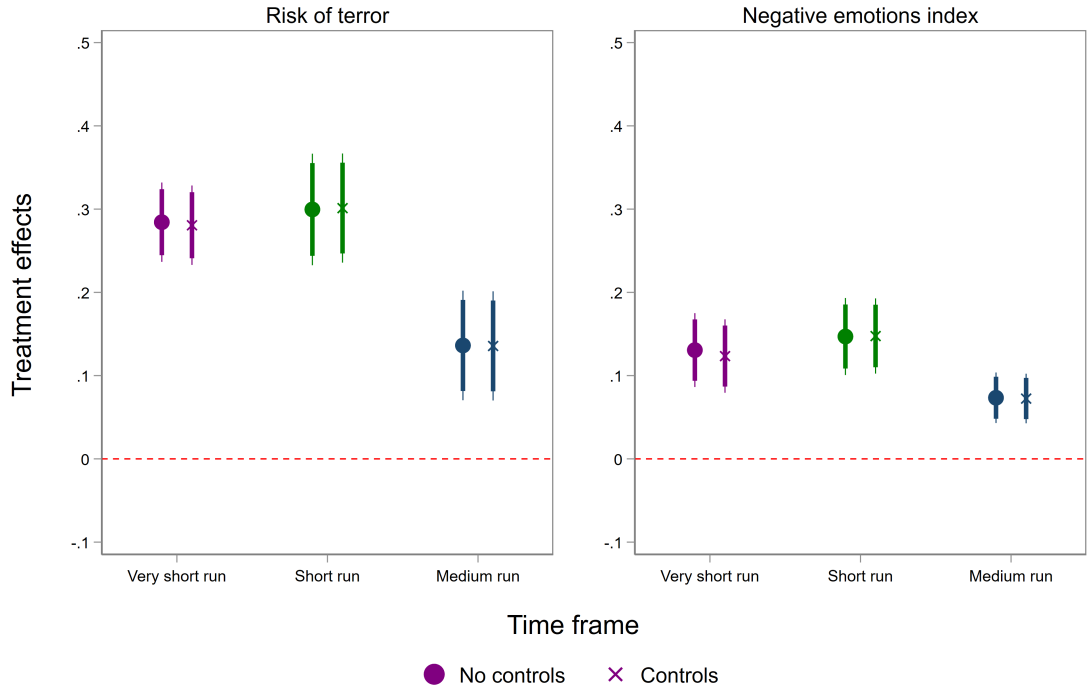

Note: the treatment effects are estimated using a linear probability model, controlling for attack-by-region fixed effects. Standard errors are clustered at the attack-by-region level. Fat (thin) lines signify the 90% (95%) confidence interval.

Source: authors' calculations based on CMS data.

Table B1: Covariate balanace—very short run

	Pre-attack		Post-attack			
	Mean	Std. dev.	Mean	Std. dev.	Diff.	p-value
Female	0.52	0.50	0.47	0.50	0.05	(0.01)
Age	49.07	14.91	54.36	13.93	-5.22	(0.00)
Age squared	2629.80	1477.75	3149.24	1491.80	-510.03	(0.00)
Has children	0.40	0.49	0.38	0.48	0.02	(0.18)
Education: 14 or under	0.01	0.11	0.01	0.10	0.00	(0.49)
Education: 15	0.13	0.33	0.11	0.31	0.02	(0.08)
Education: 16	0.22	0.42	0.22	0.42	0.00	(0.75)
Education: 17-18	0.21	0.41	0.20	0.40	0.01	(0.65)
Education: 19-20	0.08	0.27	0.07	0.26	0.01	(0.40)
Education: 21 or over	0.35	0.48	0.38	0.49	-0.04	(0.02)
White	0.96	0.19	0.95	0.23	0.01	(0.08)
Income: less than or £5,000	0.04	0.19	0.03	0.17	0.01	(0.21)
Income: £5,000 to £9,999	0.08	0.27	0.07	0.26	0.01	(0.46)
Income: £10,000 to £14,999	0.12	0.32	0.11	0.32	0.00	(0.81)
Income: £15,000 to £19,999	0.13	0.33	0.14	0.34	-0.01	(0.35)
Income: £20,000 to £24,999	0.12	0.33	0.11	0.32	0.01	(0.33)
Income: £25,000 to £29,999	0.11	0.32	0.12	0.33	-0.01	(0.43)
Income: £30,000 to £39,999	0.16	0.37	0.18	0.39	-0.03	(0.13)
Income: £40,000 to £49,999	0.11	0.32	0.12	0.32	-0.01	(0.66)
Income: £50,000 or more	0.13	0.34	0.11	0.32	0.02	(0.12)
Observations	3,253		953		4,351	

Note: this table shows the mean of covariates across treatment and control units, together with conventional t-tests for differences in means across the two groups.

Source: authors' calculations based on CMS data.

Table B2: Covariate balance—short run

	Pre-attack		Post-attack			
	Mean	Std. dev.	Mean	Std. dev.	Diff.	p-value
Female	0.52	0.50	0.51	0.50	0.01	(0.25)
Age	49.07	14.91	47.16	14.33	1.89	(0.00)
Age squared	2629.80	1477.75	2429.81	1373.44	197.77	(0.00)
Has children	0.40	0.49	0.41	0.49	-0.01	(0.26)
Education: 14 or under	0.01	0.11	0.02	0.13	-0.01	(0.07)
Education: 15	0.13	0.33	0.14	0.34	-0.01	(0.35)
Education: 16	0.22	0.42	0.24	0.42	-0.01	(0.39)
Education: 17-18	0.21	0.41	0.21	0.40	0.00	(0.72)
Education: 19-20	0.08	0.27	0.08	0.27	0.00	(0.91)
Education: 21 or over	0.35	0.48	0.33	0.47	0.02	(0.13)
White	0.96	0.19	0.97	0.17	-0.01	(0.13)
Income: less than or £5,000	0.04	0.19	0.03	0.16	0.01	(0.09)
Income: £5,000 to £9,999	0.08	0.27	0.08	0.27	-0.00	(0.93)
Income: £10,000 to £14,999	0.12	0.32	0.12	0.32	-0.00	(0.88)
Income: £10,000 to £19,999	0.13	0.33	0.13	0.33	-0.00	(0.65)
Income: £20,000 to £24,999	0.12	0.33	0.13	0.34	-0.01	(0.59)
Income: £20,000 to £29,999	0.11	0.32	0.11	0.32	-0.00	(0.90)
Income: £30,000 to £39,999	0.16	0.37	0.16	0.37	-0.00	(0.72)
Income: £40,000 to £49,999	0.11	0.32	0.11	0.31	0.01	(0.31)
Income: £50,000 or more	0.13	0.34	0.14	0.34	-0.00	(0.88)
Observations	3,253		2,633		6,089	

Note: this table shows the mean of covariates across treatment and control units, together with conventional t-tests for differences in means across the two groups.

Source: authors' calculations based on CMS data.

Table B3: Covariate balance—medium run

	Pre-attack		Post-attack			
	Mean	Std. dev.	Mean	Std. dev.	Diff.	p-value
Female	0.52	0.50	0.51	0.50	0.01	(0.17)
Age	49.07	14.91	49.30	14.58	-0.33	(0.27)
Age squared	2629.80	1477.75	2643.02	1440.46	-21.91	(0.45)
Has children	0.40	0.49	0.40	0.49	-0.00	(0.88)
Education: 14 or under	0.01	0.11	0.01	0.12	-0.00	(0.18)
Education: 15	0.13	0.33	0.13	0.33	-0.00	(0.98)
Education: 16	0.22	0.42	0.23	0.42	0.00	(0.99)
Education: 17-18	0.21	0.41	0.21	0.41	0.00	(0.71)
Education: 19-20	0.08	0.27	0.08	0.26	0.00	(0.53)
Education: 21 or over	0.35	0.48	0.35	0.48	-0.00	(0.73)
White	0.96	0.19	0.96	0.20	0.00	(0.81)
Income: less than or £5,000	0.04	0.19	0.03	0.16	0.01	(0.04)
Income: £5,000 to £9,999	0.08	0.27	0.08	0.27	0.00	(0.48)
Income: £10,000 to £14,999	0.12	0.32	0.12	0.32	0.00	(0.94)
Income: £15,000 to £19,999	0.13	0.33	0.12	0.32	0.01	(0.41)
Income: £20,000 to £24,999	0.12	0.33	0.13	0.33	-0.00	(0.63)
Income: £25,000 to £29,999	0.11	0.32	0.11	0.32	-0.00	(0.95)
Income: £30,000 to £39,999	0.16	0.37	0.17	0.37	-0.01	(0.13)
Income: £40,000 to £49,999	0.11	0.32	0.12	0.32	-0.00	(0.91)
Income: £50,000 or more	0.13	0.34	0.14	0.34	-0.00	(0.77)
Observations	3,253		11,704		15,432	

Note: this table shows the mean of covariates across treatment and control units, together with conventional t-tests for differences in means across the two groups.

Source: authors' calculations based on CMS data.

Figure B2: Entropy balancing

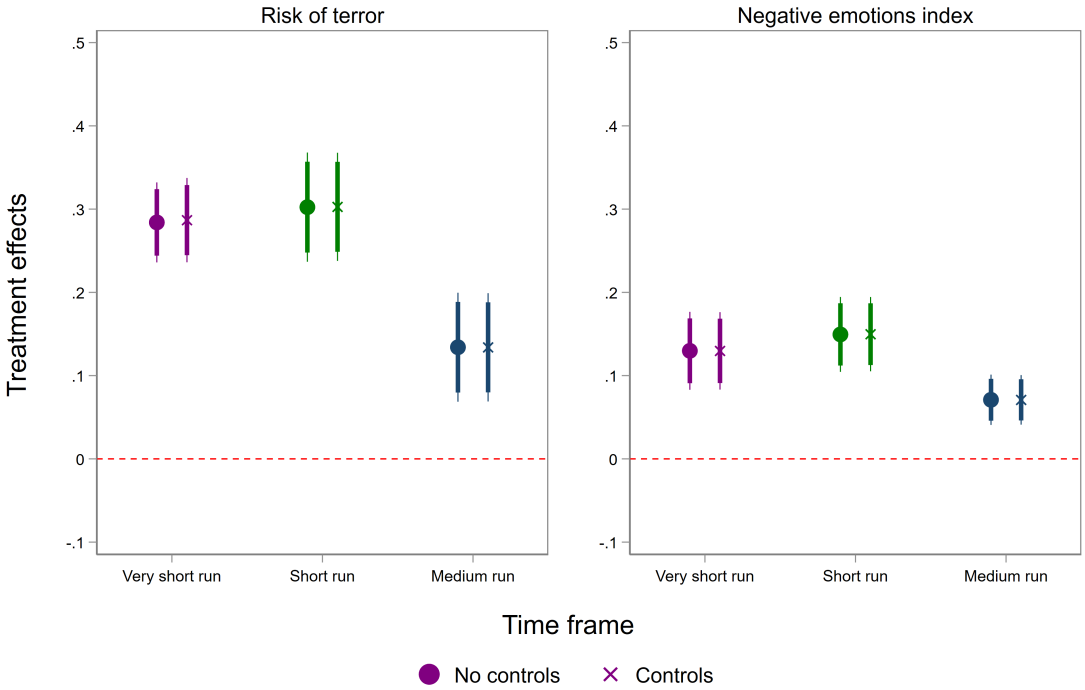

Note: the treatment effects are estimated using a linear probability model, controlling for attack-by-region fixed effects. The estimates are balanced using entropy weights that match the mean, variance, and skewness of covariates across the treatment and control units. Standard errors are clustered at the attack-by-region level. Fat (thin) lines signify the 90% (95%) confidence interval.

Source: authors' calculations based on CMS data.

Figure B3: Coarsened-exact matching

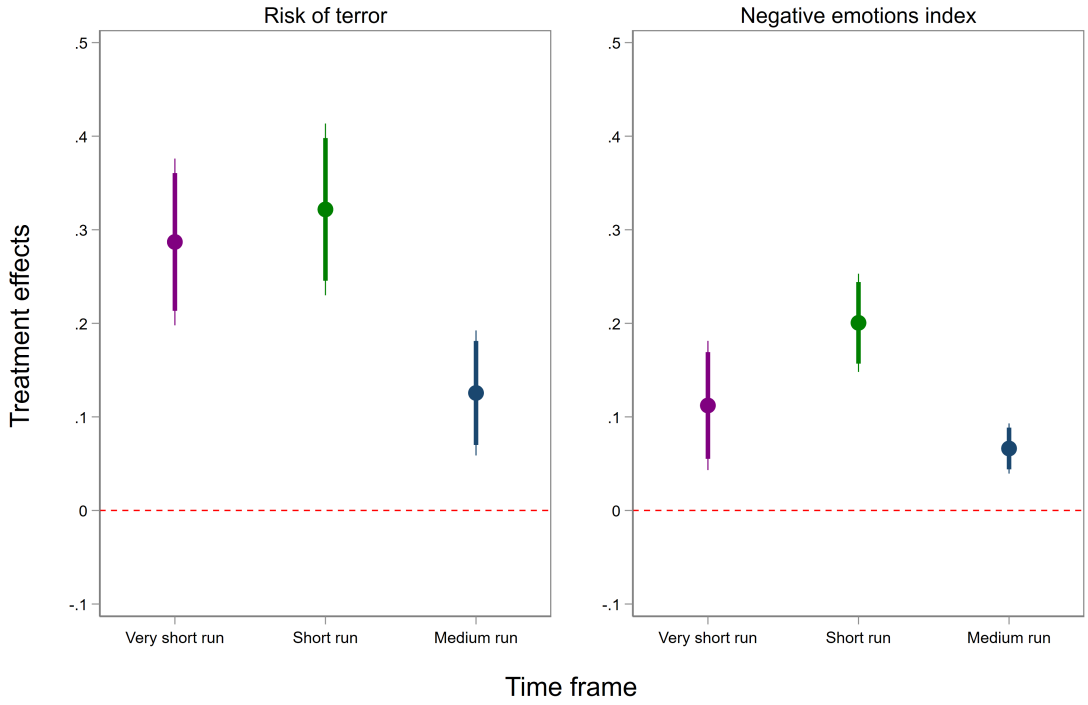

Note: this figure shows the treatment effects after performing coarsened-exact matching. To locate matches, we use the full set of control variables and restrict the matched control units to come from the same attack-by-region group as the treated units. Standard errors are clustered at the attack-by-region level. Fat (thin) lines signify the 90% (95%) confidence interval.

Source: authors' calculations based on CMS data.

B2 Identification validity tests

To strengthen our causal inference, we need to address two additional issues. The first relates to pre-existing trends; the second relates to the failed terrorist attack in July 2005.

It is possible that our estimates capture pre-existing time trends in the outcome variables, which are unrelated to the timing of the attacks. To test for the presence of such trends, we consider placebo treatments at an arbitrary time point at the left of the cut-off points, as recommended by Muñoz et al. (2020). More precisely, we focus on the pre-attack samples and set the placebo attack dates to be in the middle of the pre-attack periods. In this way, the 'placebo control' group includes the individuals interviewed 16 to 30 days before the actual attacks, and the 'placebo treatment' group includes the individuals interviewed one to 15 days before the actual attacks. We then run the same regression set-up as before. The results are presented in Figure B4. As expected, these placebo treatments have no significant effect on people's risk assessments and negative feelings.

On 21 July 2005, two weeks after the 2005 London bombings, there was a failed plot in which terrorists re-targeted the London underground network. The bombs failed to explode and there were no fatalities. This 'collateral' event could jointly affect our outcome variables, and thus bias our estimates.[3] To test for this, we focus on the original treatment group of the successful attack, and we compare individuals interviewed in the week after the failed attack with those interviewed in the week before this attack. The results are reported in Figure B5. We can see that, for both outcome variables, the 'post-failed-attack' estimate is very close to zero and fails to reach statistical significance, which indicates that this collateral event is not driving our main effects. This is likely because the original shock was so large and persistent that there was no room for a further increase in risk perceptions and negative feelings.

[3] As pointed out by Muñoz et al. (2020), this can be seen as a problem of an imprecise treatment, as it makes it difficult to narrowly interpret the effect as a consequence of the treatment (event) itself.

Figure B4: Testing for pre-existing time trends

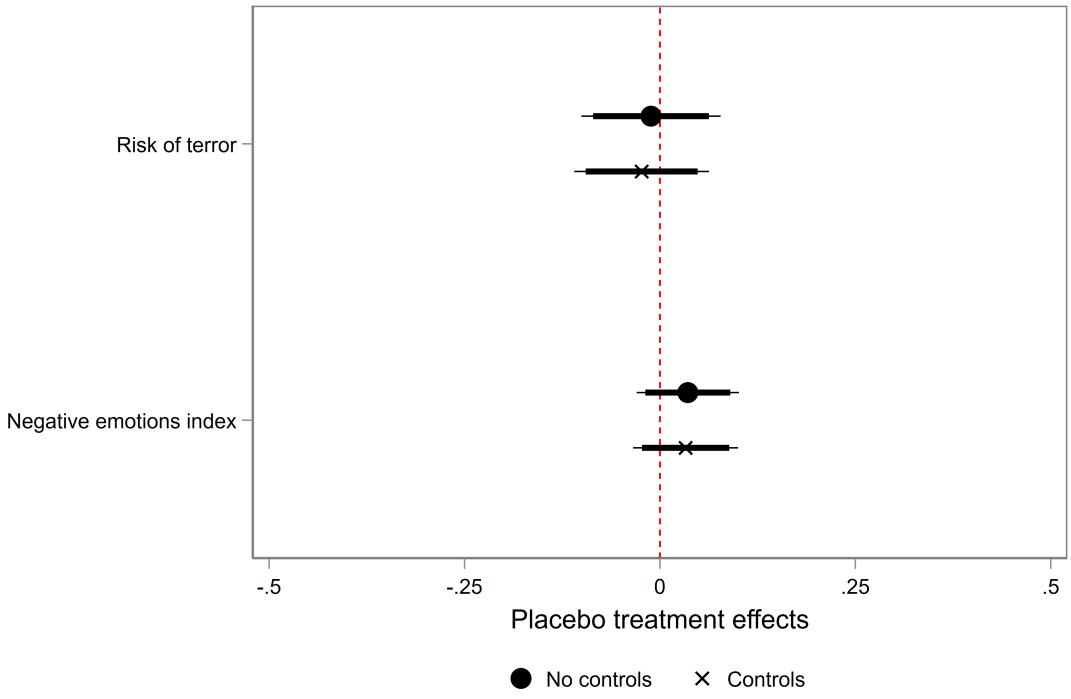

Note: the treatment effects are estimated using a linear probability model, controlling for attack-by-region fixed effects. Standard errors are clustered at the attack-by-region level. Fat (thin) lines signify the 90% (95%) confidence interval.

Source: authors' calculations based on CMS data.

Figure B5: Collateral event—failed 21 July 2005 attack

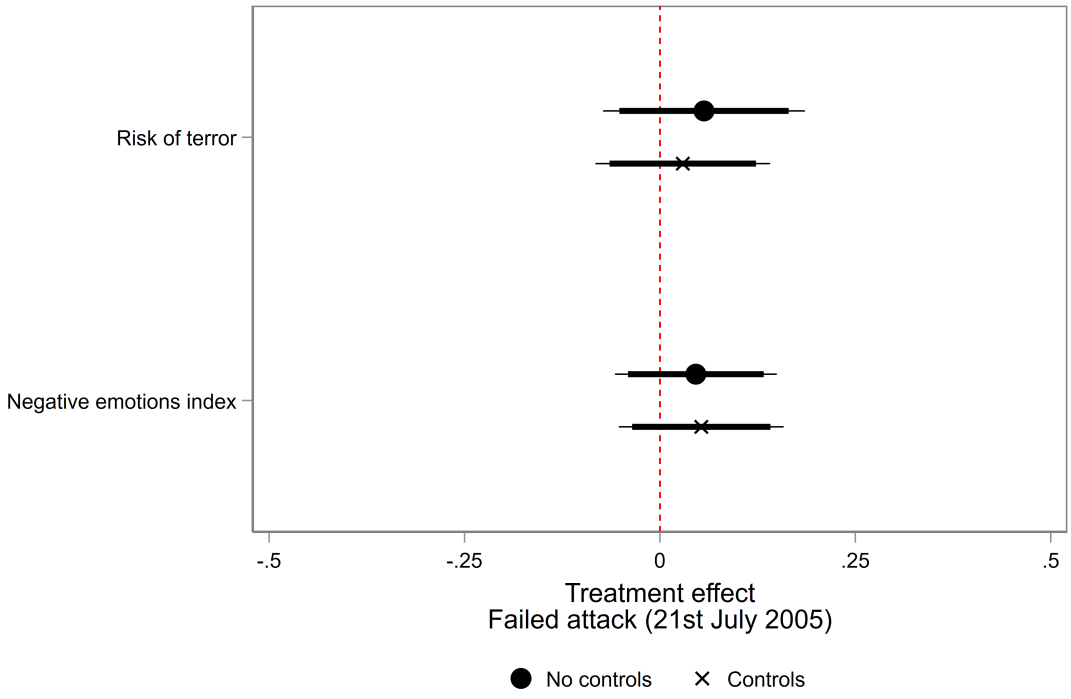

Note: the treatment effects are estimated using a linear probability model, controlling for attack-by-region fixed effects. Standard errors are clustered at the attack-by-region level. Fat (thin) lines signify the 90% (95%) confidence interval.

Source: authors' calculations based on CMS data.

B3 Placebo tests: alternative outcomes

In this section, we perform placebo tests where we examine the treatment effect on outcomes that should not be affected by terrorist incidents (or, at least, in the same way).

First, we employ measures capturing public assessments about two other key issues: crime and public healthcare. To construct these measures, we consider individuals' responses to the statements '*Do you think that the crime situation in Britain these days is...*' and '*Do you think the National Health Service in Britain these days is...*', and as in the case of terrorism risk assessments, we assign value 1 to the responses 'a little worse' and 'a lot worse' (and 0 to all the other responses). Figure B6 shows the results for these two outcomes, based on the same regression set-up as before. The treatment estimates are very close to zero and, in most of the cases, they fail to reach statistical significance. The only exception is when we exploit information from the short run, where we can observe a very small displacement effect, suggesting that exposure to terrorism sways public opinion away from other popular issues. At the same time, the absence of positive and statistically significant effects for the crime-related outcome confirms that the terrorist incidents are correctly perceived by the large audience as acts as terrorism rather than violent crime.[4]

Second, we employ measures capturing negative emotions about the state of the economy, which is often ranked as a top national concern by the British public. As in the case of terrorism, we consider individuals' responses to the question '*Which, if any, of the following words describe your feelings about the country's general economic situation?*' and construct dummy variables for the four negative emotions (anger, disgust, unease, and fear), together with a composite index. Once again, we can see that the resulting estimates are very small in magnitude, statistically insignificant, or in the opposite direction; i.e. people reporting less negative feelings about the economy in the immediate aftermath of a terrorist attack (see Figure B7).

[4] See Brück and Müller (2010) on what drives concern about terrorism vis-a-vis crime.

Figure B6: Public assessments about crime and public healthcare

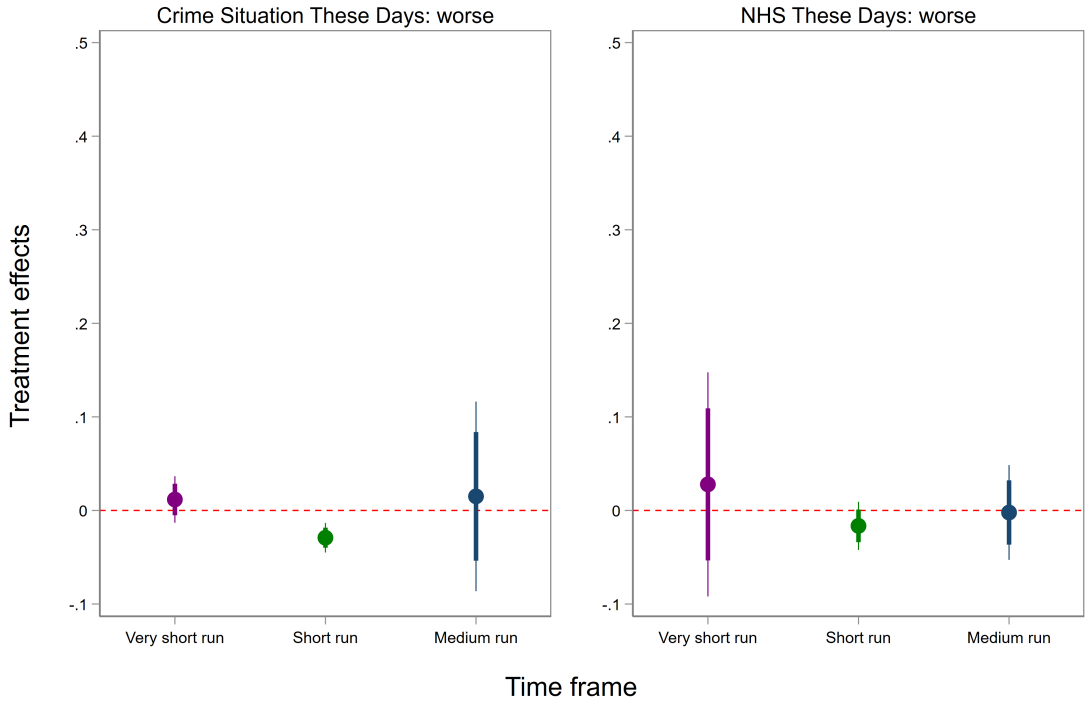

Note: the treatment effects are estimated using a linear probability model, controlling for attack-by-region fixed effects. Standard errors are clustered at the attack-by-region level. Fat (thin) lines signify the 90% (95%) confidence interval.

Source: authors' calculations based on CMS data.

Figure B7: Negative emotions about the state of the economy

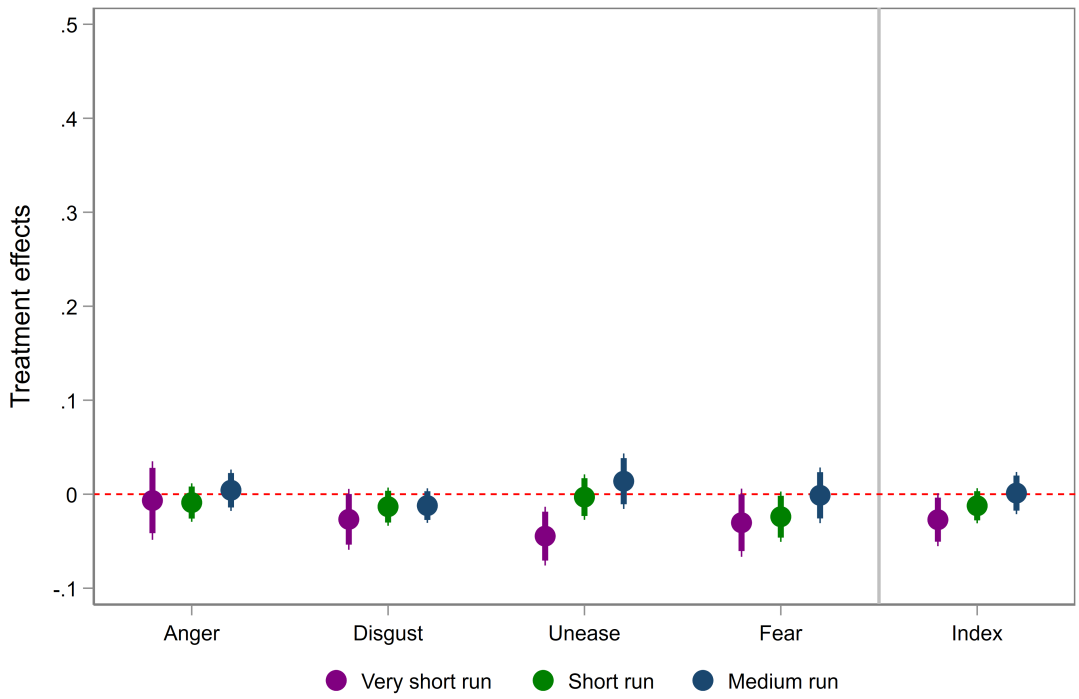

Note: the treatment effects are estimated using a linear probability model, controlling for attack-by-region fixed effects. Standard errors are clustered at the attack-by-region level. Fat (thin) lines signify the 90% (95%) confidence interval.

Source: authors' calculations based on CMS data.

B4 Positive emotions about the risk of terrorism

In this section, we examine how positive emotions about the risk of terrorism evolve over time in response to a terrorist attack. To do so, we run the same regressions as in Figure 3, but we now focus on the four positive emotions: happiness, hope, confidence, and pride. The results are displayed in Figure B8. Generally speaking, we observe the opposite patterns to those of negative emotions: after a terrorist attack, individuals are less likely to report positive feelings about the risk of terrorism—though the corresponding effects appear to be very small in magnitude and are mostly driven by a reduction in 'hope' and 'confidence' in the very short run and short run.

Figure B8: Positive emotions about the risk of terrorism

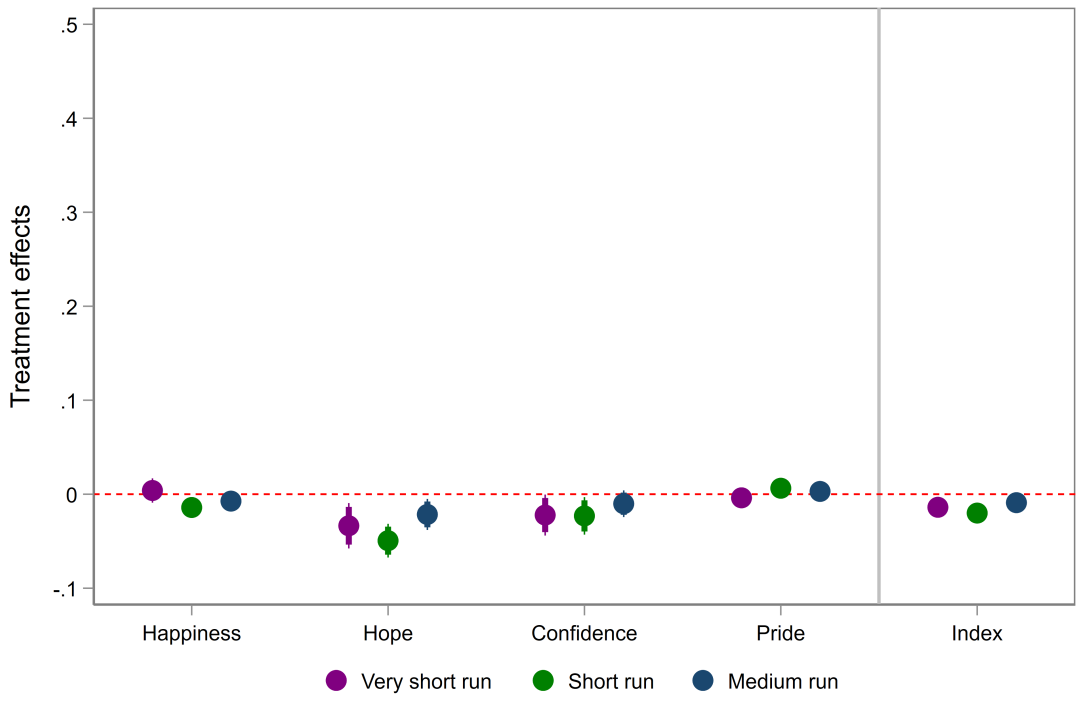

Note: the treatment effects are estimated using a linear probability model, controlling for attack-by-region fixed effects. Standard errors are clustered at the attack-by-region level. Fat (thin) lines signify the 90% (95%) confidence interval.

Source: authors' calculations based on CMS data.

B5 Heterogeneous treatment effects across individuals

We have seen, so far, that terrorism causes an increase in risk perceptions and negative emotions, which persists in the medium run. We now ask if this evidence is consistent across all population groups regardless of observed characteristics; that is, whether individuals with a certain covariate profile can exhibit the opposite patterns (e.g. report lower risk perceptions after the attacks) or be associated with shorter-lived effects. To do so, we employ a causal forest approach. Causal forest is a machine learning algorithm that automates the search for heterogeneity in the treatment effect (see Athey et al. 2019). In other words, it estimates the treatment effect for each individual in our sample as a function of their covariate profile, known as the conditional average treatment effect (CATE).

Figure B9 plots the CATEs (ordered by effect size) across the three time frames along with the 95 per cent confidence intervals. The first row reports the results for risk perceptions, whereas the second row reports the results for the negative emotions index. According to these plots, over 95 per cent of individuals in our sample have a positive treatment effect. In addition, in nearly all cases,[5] we cannot reject the null hypothesis that the CATE is significantly different from the local average treatment effect (LATE), which indicates that there are no heterogeneous effects with respect to individuals' characteristics. All in all, the analysis in this section reveals a robust degree of homogeneity in the direction (and duration) of the terrorism effects across individuals.

[5] The only exception is when we exploit information from the medium run for the negative emotions index.

Figure B9: Causal forest

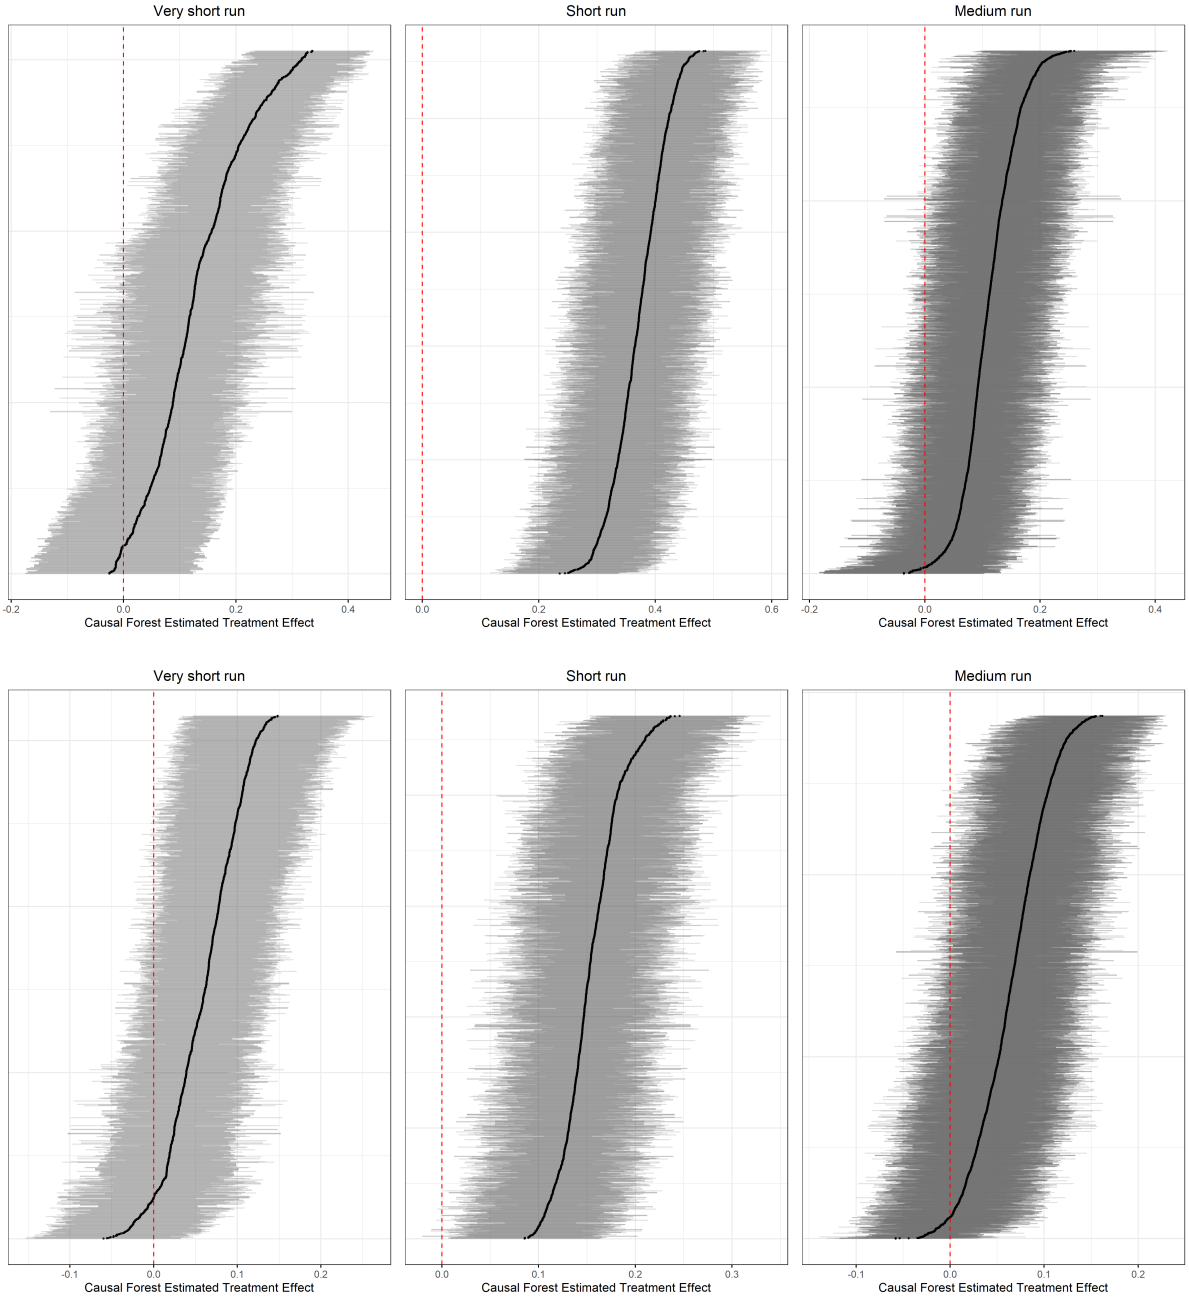

Note: the dependent variable in Row 1 is *Risk of terror*. The dependent variable in Row 2 is *Negative emotions index*. Estimated effects are obtained using the grf package for R with the recommended settings of honest splitting (i.e. sub-sample splitting) and 4000 trees. Black lines indicate estimated treatment effect for each individual, as a function of their covariate profile, ordered by effect size. Grey horizontal lines indicate 95% confidence intervals. Covariates include the full list of control variables (as reported in Section B1) and attack-by-region fixed effects.

Source: authors' calculations based on CMS data.

B6 Attacked vs non-attacked regions

Physical proximity to a terrorist attack 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 (Bove, Böhmelt, and Nussio 2021; Bove, Efthyvoulou, and Pickard 2021; Nussio et al. 2021). In line with these arguments, one would expect that distance from terrorism will act as a moderating factor whereby individuals that reside further away from an attack are less likely to report increased risk perceptions and negative feelings after attacks. Yet, the existence of a 'proximity effect' has become a debated issue, and Agerberg and Sohlberg (2021) find that individuals close to the attack do not display stronger reactions compared to less proximate individuals.

In the CMS data, the location of the respondents is only available at the region level. As such, to test whether physical proximity can influence the terrorism-induced reactions, we run the same analysis separately for individuals living in non-attacked regions and those living in attacked regions, with the latter capturing the regions in which the attacks took place. Figures B10 and B11 display the corresponding results for the two outcomes of interest. We can see that the effects on negative feelings are stronger in the attacked regions than in the non-attacked regions (especially in the very short run and short run), whereas the effects on risk perceptions are quite similar between the two samples. Overall, the analysis in this section suggests that, while physical distance 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—see also Pickard et al. (2022) for a similar finding.

Figure B10: Risk of terror—attacked vs non-attacked regions

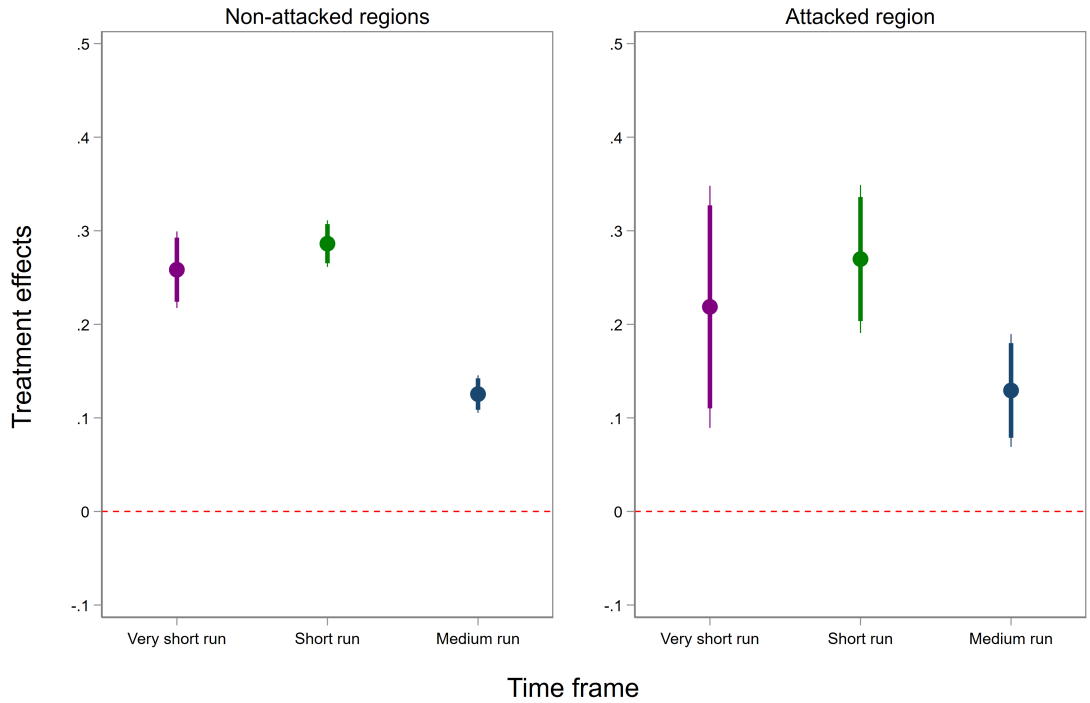

Note: the treatment effects are estimated using a linear probability model, controlling for attack-by-region fixed effects (for the non-attacked regions) and attack fixed effects (for the attacked regions). Standard errors are clustered at the attack-by-region level. Fat (thin) lines signify the 90% (95%) confidence interval.

Source: authors' calculations based on CMS data.

Figure B11: Negative emotions index—attacked vs non-attacked regions

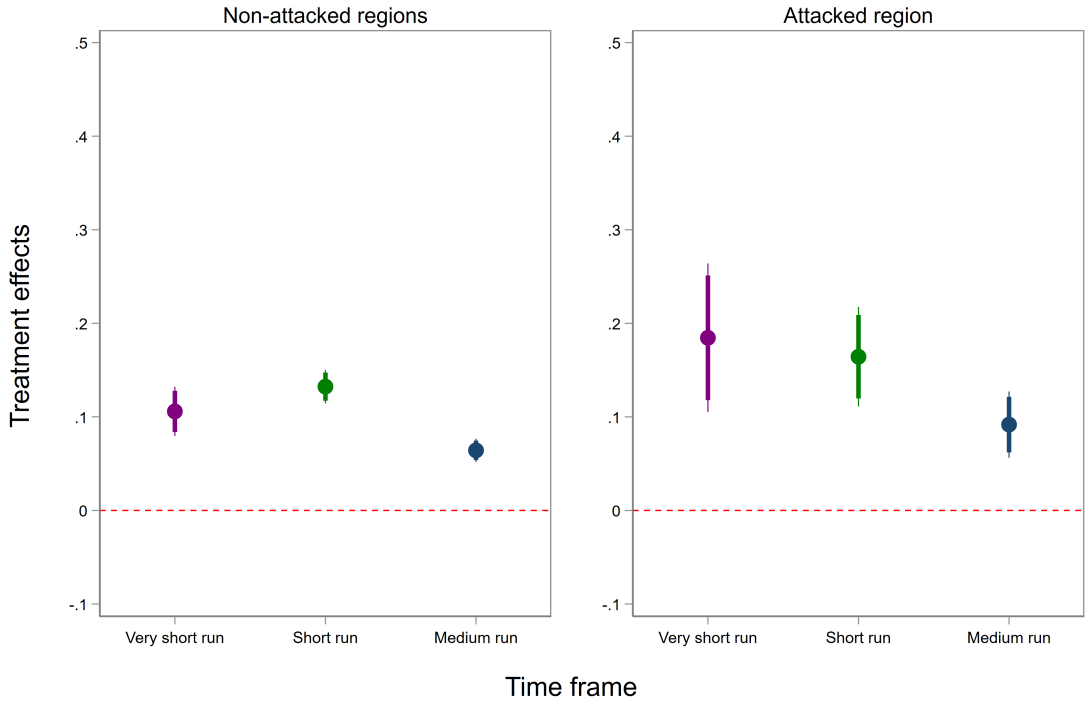

Note: the treatment effects are estimated using a linear probability model, controlling for attack-by-region fixed effects (for the non-attacked regions) and attack fixed effects (for the attacked regions). Standard errors are clustered at the attack-by-region level. Fat (thin) lines signify the 90% (95%) confidence interval.

Source: authors' calculations based on CMS data.

B7 Alternative estimation method: probit model

Throughout our main analysis, we estimate treatment effects on binary outcome variables using a linear probability model. As recently shown by Timoneda (2021), the linear probability model produces very accurate estimates both with highly common data and rare events data. Nevertheless, to address any remaining concerns about the accuracy of our chosen estimation technique, we check robustness to estimating our baseline specifications (Figures 2 and 3) using a probit model. As shown in Figure B12, the choice of the estimation model does not affect our inferences.

Figure B12: Probit estimation

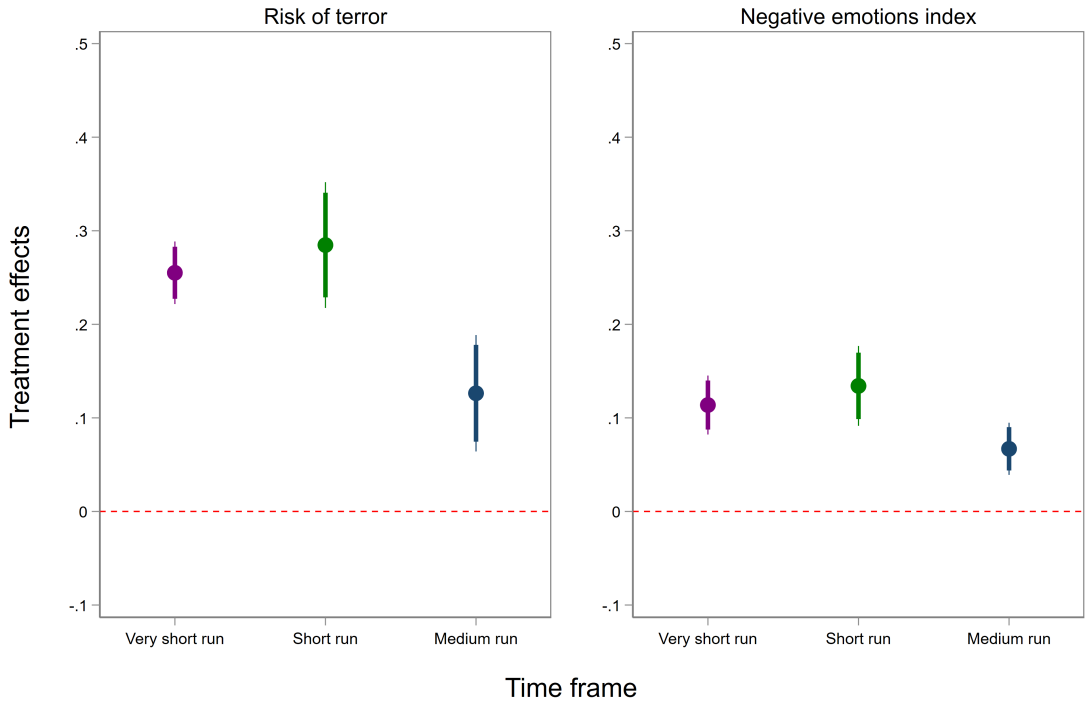

Note: the treatment effects are estimated using a probit model, controlling for attack-by-region fixed effects. Standard errors are clustered at the attack-by-region level. Fat (thin) lines signify the 90% (95%) confidence interval.

Source: authors' calculations based on CMS data.

B8 Second-order effects: migration attitudes

In Figure B13, we take our analysis one step further and test for the existence of a 'second-order echo effect' of terrorism, its influence on migration attitudes. To do so, we explore individuals' answer to the question '*How important a problem is the number of asylum seekers coming to Britain these days?*' (using a 0–10 scale) and estimate the treatment effect with OLS according to the specification of Equation (1). We find that, in the first week after the attacks, individuals are, on average, 0.25 points higher up the scale; that is, they perceive the number of asylum-seekers as a more important problem compared to before the attacks. This is in line with previous studies documenting that, in the wake of terrorist attacks, members of the broader audience are more likely to perceive foreigners and out-groups in general as a threat to the homogeneity of the nation-state population (Abou-Chadi 2016; Böhmelt et al. 2020; Bove, Böhmelt, and Nussio 2021; Helbling and Kalkum 2018; Helbling and Meierrieks 2022). However, our results also reveal that terrorism can cause a more permanent shift in such perceptions: the initial surge is followed by a slight decrease in the short run and then a stabilization at the same levels in the medium run. That said, it must be acknowledged that the second-order terrorism effects (e.g. on attitudes not directly elicited by terrorism) are likely to be subject to bias arising from the occurrence of other unrelated events, especially when we exploit information from longer time intervals, such as the short and medium runs (see also discussion in Section 3).

Figure B13: Asylum seekers as a problem

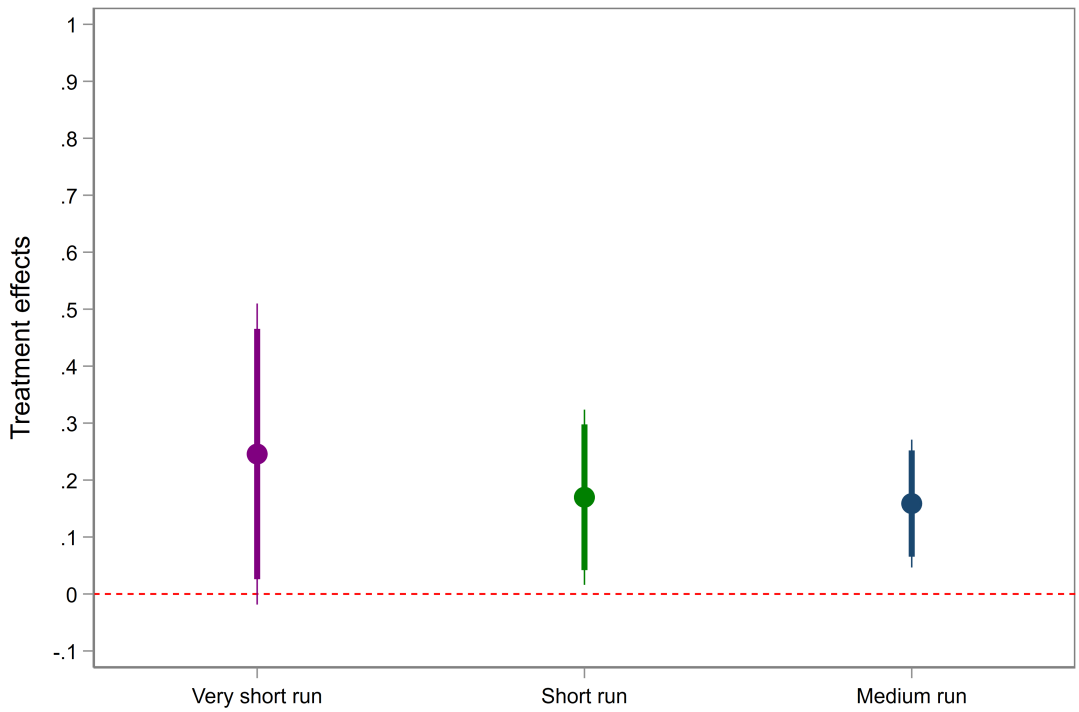

Note: the treatment effects are estimated using OLS, controlling for attack-by-region fixed effects. Standard errors are clustered at the attack-by-region level. Fat (thin) lines signify the 90% (95%) confidence interval. N (very short run) = 4,207; N (short run) = 5,871; and N (medium run) = 14,946.

Source: authors' calculations based on CMS data.

C Full regression results

See Figure C1 and Tables C1 and C2 below.

Figure C1: Negative emotions index—non-parametric estimates

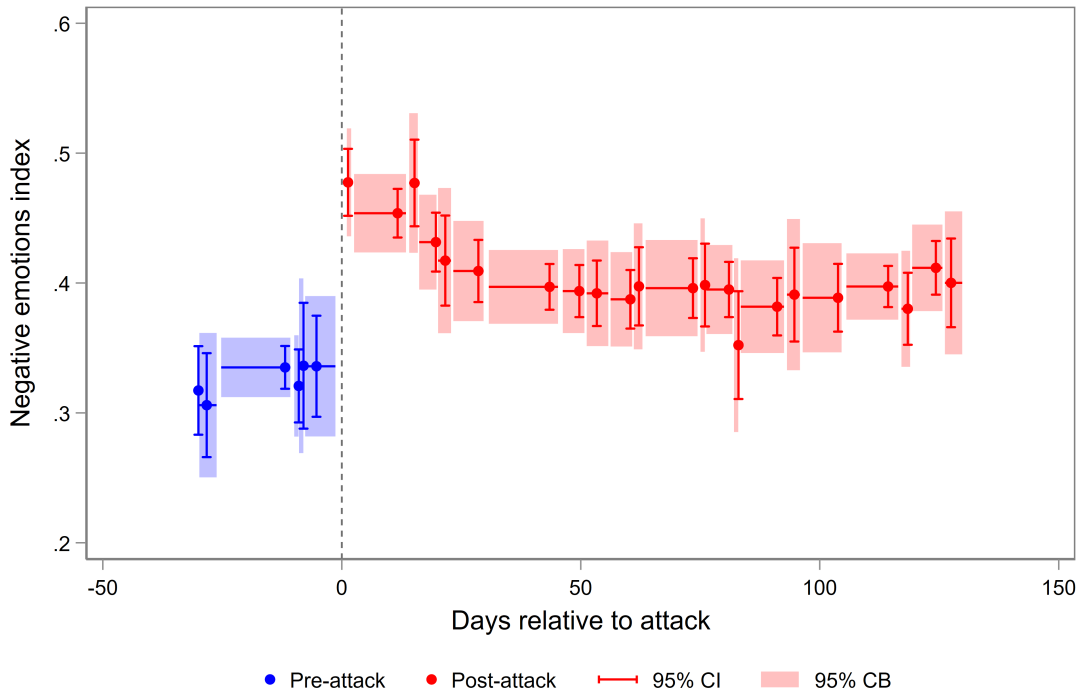

Note: this figure displays a binned scatterplot and the corresponding confidence intervals and confidence bands, as described in Cattaneo et al. (2019b) and implemented using the `binsreg` package. We choose the number of bins by minimizing the integrated mean squared error of the binned scatterplot in the pre- and post-attack periods, as in Cattaneo et al. (2019a). The estimation includes attack-by-region fixed effects.

Source: authors' calculations based on CMS data.

Table C1: Main results—risk of terror

	Risk of terror								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Very short run	0.255***	0.281***	0.284***						
	(0.016)	(0.023)	(0.023)						
Short run				0.285***	0.301***	0.300***			
				(0.033)	(0.032)	(0.033)			
Medium run							0.126***	0.136***	0.136***
							(0.031)	(0.032)	(0.032)
Female		0.048***			0.048***			0.043***	
		(0.015)			(0.012)			(0.010)	
Age		-0.001			-0.001			0.001	
		(0.003)			(0.003)			(0.002)	
Age squared		0.000			0.000			-0.000	
		(0.000)			(0.000)			(0.000)	
Has children		0.002			-0.006			-0.002	
		(0.019)			(0.015)			(0.011)	
Education: 15		0.116			0.068			0.019	
		(0.079)			(0.058)			(0.034)	
Education: 16		0.129			0.079			0.002	
		(0.081)			(0.060)			(0.035)	
Education: 17–18		0.131			0.071			-0.001	
		(0.079)			(0.056)			(0.034)	
Education: 19–20		0.117			0.050			-0.017	
		(0.077)			(0.059)			(0.039)	
Education: 21 or over		0.071			0.040			-0.050	
		(0.081)			(0.057)			(0.034)	
White		0.058			0.039			0.044	
		(0.051)			(0.045)			(0.029)	
Income: £5,000 to £9,999		-0.021			0.010			0.006	
		(0.057)			(0.039)			(0.027)	
Income: £10,000 to £14,999		-0.082			-0.060			0.007	
		(0.049)			(0.040)			(0.025)	
Income: £15,000 to £19,999		-0.016			-0.002			0.033	
		(0.046)			(0.032)			(0.025)	
Income: £20,000 to £24,999		-0.024			-0.017			0.038	
		(0.043)			(0.032)			(0.027)	
Income: £25,000 to £29,999		-0.045			-0.031			0.028	
		(0.048)			(0.038)			(0.029)	
Income: £30,000 to £39,999		-0.091*			-0.052			0.022	
		(0.047)			(0.032)			(0.028)	
Income: £40,000 to £49,999		-0.059			-0.030			0.011	
		(0.048)			(0.032)			(0.032)	
Income: £50,000 or more		-0.087*			-0.029			0.011	
		(0.044)			(0.033)			(0.025)	
Attack × Region FEs	✓	✓	✓	✓	✓	✓	✓	✓	✓
R-squared	0.074	0.092	0.080	0.155	0.169	0.162	0.086	0.088	0.081
Observations	4,186	3,052	3,052	5,886	4,594	4,594	14,957	11,220	11,220

Note: this table reports the full regression results for the variable *Risk of terror*. For each time frame, we present the results of three specifications: (i) without controls; (ii) with controls; (iii) without controls but based on the same sample as in the full control specification. Standard errors are clustered at the attack-by-region level and are reported in parentheses; * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Source: authors' calculations based on CMS data.

Table C2: Main results—negative emotions index

	Negative emotions index								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Very short run	0.114***	0.124***	0.131***						
	(0.015)	(0.022)	(0.022)						
Short run				0.134***	0.148***	0.147***			
				(0.021)	(0.022)	(0.023)			
Medium run							0.067***	0.073***	0.073***
							(0.014)	(0.015)	(0.015)
Female		0.034**			0.054***			0.035***	
		(0.015)			(0.009)			(0.006)	
Age		0.005**			0.002			0.003*	
		(0.002)			(0.002)			(0.002)	
Age squared		-0.000			-0.000			-0.000	
		(0.000)			(0.000)			(0.000)	
Has children		0.001			-0.001			0.013**	
		(0.010)			(0.009)			(0.006)	
Education: 15		0.091*			0.022			0.001	
		(0.046)			(0.040)			(0.030)	
Education: 16		0.087			0.007			-0.024	
		(0.053)			(0.039)			(0.033)	
Education: 17–18		0.058			-0.029			-0.044	
		(0.050)			(0.040)			(0.031)	
Education: 19–20		-0.003			-0.064			-0.062*	
		(0.053)			(0.045)			(0.031)	
Education: 21 or over		-0.007			-0.079**			-0.098***	
		(0.048)			(0.039)			(0.031)	
White		-0.025			0.006			-0.017	
		(0.022)			(0.027)			(0.016)	
Income: £5,000 to £9,999		-0.011			0.027			0.012	
		(0.033)			(0.025)			(0.022)	
Income: £10,000 to £14,999		-0.004			0.021			0.003	
		(0.037)			(0.028)			(0.022)	
Income: £15,000 to £19,999		-0.014			0.018			0.001	
		(0.033)			(0.026)			(0.024)	
Income: £20,000 to £24,999		-0.016			0.036			0.005	
		(0.032)			(0.026)			(0.023)	
Income: £25,000 to £29,999		-0.011			0.032			0.006	
		(0.029)			(0.025)			(0.023)	
Income: £30,000 to £39,999		-0.012			0.035			-0.003	
		(0.033)			(0.023)			(0.023)	
Income: £40,000 to £49,999		-0.034			0.004			-0.010	
		(0.034)			(0.026)			(0.021)	
Income: £50,000 or more		-0.052			0.022			-0.028	
		(0.034)			(0.030)			(0.025)	
Attack × Region FEs	✓	✓	✓	✓	✓	✓	✓	✓	✓
R-squared	0.042	0.087	0.051	0.076	0.107	0.078	0.049	0.072	0.043
Observations	4,350	3,148	3,148	6,089	4,715	4,715	15,432	11,488	11,488

Note: this table reports the full regression results for the variable *Negative emotions index*. For each time frame, we present the results of three specifications: (i) without controls; (ii) with controls; (iii) without controls but based on the same sample as in the full control specification. Standard errors are clustered at the attack-by-region level and are reported in parentheses; * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Source: authors' calculations based on CMS data.

D Theoretical model

The model follows the basic set-up in Becker and Rubinstein (2011), and we offer some extensions for our setting. Consider an economy that consists of individuals who consume a good (x) and are exposed to a terrorist attack. The attack provides disutility itself and via the creation of fear and anxiety, which in turn exaggerates subjective beliefs about the probability of surviving future attacks. This fear is driven by both media coverage and the severity of the attack. Importantly, we show how fear can (or cannot) vary over time and space in response to the attack.

Individual's expected utility is given by:

$$W = p(\tau, F) + V(x) \tag{D1}$$

where p is the subjective probability of surviving a terrorist attack and V is the utility from consumption of good x. The subjective probability is adversely affected by the degree of terrorism, τ, and negative emotions such as fear, F. It is also reasonable to assume that the severity of an attack of terrorism and fear are mutually reinforcing with respect to the subjective probability of survival:

$$p_\tau \leq 0, p_F \leq 0, p_{\tau F} \leq 0 \tag{D2}$$

The amount of fear one experiences is given by:

$$F(\tau, m) = f(\tau, m)(1 - T) \tag{D3}$$

where m represents media coverage of the terrorist attack and T is a variable that represents temporal distance from an attack, such that there is a linear decay in fear over time ($0 \leq T < 1$). Fear rises with the degree of terrorism ($f_\tau > 0$), and it is amplified by the attention drawn to the consequences of threat through propaganda or media coverage ($f_m > 0$). And, in the absence of terrorism, there is no fear, $f(0, m) = 0$. Indeed, we can also define an alternative equation for fear that accounts for non-linearities in the response to terror:

$$F(\tau, m) = f(\tau, m)h(T) \tag{D4}$$

where $h(T)$ captures a non-linear response (decay) of fear, which is possible because of framing effects or the responses of politicians, for instance. We can also introduce further shift parameters:

$$F(\tau, m) = f(\tau, m)(1 - T)(1 - D) \tag{D5}$$

where D represents the geographic distance from the terrorist attack. Now, fear is moderated by the individuals temporal and geographic distance from a terrorist incident. Similarly, it is reasonable to consider aggravating factors. Specifically, it is reasonable to assume that some attacks are so severe that their impacts transcend space and time:

$$F(\tau, m) = \begin{cases} f(\tau, m)(1 - T)(1 - D) & \text{if } \tau \neq 1 \\ f(\tau, m) & \text{if } \tau = 1 \end{cases} \tag{D6}$$

when τ is equal to 1, the most severe possible attack, the level of fear is not moderated by distance; i.e. the effects of the attack are homogeneous through space and time. Assuming a simple model of fear, as in Equation D3, expected utility is given by:

$$W^{0<T<1} = p(\tau, [f(\tau, m)(1 - T)]) + V(x), \quad W^{T=1} = p(\tau) + V(x) \tag{D7}$$

Therefore, expected utility is lower when an individual is temporally proximate to the terrorist attack due to the presence of fear:

$$W^{0<T<1} < W^{T=1} \tag{D8}$$

References

Abou-Chadi, T. (2016). 'Political and Institutional Determinants of Immigration Policies'. *Journal of Ethnic and Migration Studies*, 42(13): 2087–110. https://doi.org/10.1080/1369183X.2016.1166938

Agerberg, M., and Sohlberg, J. (2021). 'Personal Proximity and Reactions to Terrorism'. *Comparative Political Studies*, 54(14): 2512–45. https://doi.org/10.1177/0010414021997162

Athey, S., Tibshirani, J., and Wager, S. (2019). 'Generalized Random Forests'. *The Annals of Statistics*, 47(2): 1148–78. https://doi.org/10.1214/18-AOS1709

Becker, G. S., and Rubinstein, Y. (2011). 'Fear and the Response to Terrorism: An Economic Analysis'. CEP Discussion Paper 1079. London: Centre for Economic Performance, LSE.

Blackwell, M., Iacus, S., King, G., and Porro, G. (2009). 'cem: Coarsened Exact Matching in Stata'. *The Stata Journal*, 9(4): 524–46. https://doi.org/10.1177/1536867X0900900402

Böhmelt, T., Bove, V., and Nussio, E. (2020). 'Can Terrorism Abroad Influence Migration Attitudes at Home?' *American Journal of Political Science*, 64(3): 437–51. https://doi.org/10.1111/ajps.12494

Bove, V., Böhmelt, T., and Nussio, E. (2021). 'Terrorism Abroad and Migration Policies at Home'. *Journal of European Public Policy*, 28(2): 190–207. https://doi.org/10.1080/13501763.2020.1729227

Bove, V., Efthyvoulou, G., and Pickard, H. (2021). 'Did Terrorism Affect Voting in the Brexit Referendum?' *British Journal of Political Science*, First view(-): 1–18. https://doi.org/10.1017/S0007123420000927

Brück, T., and Müller, C. (2010). 'Comparing the Determinants of Concern about Terrorism and Crime'. *Global Crime*, 11(1): 1–15. https://doi.org/10.1080/17440570903475634

Cattaneo, M. D., Crump, R. K., Farrell, M. H., and Feng, Y. (2019a). *Binscatter Regressions.* https://doi.org/10.48550/arXiv.1902.09615

Cattaneo, M. D., Crump, R. K., Farrell, M. H., and Feng, Y. (2019b). *On Binscatter.* https://doi.org/10.48550/arXiv.1902.09608

Hainmueller, J. (2012). 'Entropy Balancing for Causal Effects: A Multivariate Reweighting Method to Produce Balanced Samples in Observational Studies'. *Political Analysis*, 20(1): 25–46. https://doi.org/10.1093/pan/mpr025

Helbling, M., and Kalkum, D. (2018). 'Migration Policy Trends in OECD Countries'. *Journal of European Public Policy*, 25(12): 1779–97. https://doi.org/10.1080/13501763.2017.1361466

Helbling, M., and Meierrieks, D. (2022). 'Terrorism and Migration: An Overview'. *British Journal of Political Science*, 52(2): 977–96. https://doi.org/10.1017/S0007123420000587

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. https://doi.org/10.1017/pan.2019.27

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): 900–12. https://doi.org/10.1093/poq/nfab046

Pickard, H., Efthyvoulou, G., and Bove, V. (2022). 'What's Left After Right-Wing Extremism? The Effects on Political Orientation'. *European Journal of Political Research*, Early view(-): 1–14. https://doi.org/10.1111/1475-6765.12538

Timoneda, J. C. (2021). 'Estimating Group Fixed Effects in Panel Data with a Binary Dependent Variable: How the LPM Outperforms Logistic Regression in Rare Events Data'. *Social Science Research*, 93(-): 102486. https://doi.org/10.1016/j.ssresearch.2020.102486