UNIVERSITY OF SUSSEX

ECONOMICS

Working Paper Series

No. 12-2018

Terrorist Attacks and Immigration Rhetoric: A Natural Experiment on British MPs

Daniele Guariso

Department of Economics, University of Sussex, Brighton, UK

d.guariso@sussex.ac.uk

Abstract: We study the effects of exogenous shocks on the rhetoric of British politicians on social media. In particular, we focus on the impact of terrorist attacks on the issue of immigration. For this purpose, we collect all the immigration-related Tweets from the active Twitter accounts of MPs using Web Scraping and Machine Learning techniques. Looking at the Manchester bombing of 2017 as our main Event Study, we detect a counterintuitive finding: a substantial decrease in the expected number of immigration-related Tweets occurred after the incident. We hypothesize that this "muting effect" results from risk-averse strategic behaviour of politicians during the election campaign. However, the MPs' response shows remarkable heterogeneity according to the socio-economic characteristics of their constituencies.

JEL classification: C81; D72; Z13

Key words: political behaviour; machine learning; social media; immigration; terrorism

Acknowledgements: I wish to thank Prof Richard Dickens, Dr Vikram Pathania and the participants to the 2018 PolText Incubator Workshop in Budapest for useful comments and discussions. I am also very grateful to Prof David Weir and Simon Wibberley that introduced me to Computational Linguistics and Method52, the software developed by the Text Analysis Group at the University of Sussex used to carry out the analyses in the present study. All errors are my own.

Contents

List of Tables

List of Figures

1 Introduction

We analyse the consequences of two acts of terrorism occurred in the UK in 2017 on the rhetoric of the British Members of Parliament (MPs). In particular, we will focus on the issue of immigration, as several studies have shown how these dramatic events might shape public attitudes towards out-groups (e.g., see Legewie (2013)). If opinions and beliefs are indeed affected, it appears of interest to understand which information channels could mediate this effect. The existing literature has studied in depth how media depict immigrants and how the frame they provide can influence people's opinions (Allen and Blinder (2013); Brader et al. (2008)). However, these studies seem to neglect the role that politicians might play in the process of attitudes' formation. In fact, political leaders could get leverage on these events for their own advantage, especially when they face high stakes as during an election campaign. One of the obvious reason for this gap in the literature is the lack of suitable data, but also the challenges of conceiving an appropriate research design. In this paper, we try to overcome these hurdles. We focus on a specific information environment: the social media Twitter; given its increasing relevance as platform for news' provision and political campaining. Using Web Scraping, Machine Learning techniques and a Natural Experiment setting, we try to capture the change in the rhetoric of political elites caused by a terrorist attack. Politicians might exploit these dramatic events to foster the debate about immigration and divert attention from contextual problems. This would potentially create an implicit and dangerous link in the mind of the public between the threats posed by radical extremism and the presence of perceived outgroups in the community. Alternatively, they might seize the opportunity to signal their ideological stance or valence with respect to immigration. In both scenarios, we would expect an increase in the amount of relevant information provided by politicians in the aftermath of a terrorist attack. This in turn would have an impact on public attitudes, as the high level of anxiety induced in people by these dreadful events can enhance information-seeking (Gadarian and Albertson (2014)). However, what we observe instead is rather counterintuitive. The amount of relevant information, measured as Tweets related to immigration posted by a politician, actually decreases, on average, in the week after the terrorist attack. Moreover, when we focus on the event occurred during the election campaign, we find significant heterogeneity according to the characteristics of the MP or her constituency. The "muting effect" of the attack is more pronounced for politicians standing for marginal seats and elected in more restrictionist constituencies. In constrast, a smaller stock of foreign people and adverse economic conditions appear to lessen the "muting effect" on the expected number of immigration-related Tweets. Surprisingly, the political divide between MPs belonging to the incumbent government's party and members of the main challenger does not seem to matter for the response to the event.

The first attack we consider took place on the 22nd of March 2017. The 52 years-old Briton Khalid Masood drove a grey Hyundai Tucson into pedestrians along the pavement in proximity of the Palace of Westminster in London, the seat of the Parliament. The perpetrator killed three civilians and injured more than 40 people of different nationalities; he then left the car and stabbed to death a police officer before being shot. Another wounded woman died in hospital two days after.[1] The last message sent by the attacker stated that he was waging jihad in revenge for the Western interventions in the Middle East.[2] The Islamic State claimed responsibility for the act but no evidence emerged that backed up the allegation.[3] Prior to this attack, the last act of terrorism causing multiple casualties on the British mainland was the suicide bombing in London on the 7th of July 2005.

The second event is the bombing in Manchester occurred on the 22nd of May 2017. After the concert of the singer Ariana Grande, the 22 years-old British born Salman Ramadan Abedi detonated an explosive device in the foyer area of the Manchester Arena, causing the death of 22 people (10 of them aged under 20) and injuring more than 500.[4] ISIS claimed again responsibility, stating in a post on the social media that "one of the soldiers of the caliphate was able to place an explosive device within a gathering of the crusaders in the city of Manchester".[5] This second attack occurred after the announcement on the 18th of April of a snap election by the British Prime Minister Theresa May, whose stated purpose was to gain a large majority to strengthen her position in the upcoming Brexit negotiations.[6] The elections took place on the 8th of June. The majority Conservative government lost 13 seats (shifting to 317) and was forced to secure a confidence and supply deal with the Democratic Unionist Party.[7] The main challenger instead, the Labour party, won 262 seats, with a net gain of 30 seats from the previous elections.

There are several reasons for the choice of these two particular events. First, even if not strictly iden-

[1] Gribben, P. et al. "As it happened: Coverage of London attacks". *BBC News*. Retrieved 05 January 2018.

[2] Sengupta, K. (27 April 2017). "Last message left by Westminster attacker Khalid Masood uncovered by security agencies". *The Independent*. Retrieved 05 January 2018.

[3] Metropolitan Police. (27 March 2017). "Update: Westminster terror attack. Metropolitan Police News". Retrieved 05 January 2018.

[4] (1 November 2017). "Manchester attack: Extradition bid for Salman Abedi's brother". *BBC News*. Retrieved 05 January 2018.

Topping, A. (24 May 2017). "Go sing with the angels: families pay tribute to Manchester victims". *The Guardian*. Retrieved 05 January 2018.

[5] Yeginsu, C. and Erlanger, S. (23 May 2017). "ISIS Claims Responsibility for Manchester Concert Attack; Toll Rises to 22". *The New York Times*. Retrieved 05 January 2018.

[6] Birrell, I. (2 May 2017). "Strengthen our hand in Europe? No, a landslide for May would weaken it". *The Guardian*. Retrieved 05 January 2018.

[7] (26 June 2017). "Conservatives agree pact with DUP to support May government". *BBC News*. Retrieved 05 January 2018.

tical, they represent the same type of shock and share common characteristics, as the nature of the attack (religious radicalism) and origin of the offender (English). Second, the two acts of terrorism occurred within a relatively short period of time and the subjects of our Treatments all belong to the 56th UK Parliament, so they experienced the same institutional context. Finally, the two incidents embody two distinctive Treatment conditions, where the incentives faced by the MPs were substantially different. However, given the intrinsic interest in the high stakes faced by political elites during an election campaign, our main object of study will be the Manchester attack.

In order to assess the effects of such shocks on the behaviour of politicians, we will first revise the literature related to the determinants and correlates of public attitudes towards migration and the emotional and behavioural responses to terrorist attacks. We will also mention the increasing application of Machine Learning techniques in Economics, and how this study contributes to this growing body of research. We will then describe the methodology employed, how the data were gathered and the dataset constructed. We will present the features of our Classifier, the statistical model chosen and our Identification Strategy. We will subsequently provide descriptive statistics on the data used, together with the time trends and general tweeting behaviour of politicians. Next, we will move to the core of our paper with the analysis of the impact of terrorist attacks on the rhetoric of British politicians. In addition, we will explore several channels that might mediate heterogeneous effects among MPs. We will also perform Robustness Checks on our baseline estimations. Finally, we will complete the study by discussing our results and framing the direction for future research.

2 Literature Review

This study is placed at the crossroads of different bodies of research. International migration represents in our times one of the most challenging issues from a social, economic and political perspective. During the last 20 years, there has been a growing interest in the determinants and trends of public attitudes towards immigration (Ceobanu and Escandell (2010)) and how the rising inflows of out-groups might be correlated with the upsurge of populist and xenophobic movements in Europe (e.g., see Whitaker and Lynch (2011)). In a recent review of the literature, Hainmuller and Hopkins (2014) underline how perceived threats to intangible social constructs, like national economy or identity, are among the main shaping factors of attitudes' formation. In particular, the authors put emphasis on the perceptions of sociotropic threats, especially cultural, as opposed to concerns related to material self-interest.

At the same time, research has focused on exceptional circumstances that might endanger the integration of migrant people. This can be due to the emotional impact of such events or because they are perceived as signals of assimilation's failure. As a matter of fact, we observe a growing number of studies that analyse the social and psychological effects of terrorist attacks. These dreadful incidents can have substantial consequences on natives' attitudes and on perceptions of ethnic minorities and foreigners (Cohu et al. (2016); Legewie (2013); Schüller (2016)). In addition, they have major psychological repercussions. They affect risk perception and increase the willingness to trade off civil liberties for increased public security (Bozzoli and Müller (2011)), they negatively impact expectations (Coupe (2017)) and lead to high levels of anxiety and anger (Huddy et al. (2005); Vasilopoulos (2018)). Besides, these emotional reactions can pervade actual behaviour. Hanes and Machin (2014) document an exacerbation of hate crimes against Asians and Arabs following the bombings in London in 2005. Moreover, terrorist attacks, as dramatic and dreadful events, might question the effectiveness of government and the political system, leading to an impact on electoral outcomes (Montalvo (2011)). If these events have such grievous consequences on the social structure it seems sensible to try to understand through which channels they might affect public opinions. In fact, if sociotropic threats are more influential than egotropic ones, acts of terrorism could be exploited by political elites to shape mass attitudes towards immigration, appealing to social constructs such as national identity or local economic conditions. Alternatively, they can seek the opportunity to signal their position on the political spectrum about the issue. This could be even more relevant during elections; since, as shown by Kendall et al. (2014), even in the short run voters do update their beliefs when receiving new information on the ideology or valence of a candidate.

As a matter of fact, opinion leaders, and so political leaders, are likely to represent the main hubs in an information acquisition's network (Galeotti and Goyal (2010)). In this structure, they constitute core nodes with high in-degree: the central pillars in the information environment of voters. Their role becomes even more relevant in the aftermath of a terrorist attack. It has been shown through experimental evidence that the anxiety induced by unfamiliar threatening conditions triggers political response and information seeking, with a bias on threatening news (Brader et al. (2008); Gadarian and Albertson (2014)). Hence, voters in the wake of such events are likely to be much more sensitive to any information provided by opinion leaders, which in turn can flow through a variety of communication channels.

In the last decade, one of these channels, social media, has arised for its rapid diffusion and development and, at the same time, we have witnessed to the growing impact of these new media in news' provision. As an example, surveys from the Pew Research Center show that the share of U.S. adults getting news on social media has increased from 49% in 2012 to 62% in 2016 (Gottfried and Shearer (2016)). In addition, this trend has been matched with a widespread uptake of these platforms by political leaders, with a consequential effect on their electoral performance. Recent studies find a positive association between social media-based campaigning (specifically, the activity on Twitter) and voting outcomes in the U.K. (Bright et al. (2017)). Thus, it appears relevant to assess the role of these emerging information channels and their strategic use by politicians.

Research on the use of social media, and in particular Twitter, as a communication and electoral tool by political elites is still in early stages, but with a growing number of findings (see Jungherr (2016) for a survey of the literature). Interestingly, small sample studies suggest that the personal use of Twitter by politicians might diverge from what we would expect in a communication environment strategically coordinated, where members collectively advocate party policies (Adi et al. (2014)). At the same time, tweeting behaviour seems to transcend partisanship, and common patterns emerge, at least among major political parties (Evans et al. (2014)). In fact, this microblog becomes a channel for expressing individual lines on policy, due to the personalization that this particular hybrid platform allows. The construction of a blurred private/public personality by politicians is meant to induce empathy from voters and could in turn reflect strategic behaviour aimed at earning personal support (Jackson and Lilleker (2011)). This personalization of their professional figure is in line with Impression Management Theory (Jones and Pittman (1982)), which provides a taxonomy of attitudes through which individuals try to actively manage the public perception of themselves. Nonetheless, it appears that the focus of their messages has predominantly a political theme, especially during election campaigns (Evans et al. (2014)).

It has been observed that political elites use their Twitter account for constituency service: it is a convenient channel for reaching crucial audiences quickly and effectively. However, even if it seems mainly a unidirectional channel of communication, politicians do interact with the Twitter community in order to attack an opponent, debate, or taking a position on a specific issue (Graham et al. (2013)). Said that, the evidence shows that the microblog represents a powerful way of self-promotion, leveraged to maximize the impact on the electorate (Jackson and Lilleker (2011)).

At the same time, the influence of politicians on the information environment of the public is indirectly amplified. This peculiar channel offers the possibility of manipulating the flow of the national dialogue through its impact on the agenda of traditional mass media (Kreiss (2016)). Qualitative research shows how professional journalists do use Tweets from political leaders to shape their coverage in terms of issues and events. They also obtain from them background information, polling data and quotes that subsequently include in their articles. As a matter of fact, journalists are led to focus on topics that are salient for these opinion leaders, and they tend to incorporate the stance of the politicians on those issues (Parmelee (2014)). However, it is important to underline that the role that Twitter might play in traditional media's agenda-building is very context-dependent, and it is likely to change according to the institutional setting under analysis.

In any case, the freedom of expressing personal beliefs and opinions offered by this social media might turn out to be a double-edged sword. Since journalists do rely on this microblog as a way to monitor politicians' view and inform their agenda, the exposure of political leaders to criticism and attacks is magnified, leading to a careful use of the platform. In fact, it is not uncommon in the UK context that hasty Tweets led to subsequent public condemnation and pillory, requiring formal excuses. As an example, in 2013 the Prime Minister David Cameron had to face open criticism after that a member of his staff endorsed by error an offensive Twitter account (Adi et al. (2014)). It is clear that in such a setting unexpected events might engender strategic responses, especially when stakes are high, as during an election campaign.

While the role of traditional mass media in depicting migrants has been analysed extensively (e.g., see Allen and Blinder (2013)), to the best of our knowledge it appears that poor quantitative research has been conducted on the role and behaviour of political elites and their strategic use of social media.

This paper thus tries to partially fill this gap by proposing one of the first empirical studies on the effects of terrorist attacks on the immigration rhetoric of politicians.

From a methodological perspective, this work adds to the growing literature on applications of Machine Learning techniques in economic and social research. In recent years, we have witnessed to the emergence of a data revolution and increasing availability of large-scale granular information on previously unmeasured activities. Economics has thus expanded the traditional Econometrics techniques with Data Mining methods (e.g. Learning Algorithms) that often can complement the standard toolkit of an empirical economist (Einav and Levin (2014)). Nowadays, Machine Learning Algorithms are spreading in different fields of Economics and Political Science, often with the aim of selecting the relevant covariates in an empirical model (Belloni et al. (2014)) or capturing heterogeneous treatment effects (Wager and Athey (2017)). Applications range from predicting consumer demand (Bajari et al. (2015)) to test theories of risky and ambiguous behaviour (Peysakhovich and Naecker (2017)), with an increasing emphasis on estimating causal effects (Athey and Imbens (2015)). This paper is thus an attempt to combine what Leo Breiman called the two cultures of statistical modeling (Breiman et al. (2001)). The first one, based on stochastic data models, aims to capture causal relationships between variables. The second one employs Learning Algorithms to maximize the accuracy of out-of-sample predictions. In our analysis, we leverage the latter to improve the quality of the data used. In addition, the granularity of the information retrieved is exploited to frame a Natural Experiment design that is likely to allow the interpretation of the parameter of interest as a causal effect. However, this relationship is estimated through a standard statistical model, which is meant to describe the underlying data-generating process.

3 Methodology

The methodology employed in this paper can be divided into three parts. The first illustrates the process of collecting the relevant information and constructing the dataset, whereas the second one presents the statistical model preferred to carry out the empirical analysis. We then conclude by describing the Natural Experiment setting of the study.

3.1 Construction of the Dataset

The construction of the dataset can be further broken down into two steps. In the first one we collected all the most recent Tweets of British MPs with an active Twitter account through the Twitter API. One limitation is that only the latest 3,200 Tweets can be collected per account.[8] However, as explained later on this Section, only the Members of the Parliament for whom we have complete information for the different time periods were considered in the analyses. The collection was executed on the 27th of September 2017, bringing 1,504,088 Tweets. We then started to select all the relevant Tweets through a Boolean Search. Relevant Tweets are defined as all the Tweets containing one or more of the words listed in Appendix A.[9] The choice was mainly informed by the report of Allen and Blinder (2013) that documents all the major correlates of the terms immigrants, migrants, asylum seekers and refugees (or variations) on the 20 main British national newspapers between 2010 and 2012. These terms plus synonyms of the dominant correlates and other current relevant words (e.g. free movement) form the final list. The amount of data was thus reduced to around 20,600 Tweets. However, the dataset to this point still contained lots of Tweets irrelevant for our analysis, due to the variability in the semantics of the chosen words in different textual corpora. Figure 7 in Appendix C provides some examples of these problematic Tweets. Hence, we further improved the quality of our dataset by relying on Machine Learning (ML) techniques. In the subsequent step, we trained a Classifier that was able to effectively reproduce our decision-making process and distinguish between Tweets that were actually relevant to our research from those that were not. The underlying predictive model is a Semi-supervised Multinomial Naïve Bayes coupled with the Feature Marginals (FM) algorithm, as proposed by Lucas and Downey (2013).[10] A Naïve Bayes model was preferred as it is relatively fast to train, it has been

[8] Native Retweets are counted in this total (https://developer.twitter.com/en/docs/tweets/timelines/api-reference/get-statuses-user_timeline.html).

[9] The Tweets containing words for which those chosen are substrings were retained as well (e.g. all the Tweets with the word immigrant were preserved, since they contain the word migrant, that is present in our list).

[10] Semi-supervised Learning is an approach that tries to leverage the information on both the unlabelled and labelled data in order to learn the target function (Lucas and Downey (2013)).

proven effective for Text Analysis and it suits well Semi-supervised Learning algorithms (Kober and Weir (2015)). FM was chosen as it has been shown to outperform other standard algorithms both in Text Topic Classification and Sentiment Analysis (Lucas and Downey (2013)). Its main specificity being that it does not have to iteratively compute multiple passes over the unlabelled data for each new task (contrary to the Expectation-Maximization algorithm, for example). It instead precomputes a set of statistics (i.e. the marginal probability of each word) over the unlabelled data in advance. These statistics are then used as constraints in the optimization problem, in order to improve the estimates of the class-conditional probability of each word. This procedure is particularly suitable for improving the estimates of words that have not been seen in the labelled data (i.e. the training set). The classifier was trained on a subset of human-coded Tweets made up by 600 items. Data were pre-processed by normalising URLs, punctuation was filtered and all tokens were made lowercase before extracting the features. Bigrams and trigrams were used in addition to unigrams as features for the classification, to capture more complex grammatical structures.

According to the ML literature the quality of an algorithm is assessed through its out-of-sample prediction performance using Cross-Validation (Varian (2014)). Hence, we manually labelled a subset of 900 Tweets (630 were relevant, 270 were not) in order to form a Gold Standard dataset against which the Classifier was tested. Table 1 reports the performance of the Classifier. The overall accuracy of the Classifier (i.e. the proportion of Tweets which were assigned to the correct category) is 0.878, well above the value of 0.7 recommended by Rijsbergen (1979) for scientific research. The Precision value states the proportion of all documents the Classifier believed belonging to a given class that were truly belonging to that category. Using standard hypothesis testing notation, this can be thought of as $(1 - \alpha)$, that is one minus the likelihood of a Type I Error. The Recall value is the proportion of all documents belonging to a particular category, which the Classifier labelled as belonging to that class. It can be thought of as the power of the test $(1 - \beta)$. The F-Score represents the harmonic average of Precision and Recall. However, for the purpose of our analysis, the most significant statistics is 1-Precision for the Relevant class, as it captures the proportion of False Positives for that category in our classification exercise. The proportion of Tweets erroneously labelled as Relevant by our Classifier was less than 10%, a rather small value (examples are provided in Figure 9 in Appendix C).[11] The final dataset was made up by 14,817 immigration-related Tweets, spanning a period of over 9 years.

[11] We highlight the significance of the proportion of False Positives for the Relevant category. However, also the False Positives for the Irrelevant class are meaningful, as they point out that our analysis can only be a lower-bound estimation. Examples of this kind of Tweets are shown in Figure 8 in Appendix C.

Table 1 – Classifier Performance

Label	Precision	Recall	F-Score	Accuracy
Relevant	0.908	0.919	0.913	
Irrelevant	0.805	0.781	0.793	
				0.878

It must be noted that our Relevant category captures all texts generally related to the issue of immigration posted by an English MP over the time period considered in the study. Thus, we do not perform a Sentiment Analysis, as we do not discriminate between Tweets that contain positive, negative or neutral polarity with respect to this topic.[12] The choice was driven by the potential deficiencies of this type of analysis in our particular setting. The performance of a Sentiment Classifier is crucially dependent on the domain's consistency of the data used (Barbieri et al. (2015); von Grünigen et al. (2017)) and the context of the words in a given textual corpus (Saif et al. (2016); Teng et al. (2016)). Even if the former is well-defined through our two-step procedure, the latter is very likely to change in the aftermath of a terrorist attack, given the emotional reactions that such events engender. Thus, a Sentiment Analysis could mistakenly interpret a shift in context and the choice of words used to convey feelings as a change in the amount of polarity-related Tweets.

3.2 Statistical Model

To model the data-generating process we opted for the Zero-Inflated Negative Binomial, as it is suitable for over-dispersed data (i.e. the conditional mean is not equal to the conditional variance) that present excess of zeros (Cameron and Trivedi (2013)). The main idea underlying this model is to include a separate component (π) that inflates the likelihood of observing a zero. Thus, the ZINB assumes that the zero observations arise from two different sources, a structural one (given byπ) and a sampling one (given by the base count density $f_2(y)$) (Hu et al. (2011)). Equation 1 presents the generalization of the model, in our application the base count density $f_2(y)$ is a NB2 (Hilbe (2011)).

$$Pr\left[y=j\right]=\begin{cases}\pi+(1-\pi)\,f_2(0) & \text{if } j=0\\ (1-\pi)\,f_2(j) & \text{if } j>0\end{cases} \tag{1}$$

The inflation factor π might be a constant or depend on a set of regressors in a binary outcome model. In our case, the inflation factor is a (logistic) function of the total number of Tweets posted by the MP

[12] See Figure 10 in Appendix C for examples of Tweets with different polarity.

in a given day. The insight is simple: the likelihood of observing a non-zero for our dependent variable of interest (the total number of immigration-related Tweets) is correlated to the daily Twitter activity of the politician. The more she tweets, the more likely she is to talk soon or later about immigration. Moreover, the ZINB has already been used in other scientific fields to model Twitter data (e.g., see Williams and Burnap (2015)). In Section 5, as a Robustness Check, we estimate our baseline model using a standard Negative Binomial and adding the total number of Tweets at the MP level as a covariate.

3.3 Natural Experiment Setting

We study the effects of terrorist attacks on the number of immigration-related Tweets posted by an English MP on a given day. In order to accomplish this purpose, we exploit the panel structure of our dataset in an Event Study framework. Due to the exogeneity in the timing of these acts of terrorism, our estimates are likely to provide the Average Effect of the Treatment (the attack) on the Treated (the MPs). However, the time windows chosen are crucial for our Identification Strategy. It has been noted that Twitter data are particularly volatile, and messages are generally triggered by specific events related to the topic under study (Wibberley et al. (2014)). Thus, we eliminate from our analysis two main events directly related to the issue of immigration that caused a peak in the discussions of politicians about this topic.[13] On the 7th of March, the amendment proposed by the Conservative MP Heidi Allen to properly audit local council capacity to house unaccompanied child refugees was defeated in the Parliament. The amendment was grounded on the Home Office's sudden abandonment of the Dubs scheme on refugees occurred in February and caused a significant contention on the issue among MPs.[14] The other triggering event was on the 29th of May: The Battle for Number 10, a live TV debate between the incumbent Prime Minister Theresa May and the leader of the Labour Party Jeremy Corbyn. During the discussion, immigration was a major theme of confrontation, prompting all genres of remarks along the political spectrum.[15] We exclude the day of the first event from the analysis of the Westminster attack, whereas The Battle for Number 10 will constitute the upper temporal limit for the study of the Manchester bombing. The main purpose of these omissions is to capture the effect on the average tweeting behaviour of the MPs. It is also worth mentioning that there

[13] See the Section 4 for an overview of the Tweets time trends.

[14] (7 March 2017). "Dubs amendment: Child migrant challenge defeated by MPs". *BBC News*. Retrieved 05 January 2018.

[15] Sparrow, A. (30 May 2017). "Labour and Tory leaders interviewed by Jeremy Paxman - as it happened". *The Guardian*. Retrieved 05 January 2018.

is no direct reason why the two terrorist attacks (both committed by British citizens) should provoke a change in the immigration rhetoric of the MPs, apart from political gain. However, even if we were to assume an effect, we would not expect a long-lasting impact: Issue-Attention Cycle Theory posits that public attention to even major social problems suddenly peaks, but then rapidly fades away (Downs (1972)). This hypothesis is consistent with the empirical results of Legewie (2013), and Williams and Burnap (2015), which document how the emotional and attitudinal effects of terrorist attacks are quite short-lived. Hence, we mainly expect a reaction from politicians only in the immediate aftermath of the incident. Thus, we will study a time interval that looks at the week after the event (including the day of the attack) and the week before. We will then expand it by looking also at two and three weeks prior to the incident. One main constraint of the analysis is that the further we extend our time interval the more likely we are to capture other, even if less known, triggering events, that might be systematically related to our response variable conditional on the attack (the temporal stability assumption of Legewie (2013)). Figure 1 shows a timeline of the relevant events considered in our analysis.

Figure 1 – Timeline

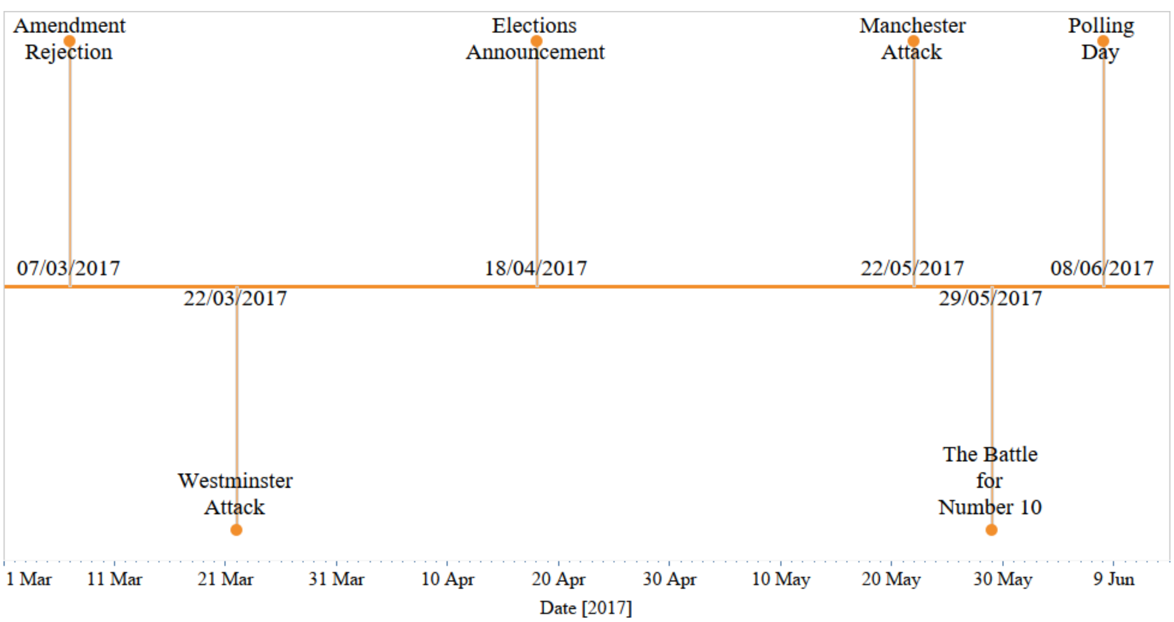

Another important point to mention is the number of cross-sectional observations considered in the following analyses. As already noted above, we were only able to gather the 3,200 most recent Tweets for each active account. This constraint is reflected in the number of observations available for the two events. When comparing the two incidents, we will only consider the MPs for whom we have full information for both the attacks (519).[16] When we will focus on Manchester, we will consider all the available active Twitter accounts for which we have information (548). In any case we look at a sizeable proportion of the members of The House of Commons.[17] Appendix B provides further descriptive statistics on the MPs included in our analyses and those excluded.

[16] Actually, in the comparison we also consider four more accounts that were created meanwhile when looking at Manchester.

[17] The Lower House of the British Parliament has 650 members. Thus, we analyse around 80% of them in the comparison and 84% when we just focus on Manchester.

4 Data Description

4.1 Twitter Data

The dataset employed in the analysis has a longitudinal structure. The cross-sectional unit of observation is a Member of the House of Commons. We record daily Twitter activity related to the issue of immigration as described in Section 3. Thus, our main dependent variable of interest is the number of immigration-related Tweets posted by an English MP in a given day. We also retain the daily number of Tweets posted by the MP and information on the account, as the number of followers, number of friends, number of statuses,[18] and the age. The characteristics of the account might be important correlates of the tweeting behaviour of the politicians, so we decide to keep them as controls. We also add demographics; as previous studies have shown how age and gender might affect the use of the microblog, in particular when considering issue-specific Tweets (Evans et al. (2014); Jackson and Lilleker (2011)).

4.2 Data at the Constituency Level

In order to explore the heterogeneity of the effect across MPs, we gather information on their constituencies. We retain the majority share of the incumbent MPs in 2015 general elections and their betting odds for the 2017 elections. We collect a proxy of the average unemployment level in 2016, measured as the proportion of economically active 16-64 years-old residents claiming Jobseeker's Allowance. These last data come from the ONS Nomis database, and are meant to capture the perception of local competition over scarce resources. From the British Election Study (2017 results) we collect the share of UK-born and the share of people of white British ethnicity as measures of intergroup contact.[19] This database also contains the estimates of the results for the 2016 EU Referendum at the constituency level computed by Hanretty (2017). This measure is meant to capture the salience of the issue of immigration in a given constituency. Unfortunately, the BES does not report information on Northern Ireland, so we have to systematically exclude its 18 constituencies (16 of them included in our dataset) when analysing the effect of these last variables.

Table 2 presents descriptive statistics for the independent variables employed in the study.

[18] These three variables change over time, but in our dataset are fixed, as they report the value on the day of the collection (27/09/2017). However, they are a good proxy for the type of node that the MP represents in the network structure of the Twitter community.

[19] It is worth noting that these two proxies might not precisely capture the same concept. In fact, the proportion of UK-born also includes second generation migrants, so it does not distinguish between multiple ethnic groups.

Table 2 – Summary Statistics (MP/Constituency Level) | 16/02/2017-01/06/2017

Variable	Mean	Std. Dev.	Min.	Max.	N
Followers Count	25330.895	82282.903	329	1538565	58512
Friends Count	1792.217	2882.238	0	38861	58512
Statuses Count	8799.687	10891.67	16	82326	58512
Male	0.67	0.47	0	1	58512
Age	50.768	10.149	22	82	58512
Age Account	5.423	2.074	0	9	58206
Majority Share (2015)	23.536	14.079	0.1	72.3	58512
Betting Odds (MP)	0.278	0.757	0.002	9	55120
Unemployment 2016 (avg.)	2.578	1.509	0.494	9.737	58512
Leave Share	51.603	11.673	20.481	75.650	56816
White British Ethnicity (%)	82.667	18.763	12.712	97.792	56816
UK Born (%)	87.960	11.639	40.728	98.018	56816

4.3 Tweeting Trends

One first important question that we might want to ask is if Twitter is a meaningful way to capture politicians' rhetoric, and if these opinion leaders really use this channel to communicate with their electorate. Figure 2 shows the time trend of the total number of Tweets for the MPs on whom we have information on the whole time period considered. We can clearly see that, after the elections announcement, the average number of Tweets substantially increases and no longer displays that seasonal pattern observed before. It instead presents small fluctuations around a higher grand level until the day of the Manchester attack, when it drops dramatically. A similar decrease seems to occur after Westminster. It appears that, at least during the election campaign, the MPs did increase their use of Twitter, presumably to get more in touch with their voters and promote themselves.

Figure 2 – Total Volume of Tweets by Day

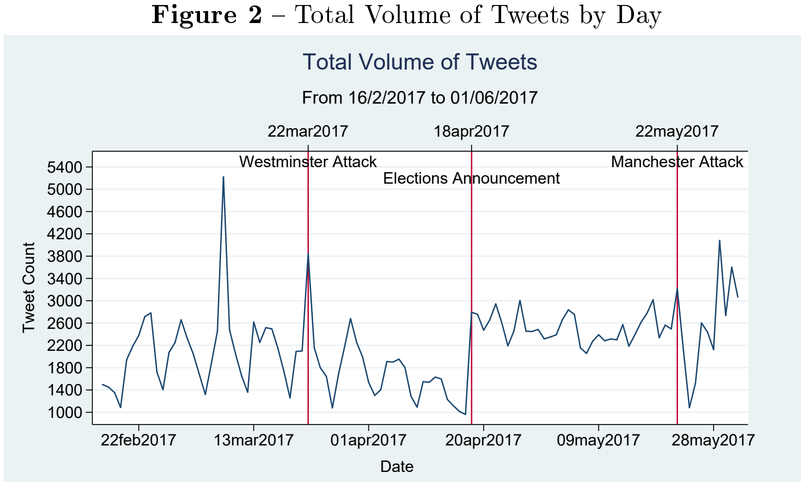

If we look at the trend of the immigration-related Tweets in Figure 3, we do not observe, on average, a significant rise during the elections. Hence, it does not seem that this was a topic particularly highlighted by the MPs in their campaign strategies, and this might be a consequence of immigration being a rather controversial and risky theme. As already mentioned in Section 3, we have two major peaks: one in correspondence of the rejection of the Dubs Scheme's amendment, the other one on the day of The Battle for Number 10. We do notice a fall after the Manchester attack, even if it does not seem as dramatic as for the total number of Tweets, but this might be due to the great difference in absolute values. The Westminster attack seems to cause a very short-lived drop, but the effect is not as clear as for the overall volume. However, these are just aggregate trends and they tell us nothing about the individual behaviour of politicians.

Figure 3 – Total Volume of Immigration-Related Tweets by Day

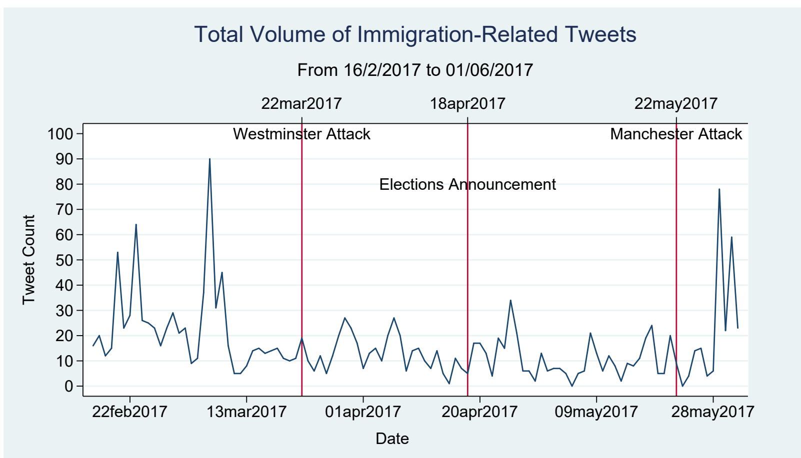

Another important point to underline is the volatility of tweeting behaviour. As crude measure, Table 3 reports the R^2 values of simple OLS regressions that capture the probability that a politician tweets. The dependent variable is a Dummy that takes the value of 1 if the MP tweets in a given day and the independent variables are a full set of individual and day Fixed Effects. This exercise is carried out for both pre and post elections announcement periods. What we observe is a rather random behaviour in the likelihood of tweeting. The full set of covariates is able to explain less than 40% of the variation in both cases. Thus, we do not expect great predictive power from our models and quite noisy estimates. However, given the number of unobservables in play, observing clearly a significant effect would mean capturing a rather relevant result.

Table 3 – Randomness in Tweeting Behaviour

	(1) Pre-Announcement	(2) Elections Campaign
R Squared	.37	.39
Observations	31,408	22,495
MP FEs	YES	YES
Day FEs	YES	YES

Note: Dependent variable is a dummy for tweeting or not in a given day.
Period before the announcement is 17/02/2017-17/04/2017.
Period during the elections is 18/04/2017-28/05/2017.

One more pattern that might be interesting to analyse is the difference in tweeting behaviour across political parties. Figure 4 presents the average daily tweeting activity for MPs by political affiliation. The values refer to 30 days before the elections announcement and 30 days after. As already noted with the time trends, we see a clear rise in the post-announcement period, and this is true for almost all political affiliations. We also observe that between the two major political parties, Labour presents systematically higher values than the Conservative. In addition, Plaid Cymru exhibits the most significant increase, with an average number of Tweets more than doubled after the announcement. These patterns seem to confirm the relevance of Twitter as a channel of communication and information exchange between the politicians and their electorate.

Figure 4 – Daily Tweets Average by Political Party

(a) Pre-Announcement Period

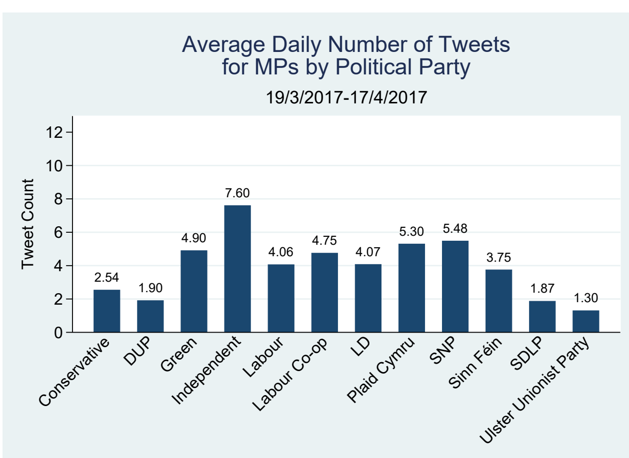

(b) Post-Announcement Period

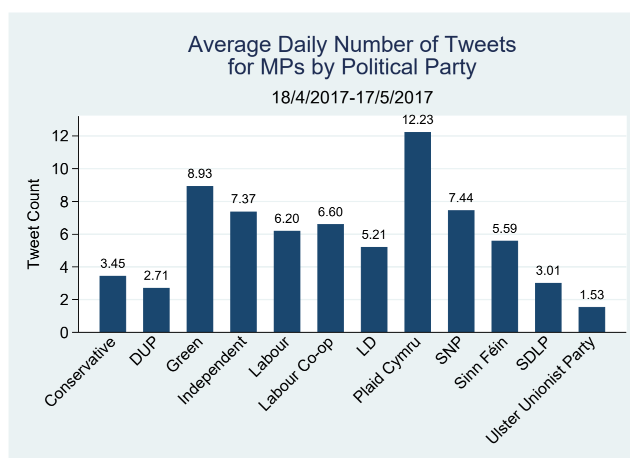

5 Analysis

5.1 Baseline Results

We now present the baseline results on the impact of the two acts of terrorism on the total number of Tweets and the number of immigration-related Tweets. In the next section, we will focus on the Manchester attack and we will explore the heterogeneity of the effect according to the different characteristics of the politician or her constituency.

5.1.1 Effect of the Terrorist Attacks on the Total Number of Tweets

In Table 4 we explore the impact of the terrorist incidents of Westminster and Manchester on the daily tweeting activity of British MPs. We look at different time intervals: 3, 2 and 1 week before the attack, but we compare them only with the first week following the event, each time. Our Treatment is a Dummy that takes the value of 1 on the day of the incident and the following six days. In each regression we add a Dummy for being male, the age of the politician, number of followers, number of friends, number of statuses, and age of the account as covariates. Day-of-the-week Fixed Effects are included, in order to capture weekly seasonality in tweeting behaviour. The models are estimated through a Negative Binomial and errors are clustered at the MP level. Table 4 shows the Incident Rate Ratios for our Treatment and the p-values are reported in parentheses. The last table shows the estimates for the Manchester attack when we consider all the available MPs. As mentioned in Section 3, the 7th of March is not considered in the estimation of the Westminster attack. The baseline regression model is presented in Equation 2, where $\mathbf{x}$ contains our control variables.

$$\mathrm{E}\left[TweetsCount_{it} | Treatment_t, \mathbf{x}_{it} \right] = \exp\left(\alpha + \beta\, Treatment_t + \mathbf{x}_{it}' \boldsymbol{\gamma} \right) \qquad (2)$$

What we observe is a clear decrease in the number of Tweets after both the events. However, the effect is definitely more pronounced for the Manchester attack (a reduction between 11% and 19% in the expected number of Tweets) and it is always highly significant in every time interval. In addition, the magnitude of the effect fades away as we extend our time period. For Westminster, the pattern is less clear. The impact appears to be not significant in the proximity of the event, but gains relevance as we enlarge the period analysed, with a magnitude that is less than 11%.[20]

Table 4 – Effect on Total Number of Tweets

	(1) 1 Week/1 Week	(2) 2 Weeks/1 Week	(3) 3 Weeks/1 Week
Westminster Attack	0.972	0.892	0.918
	(0.386)	(0.000)	(0.002)
Observations	7,266	10,892	14,000

Note: Treatment takes the value of 1 on the day of the Event and the subsequent 6 days. Day-of-the-Week Dummy, Sex, Age, Number of Followers, Number of Friends, Number of Statuses, and Age of the Account included. Errors are clustered at the MP level. IRR and p-values are reported. 07/03/2017 not included.

	(1) 1 Week/1 Week	(2) 2 Weeks/1 Week	(3) 3 Weeks/1 Week
Manchester Attack	0.818	0.876	0.885
	(0.000)	(0.000)	(0.000)
Observations	7,322	10,983	14,643

Note: Treatment takes the value of 1 on the day of the Event and the subsequent 6 days. Day-of-the-Week Dummy, Sex, Age, Number of Followers, Number of Friends, Number of Statuses, and Age of the Account included. Errors are clustered at the MP level. IRR and p-values are reported.

	(1) 1 Week/1 Week	(2) 2 Weeks/1 Week	(3) 3 Weeks/1 Week
Manchester Attack (All Obs.)	0.822	0.875	0.882
	(0.000)	(0.000)	(0.000)
Observations	7,672	11,508	15,343

Note: Treatment takes the value of 1 on the day of the Event and the subsequent 6 days. Day-of-the-Week Dummy, Sex, Age, Number of Followers, Number of Friends, Number of Statuses, and Age of the Account included. Errors are clustered at the MP level. IRR and p-values are reported. All the available MPs are used.

[20] As Incident Rate Ratios are just the exponentiated coefficients, the proportional change in the expected count is simply given by $\left(\exp^{\beta} - 1\right) \times 100\%$, where β represents the coefficient of interest. Taking as an example the third column in the first section of Table 4, the proportional change is computed as $(0.918 - 1) \times 100\% = -8.2\%$.

5.1.2 Effect of the Terrorist Attacks on the Number of Immigration-Related Tweets

Now we move to the core of the analysis; the dependent variable of interest is the daily number of immigration-related Tweets posted by a politician. Our goal is to capture the impact of terrorist attacks on the rhetoric of British political elites. The features of the analysis are the same as in the previous Subsection, but we now estimate our models through a Zero-Inflated Negative Binomial, where the inflation factor π is a function of the total number of Tweets posted by a politician in a given day. The results are presented in Table 5. It is important to reiterate that, as these dramatic incidents are not directly related with immigration, there is no reason why we should observe a distinct effect at all, apart from strategic behaviour. Moreover, if politicians wanted to seek the opportunity to signal their valence or ideology with respect to this issue, we should expect an increase in the number of immigration-related Tweets. We would also expect the same outcome if MPs were trying to shape public attitudes on the theme. Considering the Westminster attack, the effect is rather imprecisely estimated, and it is only marginally significant when we look at the three weeks before the incident, but always implying a reduction. The results for the Manchester attack appear more accurate. The effect is quite substantial in magnitude: a reduction of around 30% in the expected number of immigration-related Tweets when comparing one week before to one week after, that decreases to around 27% when considering the two weeks prior to the event. Both the effects are statistically significant at the 5% level. The effect increases substantially when we consider all the available MPs for the Manchester attack: a decrease of around 38% compared to the week before, slightly less (around 36%) when considering the two weeks previous to the incident. Both the effects are statistically significant at the 1% level. The pattern appears to be quite similar to that of the total number of Tweets: the impact fades away as we extend our time interval. However, the magnitude of the shock is proportionally greater. Table 6 reports the results of a one-sided Z-test under the null hypothesis that the coefficient of the Treatment effect on the immigration-related Tweets is less than or equal to the coefficient of the effect on the total number of Tweets. We cannot reject the null hypothesis for every time interval at any standard significance level: the impact of the attack appears to be more negative (i.e. greater in absolute value) for the immigration-related Tweets.[21] Hence, what we seem to capture is a rather counterintuitive "muting effect": a substantial reduction in the number of immigration-related Tweets following an act of terrorism, and this appears to be particularly true during elections.

However, it is important to underline at this point that the two terrorist incidents are not strictly

[21] Here we are comparing the results for the Manchester attack with all the available MPs.

comparable. The Westminster attack was the first act of terrorism on British mainland after almost twelve years, whereas the Manchester one had definitely a greater death toll, and many of the individuals involved were young people, so the emotional reactions are likely to be different. Moreover, the second attack occurred during an election campaign.

A possible explanation for the observed behaviour is a risk-averse strategy adopted by the politicians. Being aware of the unpredictable reactions and emotional distress of their electorate, and knowing the potential link between terrorist attacks and attitudes towards immigration, they prefer not to expose themselves and being on the safe side by neglecting the topic in the aftermath of the event. Moreover, the difference in the estimated effect between the two episodes derserves further considerations. It appears that the "muting effect" on the immigration-related Tweets is clearly observed only for the second attack. This could be the result of a dynamic process, in which political leaders learn to avoid risky issues and tend to maximize this behaviour in high stakes situations, like elections. We will further examine this hypothesis in Section 6.

In order to study the heterogeneity of the impact across different characteristics of the MPs or their constituencies, we will now focus on the Manchester attack, as it occured during an election campaign and it is thus more suitable to analyse the different incentives that politicians might face.

Table 5 – Effect on Total Number of Immigration-Related Tweets

	(1) 1 Week/1 Week	(2) 2 Weeks/1 Week	(3) 3 Weeks/1 Week
Westminster Attack	0.966	0.895	0.785
	(0.852)	(0.457)	(0.099)
Observations	7,266	10,892	14,000

Note: Treatment takes the value of 1 on the day of the Event and the subsequent 6 days. Day-of-the-Week Dummy, Sex, Age, Number of Followers, Number of Friends, Number of Statuses, and Age of the Account included. Errors are clustered at the MP level. IRR and p-values are reported. 07/03/2017 not included.

	(1) 1 Week/1 Week	(2) 2 Weeks/1 Week	(3) 3 Weeks/1 Week
Manchester Attack	0.704	0.734	0.913
	(0.030)	(0.048)	(0.543)
Observations	7,322	10,983	14,643

Note: Treatment takes the value of 1 on the day of the Event and the subsequent 6 days. Day-of-the-Week Dummy, Sex, Age, Number of Followers, Number of Friends, Number of Statuses, and Age of the Account included. Errors are clustered at the MP level. IRR and p-values are reported.

	(1) 1 Week/1 Week	(2) 2 Weeks/1 Week	(3) 3 Weeks/1 Week
Manchester Attack (All Obs.)	0.623	0.638	0.797
	(0.002)	(0.002)	(0.103)
Observations	7,672	11,508	15,343

Note: Treatment takes the value of 1 on the day of the Event and the subsequent 6 days. Day-of-the-Week Dummy, Sex, Age, Number of Followers, Number of Friends, Number of Statuses, and Age of the Account included. Errors are clustered at the MP level. IRR and p-values are reported. All the available MPs are used.

Table 6 – One-sided Z-Test $H_0 : \beta_{Immigration} \leq \beta_{Total}$

	1 Week/1 Week	2 Weeks/1 Week	3 Weeks/1 Week
Z Statistic	-1.75	-2.16	-0.72
P-Value	0.96	0.98	0.76

Note: One-sided z test. Null hypothesis: Coefficient of the Treatment effect on Immigration-related Tweets equal or less than the coefficient of the Treatment effect on Total Tweets. Manchester attack, all the available MPs are used.

5.2 Heterogeneity

We now focus on the Manchester bombing and try to capture potential channels of heterogeneity in the effect among the MPs. We will use all the available MPs as we no longer consider the two events. We will explore the following factors that could mediate the effect: the "safety" of a politician's seat, intergroup contact, competition over scarce resources, incumbency and the salience of the topic in the constituency.

The first channel that we are going to analyse is the relative strength of a MP's position in her constituency. It should be underlined that the expected sign of our Treatment effect is not straightforward. On one hand, we might think that the marginal utility coming from an additional Tweet for those MPs with a safe seat is lower, so they will tend to ignore the issue of immigration. On the other hand, these politicians might be willing to take a stance even on the riskier topics, due to the strength of their position. The same reasoning, but with opposite effects, applies to MPs standing in marginal seats. In order to test these two contrasting hypoteses, we use two proxies for the relative risk of a politician's position: the majority share in the 2015 general elections and the last available betting odds at the constituency level.[22] We will focus just on those MPs standing in marginal constituencies for the 2017 elections, and perform our analysis on two different subsamples. The first one is defined by all MPs standing for re-election in those constituencies where their marginal share of votes in 2015 was less than 10%. The second one is restricted to only those MPs standing in constituencies where their betting odds were greater than 0.1. Thus, we are only considering those politicians with a risky seat. Table 7 and Table 9 report the results of our exercise. What we observe is a substantial increase in the absolute magnitude of the effect, and this seems to hold for both proxies and in every time interval. For instance, if we look at the narrowest time window for the subsample defined by the ma-

[22] Data on betting odds were retrieved on the 16th of January 2018 from BetOnPolitics.co.uk (now bettingpro.com).

jority share, the resulting reduction in the expected number of immigration-related Tweets is greater by around 20 percentage points compared to our baseline results. As in Subsection 5.1.2, we perform a one-sided Z-test to compare the size of the effect in the subsample considered with the full sample. The null hypothesis is that the coefficient of the Treatment effect on the immigration-related Tweets in the subsample is less than or equal to the coefficient of the effect in the full sample. Table 8 reports the results for the subsample defined by the majority share and Table 10 reports the results for the subsample defined by the betting odds. For both subsamples we cannot reject the null hypothesis for every time interval at any standard significance level. Hence, it appears that this risk-averse behaviour does depend on the relative strength of the politician and the "muting effect" of the terrorist attack is magnified for those leaders with a marginal seat. MPs tend to be even more cautious in their rhetoric when their position is not safe.

Table 7 – Heterogeneity by Majority Share

	(1) 1 Week/1 Week	(2) 2 Weeks/1 Week	(3) 3 Weeks/1 Week
Manchester Attack	0.420	0.488	0.661
	(0.013)	(0.033)	(0.135)
Observations	1,456	2,184	2,912

Note: Treatment takes the value of 1 on the day of the Event and the subsequent 6 days. Day-of-the-Week Dummy, Sex, Age, Number of Followers, Number of Friends, Number of Statuses, and Age of the Account included. Errors are clustered at the MP level. IRR and p-values are reported. Sample restricted to MPs with a Majority Share in 2015 Elections of less than 10 percentage points. Only MPs standing for 2017 Elections are considered.

Table 8 – One-sided Z-Test $H_0 : \beta_{Subsample} \leq \beta_{FullSample}$

	1 Week/1 Week	2 Weeks/1 Week	3 Weeks/1 Week
Z Statistic	-1.20	-0.84	-0.71
P-Value	0.88	0.80	0.76

Note: One-sided z test. Null hypothesis: Coefficient of the Treatment effect on Immigration-related Tweets in subsample equal or less than the coefficient of the Treatment effect on the full sample. Manchester attack, all the available MPs are used.

We now analyse a different channel through which our Treatment might have a heterogeneous impact: the presence of a relevant stock of migrant people in the constituency. In order to explore this hy-

Table 9 – Heterogeneity by Betting Odds

	(1) 1 Week/1 Week	(2) 2 Weeks/1 Week	(3) 3 Weeks/1 Week
Manchester Attack	0.537	0.555	0.668
	(0.007)	(0.011)	(0.088)
Observations	2,394	3,591	4,787

Note: Treatment takes the value of 1 on the day of the Event and the subsequent 6 days. Day-of-the-Week Dummy, Sex, Age, Number of Followers, Number of Friends, Number of Statuses, and Age of the Account included. Errors are clustered at the MP level. IRR and p-values are reported. Sample restricted to MPs with Betting Odds greater than 0.1. Only MPs standing for 2017 Elections are considered.

Table 10 – One-sided Z-Test $H_0 : \beta_{Subsample} \leq \beta_{FullSample}$

	1 Week/1 Week	2 Weeks/1 Week	3 Weeks/1 Week
Z Statistic	-0.73	-0.69	-0.86
P-Value	0.77	0.76	0.81

Note: One-sided z test. Null hypothesis: Coefficient of the Treatment effect on Immigration-related Tweets in subsample equal or less than the coefficient of the Treatment effect on the full sample. Manchester attack, all the available MPs are used.

pothesis, we will add an interaction of our Treatment with the variable of interest, keeping the latter as a covariate to account for differences in levels. According to Intergroup Contact Theory, increased intergroup relations reduce the conventional image of outgroups and enhance attitudes toward them (Legewie (2013)). Hence, we should expect that, if the share of migrant population in a constituency is relatively small (i.e. the share of native people is large), voters might be more worried about the issue of immigration and the politician could exploit the event to signal her position on the political spectrum. Thus, we would expect those MPs to be more prone to expose themselves in the aftermath of the incident. Our proxies for intergroup contact are the share of UK-born people and the share of people of white British ethnicity at the constituency level. However, these two measures present significant drawbacks. First, they come from the 2011 Census, so they do not reflect the constituency's condition at the time of the event. Second, there might be concerns on how well these variables represent the same concept. In fact, the share of UK-born people also includes second-generation migrants, who might still be considered outsiders from the other natives, so it does not discriminate between different ethnicities. Moreover, as these variables come from the British Election Study 2017, we do

Table 11 – Heterogeneity by Share of UK-Born People

	(1) 1 Week/1 Week	(2) 2 Weeks/1 Week	(3) 3 Weeks/1 Week
Treatment	0.130	0.135	0.194
	(0.021)	(0.025)	(0.063)
Interaction	1.019	1.019	1.017
	(0.079)	(0.074)	(0.098)
UKBorn Share	0.984	0.985	0.986
	(0.036)	(0.011)	(0.026)
Observations	7,448	11,172	14,895

Note: Treatment takes the value of 1 on the day of the Event and the subsequent 6 days. Day-of-the-Week Dummy, Sex, Age, Number of Followers, Number of Friends, Number of Statuses, and Age of the Account included. Errors are clustered at the MP level. IRR and p-values are reported. All the available MPs are used.

Table 12 – Wald Test for Joint Significance (UK-Born)

	1 Week/1 Week	2 Weeks/1 Week	3 Weeks/1 Week
Wald Statistic	5.56	7.68	6.41
P-Value	0.06	0.02	0.04

Note: Wald test for the joint significance of the Interaction and UKBorn Share.

not have information on the 16 constituencies of Northern Ireland present in our dataset. Table 11 and 13 present the results of the estimations, we report also a Wald test for the joint significance of the interaction term and the variable considered (Table 12 and 14). The terms are jointly statistically significant in every specification, and in all periods. The magnitudes do not differ substantially and the cumulated effects have positive sign. Yet, the estimated impact is quite small. For instance, considering Table 11 and the closest time interval, if we raise the share of UK-born individuals in a constituency by 20 percentage points, the elected politician is predicted to increase its expected number of immigration-related Tweets by only 4.5% after the attack, compared to the others.[23]

The next assumption that we would like to test is related again to the contextual factors that might shape politicians' behaviour. Material concerns and perceived group deprivation could increase intergroup hostility. Adverse economic conditions might reduce collective resources and enhance out-group

[23] The cumulated impact is computed as $(1.018526 \times .9839869)^{20}$, as the effect is multiplicative. The comparison group is represented by politicians affected by the Treatment (i.e. after the attack), but belonging to constituencies with a share of UK-born people lower by 20 percentage points.

Table 13 – Heterogeneity by Share of White British People

	(1) 1 Week/1 Week	(2) 2 Weeks/1 Week	(3) 3 Weeks/1 Week
Treatment	0.222	0.229	0.306
	(0.004)	(0.006)	(0.027)
Interaction	1.013	1.013	1.013
	(0.050)	(0.045)	(0.056)
WhiteBritish Share	0.991	0.992	0.993
	(0.073)	(0.038)	(0.068)
Observations	7,448	11,172	14,895

Note: Treatment takes the value of 1 on the day of the Event and the subsequent 6 days. Day-of-the-Week Dummy, Sex, Age, Number of Followers, Number of Friends, Number of Statuses, and Age of the Account included. Errors are clustered at the MP level. IRR and p-values are reported. All the available MPs are used.

Table 14 – Wald Test for Joint Significance (White British)

	1 Week/1 Week	2 Weeks/1 Week	3 Weeks/1 Week
Wald Statistic	5.31	6.66	5.90
P-Value	0.07	0.04	0.05

Note: Wald test for the joint significance of the Interaction and White British Share.

threat, as attitudes are likely to be shaped by the perceived impact of the outsiders at the community rather than at the individual level (Hainmuller and Hopkins (2014)). Thus, we could expect that in constituencies facing downturns a politician would be more prone to exploit a terrorist event to highlight the issue of immigration and shift public attention toward this topic, using outgroups as scapegoats for the recession. Hence, we would anticipate a relatively higher number of immigration-related Tweets after the attack for those politicians elected in constituencies facing worse economic conditions. Our proxy for competition over scarce resources is the average unemployment level in 2016, measured as the share of economically active residents aged between 16-64 years-old claiming Jobseeker's Allowance. Results are shown in Table 15. The estimated effect is quite large in magnitude, but it is only jointly statistically significant in the closest interval (see Table 16). For instance, a politician elected in a constituency with an unemployment rate 2 percentage points higher, is predicted to increase the expect number of immigration-related Tweets by around 28% after the incident, compared to the others.[24] Thus, it seems that we do observe a different impact, but just when taking into account the interval in the immediate proximity of the event.

Table 15 – Heterogeneity by Level of Unemployment

	(1) 1 Week/1 Week	(2) 2 Weeks/1 Week	(3) 3 Weeks/1 Week
Treatment	0.277	0.378	0.499
	(0.000)	(0.005)	(0.034)
Interaction	1.325	1.193	1.170
	(0.007)	(0.122)	(0.138)
Average Unemployment	0.854	0.930	0.950
	(0.040)	(0.363)	(0.437)
Observations	7,672	11,508	15,343

Note: Treatment takes the value of 1 on the day of the Event and the subsequent 6 days. Day-of-the-Week Dummy, Sex, Age, Number of Followers, Number of Friends, Number of Statuses, and Age of the Account included. Errors are clustered at the MP level. IRR and p-values are reported. All the available MPs are used.

Another important channel of heterogenous effect is that of incumbency. In particular, it appears of interest to understand if politicians belonging to the party of the incumbent government act differently from the main challengers. If the "muting effect" is due to politicians strategically being cautious

[24] The cumulated effect is computed as $(1.325405 \times .854322)^2$. The comparison group is represented by politicians affected by the Treatment (i.e. after the attack), but belonging to constituencies with an unemployment rate lower by 2 percentage points.

Table 16 – Wald Test for Joint Significance (Unemployment)

	1 Week/1 Week	2 Weeks/1 Week	3 Weeks/1 Week
Wald Statistic	7.83	2.40	2.20
P-Value	0.02	0.30	0.33

Note: Wald test for the joint significance of the Interaction and Average Unemployment.

and avoiding a risky topic, we should expect a greater reduction for the members of the incumbent government's party compared to the challenging one, as they might be deemed responsible for the current immigration policy. We thus select a subsample of the MPs: those only belonging to either the Conservative (the incumbent) or the Labour party (the main challenger). We re-estimate our model by adding a Dummy variable for belonging to the Tories and an interaction with our Treatment. Results are presented in Table 17. The effects are rather imprecisely estimated, and the two terms are jointly statistically significant (at the conventional levels) only when considering the largest time period (see Table 18). However, the sign of the cumulated effect is as expected, but the magnitude is rather small. If we look at the broadest interval, after the terrorist attack a Conservative MP is predicted to reduce its expected number of immigration-related Tweets by a further 5% compared to a Labour one.[25]

Table 17 – Heterogeneity by Incumbency

	(1) 1 Week/1 Week	(2) 2 Weeks/1 Week	(3) 3 Weeks/1 Week
Treatment	0.565	0.553	0.654
	(0.018)	(0.009)	(0.047)
Interaction	1.695	1.566	1.624
	(0.242)	(0.233)	(0.181)
Conservative	0.509	0.617	0.588
	(0.037)	(0.045)	(0.014)
Observations	6,160	9,240	12,319

Note: Treatment takes the value of 1 on the day of the Event and the subsequent 6 days. Day-of-the-Week Dummy, Sex, Age, Number of Followers, Number of Friends, Number of Statuses, and Age of the Account included. Errors are clustered at the MP level. IRR and p-values are reported. Only observations for Conservative and Labour are used.

[25] The cumulated effect is computed as ($1.623777 \times .5876088$). The impact is larger when considering the closest interval to the event (a reduction of 14%), but it is only jointly significant at 11%. The comparison group is represented by Labour MPs affected by the Treatment (i.e. after the attack).

Table 18 – Wald Test for Joint Significance (Incumbency)

	1 Week/1 Week	2 Weeks/1 Week	3 Weeks/1 Week
Wald Statistic	4.38	4.20	6.52
P-Value	0.11	0.12	0.04

Note: Wald test for the joint significance
of the Interaction and Conservative Dummy.

One last factor that might mediate politicians' strategic behaviour is the salience of the issue among their voters. However, the sign of the resulting effect is not straightforward. On one side, we could think that hostility towards immigrants in the electorate can be hazardously exploited by a politician to signal her ideology or valence after the incident. On the other side, we might expect that, if voters are particularly sensitive to immigration and the politician adopts a risk-averse stance, she would avoid dealing with that issue in the aftermath of the attack, given the emotional impact that such events have on the public. Our proxy for the salience of the topic is the share of votes for Leave in the 2016 EU Referendum. We assume that a higher proportion of Leave is suggestive of restrictionism in immigration policy, and so greater concerns about free movement of people. Unfortunately, the results of the Referendum are not available at the constituency level. Hence, we use the estimates computed by Hanretty (2017). As the data come from the British Election Study, we lose again information on the 16 constituencies of Northern Ireland present in our dataset. Results are shown in Table 19. The share of votes for Leave and its interaction with the Treatment are jointly statistically significant in all time periods (see Table 20) and their cumulated effect is rather substantial. Looking at the closest interval, if we increase in a constituency the share of vote for Leave by 10 percentage points, the elected politician is predicted to reduce the expected number of immigration-related Tweets by an additional 23% after the attack, compared to the others.[26]

This last result motivates a closer look to the strategic behaviour of MPs belonging to those constituencies where the Leave vote scored high in the EU Referendum. Figure 5 displays the cumulated number of immigration-related Tweets during the election campaign for the twenty constituencies with the highest Leave share. From the chart it does not appear that politicians belonging to those areas were particularly keen on approaching the topic: 13 out of 20 did not touch upon the immigration

[26] The cumulated effect is computed as $(1.030285 \times .9459803)^{10}$. The comparison group is represented by politicians affected by the Treatment (i.e. after the attack), but belonging to constituencies with a Leave share lower by 10 percentage points.

Table 19 – Heterogeneity by EU Referendum Results

	(1) 1 Week/1 Week	(2) 2 Weeks/1 Week	(3) 3 Weeks/1 Week
Treatment	0.173	0.324	0.485
	(0.002)	(0.025)	(0.148)
Interaction	1.030	1.017	1.012
	(0.019)	(0.135)	(0.261)
Leave Share	0.946	0.959	0.963
	(0.000)	(0.000)	(0.000)
Observations	7,448	11,172	14,895

Note: Treatment takes the value of 1 on the day of the Event and the subsequent 6 days. Day-of-the-Week Dummy, Sex, Age, Number of Followers, Number of Friends, Number of Statuses, and Age of the Account included. Errors are clustered at the MP level. IRR and p-values are reported. All the available MPs are used.

Table 20 – Wald Test for Joint Significance (EU Referendum)

	1 Week/1 Week	2 Weeks/1 Week	3 Weeks/1 Week
Wald Statistic	34.28	25.90	23.19
P-Value	0.00	0.00	0.00

Note: Wald test for the joint significance of the Interaction and Leave Share.

issue at all during the elections. In fact, when considering all constituencies, the Pearson's correlation coefficient between the two variables is negative, even if not dramatically large (-0.2798). Thus, the pattern that seems to emerge is an inverse association: the greater the demand for restrictionism, the less the MP covers the issue of immigration in her rhetoric. However these findings are just suggestive, as MPs belonging to "Leave constituencies" might also tweet systematically less. We address this issue in Table 21, which shows a regression of the daily number of immigration-related Tweets posted by a MP during the elections (19/04/2017-07/06/2017) on the Leave share in her constituency. We estimate the model through a ZINB with the inflation factor π given by a logistic function of the daily number of Tweets, thus keeping into account the everyday use of the microblog by the politician. We cluster the errors at the MP level and use the same covariates as in our baseline. The result suggests again a negative relationship, and the coefficient is highly significant (<1%). Taken together, these findings provide further evidence for politicians adopting a risk-averse attitude on immigration when their electorate is more sensitive to the issue.

Figure 5 – Leave Share and Immigration-Related Tweets During Elections

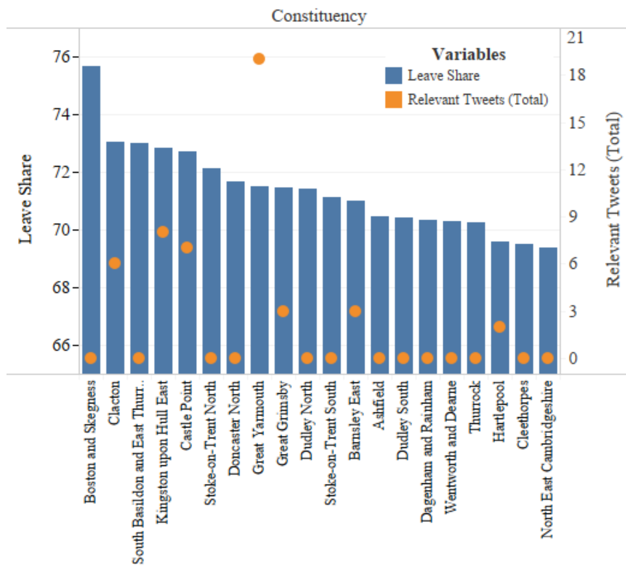

Table 21 – Leave Share and Immigration-Related Tweets During Elections

	(1) Election Campaign Immigration-Related Tweets
Leave Share	0.975
	(0.001)
Observations	23,069

Note: Day-of-the-Week Dummy, Sex, Age, Number of Followers, Number of Friends, Number of Statuses, and Age of the Account included. Errors are clustered at the MP level.
IRR and p-value are reported. Only MPs for whom we have complete information are used.

5.3 Robustness Checks

We now present a series of Robustness Checks for our baseline results on the Manchester attack, when we consider all the available MPs for whom we have complete information over the time periods considered. In the first ones, we perform different estimation strategies. We estimate our model with a standard Negative Binomial adding the total number of Tweets posted by the politician in a given day as a covariate. Results are reported in Table 22. The effect is less precisely estimated, but the magnitudes do not change substantially: the difference in the IRRs remains between 2 and 4 percentage points from our baseline.

Table 22 – Robustness Check 1: NB2 with Total Tweets as Covariate

	(1) 1 Week/1 Week	(2) 2 Weeks/1 Week	(3) 3 Weeks/1 Week
Manchester Attack (All Obs.)	0.646	0.671	0.820
	(0.003)	(0.008)	(0.173)
Observations	7,672	11,508	15,343

Note: Treatment takes the value of 1 on the day of the Event and the subsequent 6 days. Day-of-the-Week Dummy, Sex, Age, Total Tweets, Number of Followers, Number of Friends, Number of Statuses, and Age of the Account included. Errors are clustered at the MP level. IRR and p-values are reported. All the available MPs are used.

Next we exploit the panel structure of our dataset to take into account unobserved individual heterogeneity. We estimate our baseline with a Poisson Random Effects (RE) model (Cameron et al. (2013)). This model is less demanding in terms of distributional assumptions than a NB RE, but it is more efficient than a Pooled Poisson when overdispersion is of the NB2 form (as we have assumed in our baseline). The choice of RE is justified by the fact that, given the exogeneity in the timing of the event, it is unlikely for the time-constant individual effect to be correlated with our Treatment variable. We use the same covariates as in the baseline, adding the daily total number of Tweets posted by a MP as in the previous Robustness Check. Results are reported in Table 23. The IRRs are very close to our baseline regression and the effect is even more precisely estimated for the largest time interval.

Table 23 – Robustness Check 2: Poisson RE with Total Tweets as Covariate

	(1) 1 Week/1 Week	(2) 2 Weeks/1 Week	(3) 3 Weeks/1 Week
Manchester Attack (All Obs.)	0.615	0.632	0.777
	(0.001)	(0.001)	(0.066)
Observations	7,672	11,508	15,343

Note: Treatment takes the value of 1 on the day of the Event and the subsequent 6 days. Day-of-the-Week Dummy, Sex, Age, Total Tweets, Number of Followers, Number of Friends, Number of Statuses, and Age of the Account included. Errors are clustered at the MP level. IRR and p-values are reported. All the available MPs are used.

The second exercise that we are going to perform is to exclude all MPs elected in London's constituencies, to control that our results are not driven by what is happening in the capital.[27] Results are displayed in Table 24. Again, our main conclusions are unaffected by this test: we still observe a substantial and significant decrease in the closest time intervals.

Table 24 – Robustness Check 3: London's Constituencies Excluded

	(1) 1 Week/1 Week	(2) 2 Weeks/1 Week	(3) 3 Weeks/1 Week
Manchester Attack (All Obs.)	0.644	0.669	0.832
	(0.008)	(0.011)	(0.225)
Observations	6,776	10,164	13,551

Note: Treatment takes the value of 1 on the day of the Event and the subsequent 6 days. Day-of-the-Week Dummy, Sex, Age, Number of Followers, Number of Friends, Number of Statuses, and Age of the Account included. Errors are clustered at the MP level. IRR and p-values are reported. All the available MPs are used. 65 constituencies of London (out of 73) are excluded from the analysis.

Subsequently, we slightly change the nature of our dependent variable in Table 25: we construct a Dummy at the MP level for posting an immigration-related Tweet in a given day. Hence, now we are looking at the likelihood of tweeting about immigration in the days following the attack. The results are consistent with our previous findings:[28] focusing on the narrowest time interval, the probability of writing an immigration-related Tweet was almost 39% less during the week after the incident, compared to the week before.

Table 25 – Robustness Check 4: Probability of Immigration-Related Tweets

	(1) 1 Week/1 Week	(2) 2 Weeks/1 Week	(3) 3 Weeks/1 Week
Manchester Attack (All Obs.)	0.618	0.668	0.812
	(0.003)	(0.007)	(0.152)
Observations	7,672	11,508	15,343

Note: Treatment takes the value of 1 on the day of the Event and the subsequent 6 days. Day-of-the-Week Dummy, Sex, Age, Total Tweets, Number of Followers, Number of Friends, Number of Statuses, and Age of the Account included. Errors are clustered at the MP level. OR and p-values are reported. All the available MPs are used.

Another interesting question that we might want to ask is if the effect captured for the Manchester

[27] Out of the 73 constituencies of London, 65 are present in our dataset.

[28] The model is estimated through a Logit using the same covariates as in the baseline estimations, but keeping also the total number of Tweets by day per MP as independent variable. Errors are clustered at the MP level. Odds ratios and p-values are reported.

attack is constrained to the MPs that were standing for the elections or it is instead a more generalized result that applies to all politicians in charge. In fact, it should be pointed out that these politicians, even if facing different incentives during the election campaign, might adopt a strategic response as well, since their behaviour is likely to influence the odds of the candidate of the same party standing for their constituency. In Table 26 we report our baseline estimations excluding the 23 MPs not standing in the 2017 elections that are present in our dataset. If we compare it with the last section of Table 5 we notice that the IRRs are not affected by this exercise, the difference is less than 1 percentage point in every time interval. Hence, the observed behaviour seems to hold across all MPs. However, it should be highlighted that the politicians excluded are only a small proportion of our sample, so their reaction to the terrorist attack should be substantially different to radically change the size and sign of the average effect estimated in our baseline.

Table 26 – Robustness Check 5: MPs Not Standing in 2017 Elections Excluded

	(1) 1 Week/1 Week	(2) 2 Weeks/1 Week	(3) 3 Weeks/1 Week
Manchester Attack (All Obs.)	0.616	0.644	0.804
	(0.002)	(0.003)	(0.124)
Observations	7,350	11,025	14,699

Note: Treatment takes the value of 1 on the day of the Event and the subsequent 6 days. Day-of-the-Week Dummy, Sex, Age, Number of Followers, Number of Friends, Number of Statuses, and Age of the Account included. Errors are clustered at the MP level. IRR and p-values are reported. All the available MPs are used. Politicians not standing in 2017 elections are excluded from the analysis (23 in our dataset).

Finally, we address the suspension of the election campaign that occurred in the aftermath of the attack. After the Manchester bombing, to pay tribute to the victims, the leaders of all major parties agreed on suspending the campaigning activity, which was subsequently resumed at the local level after two days.[29] Hence, we might wonder if the "muting effect" we are capturing it is not just the result of this political freeze. In order to explore this hypothesis we exclude the two days following the attack from our sample and re-estimate our model. However, this exercise implies a substantial reduction in our treatment group. In order to overcome the loss of efficiency given by a reduced sample size we narrow down our analysis to the two closest time intervals. This allows us to add 4 additional MPs

[29] Walker, P. and Phipps, C. (23 May 2017). "General election campaigning suspended after Manchester attack". *The Guardian*. Retrieved 13 March 2018.
(24 May 2017). "Political parties resume general election campaign". *BBC News*. Retrieved 13 March 2018.

for whom we have complete information over these periods. Table 27 reports the results for this last Robustness Check. The effect is less precisely estimated, but it is significant at conventional levels. We still capture a substantial proportional decrease in the expected number of immigration-related Tweets, around 30% when we compare the week before to the week following the attack. Thus, it seems that our "muting effect" lasted even after this major campaigning shock.

Table 27 – Robustness Check 6: Days of Suspended Campaigning Activity Excluded

	(1) 1 Week/1 Week	(2) 2 Weeks/1 Week
Manchester Attack (All Obs.)	0.700	0.736
	(0.048)	(0.058)
Observations	6,624	10,488

Note: Treatment takes the value of 1 on the day of the Event and the subsequent 6 days, but 23/05 and 24/05 are excluded. Day-of-the-Week Dummy, Sex, Age, Number of Followers, Number of Friends, Number of Statuses, and Age of the Account included. Errors are clustered at the MP level. IRR and p-values are reported. All the available MPs are used.

6 Discussion and Conclusion

In this study we analyse the consequences of two acts of terrorism occurred in 2017 on the immigration rhetoric of British MPs. We focus on a specific information environment: the social media Twitter. Our goal is to explore a potential channel through which these events might be exploited by political elites and in turn shape public attitudes towards immigrants. Natives' attitudes towards outgroups are crucial for the social integration of minorities and the economic success of the community as a whole. Thus, it appears relevant to look at the role that political leaders might play in this process of perceptions' formation. To answer this question, we scrape politicians' Twitter accounts using Text Analysis and Machine Learning techniques in order to gather all their Tweets related to the issue of immigration. We then frame a Natural Experiment setting exploiting the exogeneity in the timing of the events and the granularity of the data gathered. We find a significant impact during the election campaign, but the direction of the effect is rather counterintuitive. In fact, political leaders might strategically exploit these dramatic episodes to foster the debate about immigration and divert attention from contextual problems. Alternatively, they might seek the opportunity to signal their ideological stance or valence. In both cases, we would expect an increase in the amount of relevant information provided by politicians in the aftermath of a terrorist attack. In contrast, what we observe is a "muting effect": a substantial decrease, on average, in the number of immigration-related Tweets in the week following the incident. Our hypothesis is that, given the high stakes they face during elections and the emotional distress caused by these dreadful events, MPs strategically prefer not to take a stance on a risky topic.

In order to further investigate this hypothesis we construct a measure of attention on a Tweet and we analyse if in the days following the attacks the public was more sensitive about the issue of immigration. Increased attention by their followers on this theme would justify a more risk-averse attitude, due to the open exposure to attacks and criticisms that the microblog implies. We conduct this exercise for the Westminster attack, in order to motivate a learning behaviour by politicians that would help to understand the observed results for the Manchester bombing. Our attention variable is computed as $\ln(Favourites + Retweets + 1)$,[30] for any immigration-related Tweet posted by a British MP in a given day. We include in the regression day-of-the-week Dummies, sex, number of followers, number of friends, number of statuses, and another Dummy for the message shared by the politician being a Retweet itself; as these are all factors that might affect the attention on a Tweet. The impact of the event is estimated through OLS and we compute Robust S.E.. In the exercise we compare the attention on immigration-related Tweets two weeks before and two weeks after the incident, so our Treatment takes the value of 1 on the day of the attack and the subsequent 13 days. Results are reported in Table 28. We can observe that in the aftermath of the event the public was definitely more sensitive about the issue. The attention on immigration-related Tweets posted by MPs increased by approximately 60%,[31] compared to the two weeks before the attack, and the effect is statistically significant at 5%.

These last findings appear to be consistent with the experimental results of Gadarian and Albertson (2014), who, building on Affective Intelligence Theory, show how anxious individuals exhibit increased and biased information seeking. Hence, it appears that the incident caused a greater attention on the issue of immigration and this attentiveness could have been exploited by political elites. However, for those who tried to do so after the Westminster attack, this turned out to be quite a risky strategy.

[30] We add 1 to the argument of the logarithm to account for Tweets that are not favourite or shared. The adjustment however, should not be too problematic, as the percentage of zeroes in this restricted sample is around 2.3% (Wooldridge (2015)). A log function is preferred to squeeze the distribution of Favourites and Retweets.

[31] The effect is computed as $\exp^{0.466} - 1$.

Table 28 – Attention on Immigration-Related Tweets

	(1) OLS 2 Weeks/2 Weeks
Westminster Attack	0.466
	(0.021)
Observations	428

Note: Dependent Variable is ln(Favourites+Retweets+1). Treatment takes the value of 1 on the day of the Event and the subsequent 13 days. Day-of-the-Week Dummy, Sex, Number of Followers, Number of Friends, Number of Statuses, and Dummy for Retweet included. Robust S.E. are computed. P-value reported in parentheses.

Donald Trump Jr., whose racist messages in the microblog already prompted widespread backlash,[32] was publicly denounced after his Tweet criticizing London's Mayor Sadiq Khan in the aftermath of the Westminster attack. Among the critics, Wes Streeting, Labour MP for Ilford North, replied on the social media defining the US President's son "a disgrace", condemning his attempt to exploit the event for his own political gain.[33] Other European politicians, as the Front National leader Marine Le Pen or the Polish PM Beata Szydło, openly linked the attack to migration policy and borders control.[34] This generated prompt reactions from different MPs and inflamed the debate in the Twitter community, especially among those users who blamed failed multiculturalism, as Figure 6 shows. At the same time, Nigel Farage appeared on US television endorsing the hard-line immigration and anti-Muslim policies of President Trump. The former UKIP leader clearly connected the episode with British politics, blaming for the attack Tony Blair's government which encouraged mass immigration and "invited in terrorism".[35] He was then forced to draw back from his initial position and publicly admit no direct link between the event and the issue of immigration, once it was clear that the offender was actually British.[36]

Thus, it seems that taking a stance was a dangerous strategy for both sides of the political spectrum,

[32] Malkin, A. (20 September 2016). "Donald Trump Jr compares Syrian refugees to poisoned Skittles". *The Guardian*. Retrieved 06 February 2018.

[33] Levin, S. (23 March 2017). "Donald Trump Jr called 'a disgrace' for tweet goading London mayor Sadiq Khan". *The Guardian*. Retrieved 06 February 2018.

[34] Henley, J. and Jamieson, A. (23 March 2017). "Anti-immigration politicians link London attack to migrant policy". *The Guardian*. Retrieved 06 February 2018.

[35] Oppenheim, M. (23 March 2017). "Nigel Farage blames multiculturalism for London terror attack". *The Independent*. Retrieved 06 February 2018.

[36] Sharman, J. (26 March 2017). "Nigel Farage forced to admit Westminster attack had nothing to do with immigration". *The Independent*. Retrieved 06 February 2018.

Figure 6 – Example of MP Reaction to Anti-Immigration Rhetoric Following The Westminster Attack

as it exposed the leaders to attacks and criticisms by opponents and the public. Hence, the observed "muting effect" for Manchester might be a consequence of politicians learning to avoid a risky topic when the electorate is more sensitive about the theme.

Digging deeper, we find significant heterogenity in this "muting effect" according to the characteristics of the MPs or their constituencies, but also further evidence for a risk-averse attitude adopted by political elites in the aftermath of the attack.

A possible consequence of this reluctant behaviour is a potential mismatch between voters' preferences and the actual type of politicians. Due to the increased information seeking and sensitivity after the event, the electorate might be more receptive to the few opinion leaders who are willing to expose themselves, irrespective of their quality. In particular, if the political leaders in charge adopt a risk-averse attitude and are less willing to take a stance, voters might become more sensitive to the rhetoric of anti-establishment parties and movements, which are not afraid to expose themselves given their firm position on such issues. However, a potential connection between the observed behaviour of politicians and the actual electoral outcomes is not pursued in this paper and it appears to be an interesting and unanswered question for future research.

A general concern with the analysis might be that our "muting effect" is actually a result of messages with extreme negative polarity being censored by Twitter itself. However, Twitter's hateful conduct policy applies to rather extreme cases, such as "*promote violence against or directly attack or threaten other people*".[37] Hence, it seems very unlikely that the MPs will be so radical in their response to the event in such a critical juncture represented by the election campaign (given also the absence of far-right parties in the Parliament).

An interesting extension of our work would be trying to understand if the risk-averse behaviour we capture is a decision of the single politician or a strategic response coordinated by the parties. The results presented in Section 5.2 provide some suggestive evidence for the first hypothesis, as we detect greater variability at the individual level, with a stronger "muting effect" for MPs sitting on more marginal seats or belonging to constituencies where the issue of immigration is more salient. Instead, the response to the event among the two major parties does not seem to dramatically differ. However, a deeper study of the relationships between MPs on social media and the type of response to the attack among clusters of accounts should be carried out.

In addition, exploring other issues, more directly related to the nature of the attacks, as multiculturalism and Islamophobia, could provide more insights on the strategic reactions of politicians to these dreadful events.

One important limitation of the present study is that it does not explore possible changes in the polarity of the Tweets. However, as explained in Section 3, this is mainly due to our Research Design and the current state of the art in Computational Linguistics. In fact, one might wonder if the drop in the immigration-related Tweets is not capturing a reduction in the amount of messages, but a change in the wording around this theme instead. This is unlikely to be our case, as Twitter, with its 140 characters constraint, does not allow for complex phrasing or involuted circumlocution. Thus, our two-step classification should effectively cover the domain of interest (i.e. immigration). This conclusion will not be the same in a Sentiment Analysis framework. In such setting, the choice of words defining the

[37] For Twitter's hateful conduct policy, please visit https://help.twitter.com/en/rules-and-policies/hateful-conduct-policy.

polarity on the theme of interest is likely to change in response to the event and our classification exercise would not be able to capture this shift. Future work should try to overcome these constraints and capture if and how political elites shape their sentiment towards immigration following such shocks.

7 Bibliography

References

Adi, A., Erickson, K., Lilleker, D. G., 2014. Elite tweets: Analyzing the twitter communication patterns of labour party peers in the house of lords. Policy & Internet 6 (1), 1–27.

Allen, W., Blinder, S., 2013. Migration in the news: Portrayals of immigrants, migrants, asylum seekers and refugees in national british newspapers, 2010 to 2012. Migration Observatory report, COMPAS, University of Oxford.

Athey, S., Imbens, G. W., 2015. Machine learning methods for estimating heterogeneous causal effects. stat 1050 (5).

Bajari, P., Nekipelov, D., Ryan, S. P., Yang, M., 2015. Machine learning methods for demand estimation. The American Economic Review 105 (5), 481–485.

Barbieri, F., Ronzano, F., Saggion, H., 2015. How topic biases your results? a case study of sentiment analysis and irony detection in italian. In: RANLP. pp. 41–47.

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

Bozzoli, C., Müller, C., 2011. Perceptions and attitudes following a terrorist shock: Evidence from the uk. European Journal of Political Economy 27, S89–S106.

Brader, T., Valentino, N. A., Suhay, E., 2008. What triggers public opposition to immigration? anxiety, group cues, and immigration threat. American Journal of Political Science 52 (4), 959–978.

Breiman, L., et al., 2001. Statistical modeling: The two cultures (with comments and a rejoinder by the author). Statistical science 16 (3), 199–231.

Bright, J., Hale, S. A., Ganesh, B., Bulovsky, A., Margetts, H., Howard, P., 2017. Does campaigning on social media make a difference? evidence from candidate use of twitter during the 2015 and 2017 uk elections. arXiv preprint arXiv:1710.07087.

Cameron, A. C., Trivedi, P. K., 2013. Regression analysis of count data. Vol. 53. Cambridge university press.

Cameron, A. C., Trivedi, P. K., et al., 2013. Count panel data. Oxford Handbook of Panel Data, Oxford University Press, Oxford.

Ceobanu, A. M., Escandell, X., 2010. Comparative analyses of public attitudes toward immigrants and immigration using multinational survey data: A review of theories and research. Annual review of sociology 36, 309–328.

Cohu, M., Maisonneuve, C., Testé, B., 2016. The charlie-hebdo effect: Repercussions of the january 2015 terrorist attacks in france on prejudice toward immigrants and north-africans, social dominance orientation, and attachment to the principle of laïcité. International Review of Social Psychology 29 (1).

Coupe, T., 2017. The impact of terrorism on expectations, trust and happiness–the case of the november 13 attacks in paris, france. Applied Economics Letters 24 (15), 1084–1087.

Downs, A., 1972. Up and down with ecology: the issue-attention cycle. The public.

Einav, L., Levin, J., 2014. Economics in the age of big data. Science 346 (6210), 1243089.

Evans, H. K., Cordova, V., Sipole, S., 2014. Twitter style: An analysis of how house candidates used twitter in their 2012 campaigns. PS: Political Science & Politics 47 (2), 454–462.

Gadarian, S. K., Albertson, B., 2014. Anxiety, immigration, and the search for information. Political Psychology 35 (2), 133–164.

Galeotti, A., Goyal, S., 2010. The law of the few. The American Economic Review 100 (4), 1468–1492.

Gottfried, J., Shearer, E., 2016. News Use Across Social Medial Platforms 2016. Pew Research Center.

Graham, T., Broersma, M., Hazelhoff, K., Van'T Haar, G., 2013. Between broadcasting political messages and interacting with voters: The use of twitter during the 2010 uk general election campaign. Information, Communication & Society 16 (5), 692–716.

Hainmuller, J., Hopkins, D., 2014. Public attitudes towards immigrants. Annual Review of Political Science 17, 472–483.

Hanes, E., Machin, S., 2014. Hate crime in the wake of terror attacks: Evidence from 7/7 and 9/11. Journal of Contemporary Criminal Justice 30 (3), 247–267.

Hanretty, C., 2017. Areal interpolation and the uk's referendum on eu membership. Journal of Elections, Public Opinion and Parties, 1–18.

Hilbe, J. M., 2011. Negative binomial regression. Cambridge University Press.

Hu, M.-C., Pavlicova, M., Nunes, E. V., 2011. Zero-inflated and hurdle models of count data with extra zeros: examples from an hiv-risk reduction intervention trial. The American journal of drug and alcohol abuse 37 (5), 367–375.

Huddy, L., Feldman, S., Taber, C., Lahav, G., 2005. Threat, anxiety, and support of antiterrorism policies. American journal of political science 49 (3), 593–608.

Jackson, N., Lilleker, D., 2011. Microblogging, constituency service and impression management: Uk mps and the use of twitter. The Journal of Legislative Studies 17 (1), 86–105.

Jones, E. E., Pittman, T. S., 1982. Toward a general theory of strategic self-presentation. Psychological perspectives on the self 1 (1), 231–262.

Jungherr, A., 2016. Twitter use in election campaigns: A systematic literature review. Journal of information technology & politics 13 (1), 72–91.

Kendall, C., Nannicini, T., Trebbi, F., 2014. How do voters respond to information? evidence from a randomized campaign. The American Economic Review 105 (1), 322–353.

Kober, T., Weir, D. J., 2015. Optimising agile social media analysis. In: WASSA@ EMNLP. pp. 31–40.

Kreiss, D., 2016. Seizing the moment: The presidential campaigns use of twitter during the 2012 electoral cycle. New Media & Society 18 (8), 1473–1490.

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

Lucas, M., Downey, D., 2013. Scaling semi-supervised naive bayes with feature marginals. In: ACL (1). pp. 343–351.

Montalvo, J. G., 2011. Voting after the bombings: A natural experiment on the effect of terrorist attacks on democratic elections. Review of Economics and Statistics 93 (4), 1146–1154.

Parmelee, J. H., 2014. The agenda-building function of political tweets. New Media & Society 16 (3), 434–450.

Peysakhovich, A., Naecker, J., 2017. Using methods from machine learning to evaluate behavioral models of choice under risk and ambiguity. Journal of Economic Behavior & Organization 133, 373–384.

Rijsbergen, C. J. V., 1979. Information Retrieval, 2nd Edition. Butterworth-Heinemann, Newton, MA, USA.

Saif, H., He, Y., Fernandez, M., Alani, H., 2016. Contextual semantics for sentiment analysis of twitter. Information Processing & Management 52 (1), 5–19.

Schüller, S., 2016. The effects of 9/11 on attitudes toward immigration and the moderating role of education. Kyklos 69 (4), 604–632.

Teng, Z., Vo, D.-T., Zhang, Y., 2016. Context-sensitive lexicon features for neural sentiment analysis. In: EMNLP. pp. 1629–1638.

Varian, H. R., 2014. Big data: New tricks for econometrics. The Journal of Economic Perspectives 28 (2), 3–27.

Vasilopoulos, P., 2018. Terrorist events, emotional reactions, and political participation: the 2015 paris attacks. West European Politics 41 (1), 102–127.

von Grünigen, D., Weilenmann, M., Deriu, J., Cieliebak, M., 2017. Potential and limitations of cross-domain sentiment classification. SocialNLP 2017, 17.

Wager, S., Athey, S., 2017. Estimation and inference of heterogeneous treatment effects using random forests. Journal of the American Statistical Association (just-accepted).

Whitaker, R., Lynch, P., 2011. Explaining support for the uk independence party at the 2009 european parliament elections. Journal of Elections, Public Opinion & Parties 21 (3), 359–379.

Wibberley, S., Weir, D., Reffin, J., 2014. Method51 for mining insight from social media datasets. In: Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: System Demonstrations. pp. 115–119.

Williams, M. L., Burnap, P., 2015. Cyberhate on social media in the aftermath of woolwich: A case study in computational criminology and big data. British Journal of Criminology 56 (2), 211–238.

Wooldridge, J. M., 2015. Introductory Econometrics: A Modern Approach. Cengage Learning.

A List of Words for Boolean Search

List of words used to conduct the Boolean Search:

- Migrant
- Asylum Seeker
- Refugee
- Migration
- Influx
- Wave
- Not Native
- Deportation
- Border
- Foreigner
- Exodus
- Free Movement
- Confine
- Expatriate
- Displacement
- Non-native
- Flee
- Frontier

B Additional Descriptive Statistics

Here we present additional descriptive statistics on the MPs included in our analyses and those never considered. Table 29 shows the distribution by political affiliation. We can observe that the large majority (74.5%) of MPs excluded in our study belongs to the Conservative Party. Table 30 presents the comparison of demographic characteristics between the two groups. It appears that the politicians in our sample are younger than the excluded ones, and women are more represented. Finally, Table 31 shows the reasons for the exclusion of some MPs from our analyses. Most of them were not considered as they did not have a Twitter account at the time of the collection (75.5%). Four politicians had a protected account, whereas one MP was using the Commons Leader account. Three accounts were never considered because the limit on the collection from their timeline (i.e. 3,200 Tweets) was reached before the day that represents the upper temporal bound in our analyses (29th of May 2017).

Table 29 – MP Distribution by Political Party

	MPs Included		MPs Excluded	
Party	Freq.	Percent	Freq.	Percent
Conservative	257	46.56	73	74.49
Democratic Unionist Party	7	1.27	1	1.02
Green	1	0.18	-	-
Independent	3	0.54	2	2.04
Labour	184	33.33	19	19.39
Labour Co-operative	26	4.71	1	1.02
Liberal Democrats	9	1.63	-	-
Plaid Cymru	3	0.54	-	-
Scottish National Party	53	9.60	1	1.02
Sinn Féin	4	0.72	-	-
Social Democratic and Labour Party	3	0.54	-	-
Ulster Unionist Party	2	0.36	-	-
Speaker	-	-	1	1.02
Total	552	100.00	98	100.00

Table 30 – Comparison Demographic Characteristics

	MPs Included		MPs Excluded	
	Mean	Std. Dev.	Mean	Std. Dev.
Male	.67	-	.84	-
Age	51.28	10.17	58.08	12.12

Table 31 – Reasons for Exclusion

Reason	Freq.	Percent
Too Active	3	3.06
No Personal Account	1	1.02
Protected	4	4.08
No Account	75	76.53
Not Active	15	15.31
Total	98	100.00

C Examples of Tweets

Figure 7 – Examples of Irrelevant Tweets Picked Up in The Boolean Search

Figure 8 – Examples of False Positives for the Irrelevant Category

Figure 9 – Examples of False Positives for the Relevant Category

Figure 10 – Examples of Tweets with Different Polarities

(a) Example of Tweet with Positive Stance

(b) Example of Tweet with Negative Stance

(c) Example of Tweet with Neutral Stance

(d) Example of Tweet with Ambiguous Stance