Terrorism, Media, and Attitudes Toward Immigrants: A Cross-National Analysis of the Charlie Hebdo Attack

Mamadou Sacko*[1]

[1]*Laboratoire d'Economie d'Orléans, Université d'Orléans*

November 27, 2024

Abstract

This paper presents a comprehensive cross-national analysis of the impact of the January 2015 Charlie Hebdo shooting in Paris on public attitudes towards immigrants and Muslims. Utilizing a robust array of methodologies, including fixed-effects models, various matching methods, and advanced machine learning techniques, we examine how attitudes shifted in different European nations post-Charlie Hebdo, revealing a nuanced picture of regional variations in response to the attack. Additionally, we investigate the role of media in shaping these attitudes, employing two natural language processing (NLP) methods: Large Language Models (LLM) and Structural Topic Modeling (STM), to analyze news articles from five European countries and understand how immigrants and Muslims were portrayed following the attack. Our findings indicate a general decline in positive attitudes towards immigrants and Muslims in Europe, with distinct patterns emerging in specific countries that correlate with media coverage. This paper not only sheds light on the societal effects of terrorism but also highlights the critical role of media narratives in influencing public opinion. The results emphasize the need for a sensitive approach to reportage in the wake of such incidents.

Keywords: Terrorism, Immigration, Media, France

JEL Classification: F51, F22, D72

*E-mail: sacko.mamadou@univ-orleans.fr. University of Orléans and Laboratoire d'Economie d'Orléans (LEO)

1 Introduction

On January 7, 2015, a profound event shook Paris. Said and Cherif Kouachi, French citizens of Algerian descent, launched a deadly attack on the office of Charlie Hebdo, a well-known satirical newspaper. This assault, rooted in the offenders' claimed retaliation for the paper's portrayal of Prophet Mohammed, not only claimed the lives of eleven employees and a police officer but also sparked a series of related attacks over three days. This culminated in a hostage situation at a kosher supermarket by Amedy Coulibaly, another French citizen, this time of Malian origin. The attacks were followed by large-scale demonstrations and received extensive coverage in both the international and French press, with narratives linking these brutal acts to radical Islam.

In recent years, several other deadly terrorist attacks targeting civilians have been carried out in major European cities – Madrid in 2004, London in 2005, Paris in 2015, Brussels, Nice, and Berlin in 2016, London and Manchester in 2017 – committed by individuals of migrant backgrounds claiming a connection with the Islamic State of Iraq and Syria (ISIS). Previous research has examined how these terrorist attacks influence a variety of societal outcomes, including attitudes towards immigrants and Muslims. For example, anti-Arab and anti-Semitic prejudices intensified after the 2004 Madrid attack Echebarria-Echabe and Fernández-Guede (2006). Ethnic segregation also increased between Arab immigrants and native Spaniards shortly after the attack (Sandell et al., 2016), as did anti-immigrant attitudes in general after the assassination of Theo van Gogh in the Netherlands Boomgaarden and De Vreese (2007); Das et al. (2009). (Rabby and Rodgers iii, 2010) documented that the London bombings caused a decrease in the employment, real earnings, and hours worked of very young Muslim men. (Van de Vyver et al., 2016) also reported stronger prejudices toward Muslims and immigrants after the London bombings. Both (De Coninck, 2022) and (Andersen and Mayerl, 2018) show that positive attitudes toward refugees are associated with lower terrorism fear in Belgium. Attitudes towards immigrants also became more negative in some European countries after the Bali attacks in 2002 Legewie (2013) and the 2015 Paris (Charlie Hebdo) attack Castanho Silva (2018); Savelkoul et al. (2022).

While the consistent finding in this literature is that terrorist attacks, particularly those committed by 'out-group' members such as Islamic terror groups, increase negative attitudes towards Muslims in particular and by extension towards all immigrants, there are also suggestive findings pointing in the opposite direction. For example, Castanho Silva (2018) illustrates the (average) non-impact of the Paris 2015 (Charlie Hebdo) and November 2015 events on a wide range of out-group-related measures, such as anti-immigrant and anti-Muslim attitudes. Van Assche and Dierckx (2021) also found no impact of the November 2015 and 2016 Brussels terrorist attacks in France and Belgium respectively on attitudes towards Muslims, refugees, and immigrants.

A growing number of these studies have taken advantage of the coincidental overlap between the terrorist attack and survey fieldwork to estimate causal effects. This analytical approach, known as unexpected event during survey design (UESD) Muñoz et al. (2020), allows researchers to use the timing of the interview to assign respondents to either a treatment group (i.e., interviewed after the attack) or a control group (i.e., interviewed before the attack). The unexpected nature of the event means that assignment to treatment and control is exogenous to respondent characteristics, conditional on excludability and ignorability assumptions.

This paper adopts the same approach to test whether the Charlie Hebdo shooting caused Europeans to have more negative attitudes towards Muslims in particular and by extension towards all immigrants. Although the attack happened in France, we can expect other European countries to react negatively as well, since citizens of other European countries will likely consider France as part of the same broad "in-group" when it comes to terrorism. We focus on the Charlie Hebdo shooting because the attack was a direct consequence of the journal's vocal critics of Islam. The attack also happened in the middle of the fieldwork of the European Social Survey (ESS) in

France and many other European countries, offering a unique opportunity to assess the impact of the event in 12 European nations with respondents interviewed before and after the event. While the Charlie Hebdo attack has been the focal event of two studies on attitudes toward Muslims Castanho Silva (2018) and Savelkoul et al. (2022), our study is an important addition to this literature as there is still no consensus on the average impact of the attack on attitudes towards Muslims. By re-examining the impact of the Charlie Hebdo shooting on attitudes towards immigrants and Muslims using a vast number of methodologies, we hope that the study finally answers the question of whether the Charlie Hebdo shooting caused a shift in public opinion about Muslim immigrants.

Moreover, and most importantly, while the studies mentioned above are undeniably insightful, they have a different focus as none of them have explored the role of the media in shaping attitudes towards immigrants and Muslims. Indeed, Charlie Hebdo led to the global Je suis Charlie movement, which strengthened the media interest in the event with an outstanding amount of news on the issue (see figure 10). In this study, we explore the idea of the media environment resulting from the incident acting as a mediator that either exacerbates or reduces negative perceptions towards immigrants and Muslims. Only a few studies have considered the impact of certain types of news coverage on immigration attitudes, and most evidence that exists is not specific to terrorism. For example, (Sniderman et al., 2004) have shown that by simply emphasizing the ethnicity of news subjects (i.e., by making it visible), news media can increase out-group hostility in the native media audiences. (Van Klingeren et al., 2015) indicate that media visibility of immigration increases public anti-immigration attitudes, even when controlling for real-world developments (e.g., immigrant inflows) or media tone. Furthermore, many pieces of evidence indicate that the more news media report on immigration (i.e., in quantity alone), the more people tend to vote for parties with an anti-immigrant stance (see (Vliegenthart and Boomgaarden, 2007) and (Burscher et al., 2015)) and that media bias in coverage of immigrant criminality can significantly affect political outcomes and support for far-right parties (i.e. (Couttenier et al., 2021); (Djourelova, 2020) or that lack of bias can significantly reduce natives' concerns about immigration (Keita et al., 2021). The only piece of evidence specific to terrorism was conducted in a controlled laboratory setting and found that undifferentiated news about Islamic State terrorism increased participants' fear of terrorism and resulted in hostile perceptions toward Muslims in general (Von Sikorski et al., 2020). Despite this growing body of literature, a direct link between media coverage of specific terrorist acts, like the Charlie Hebdo shooting, and a decline in attitudes toward outgroups covered in the news remains under-explored. Our study that focuses specifically on the coverage of immigrants and Muslims in the Charlie Hebdo news articles is therefore an important addition to this literature.

We extend the existing literature on the effect of terror on attitudes toward outgroups in two ways. First, leveraging data from the European Social Survey's seventh wave and employing the latest causal inference methods, we offer a multifaceted analysis of the shooting's effect on attitudes toward both immigrants and Muslims. Our approach involves an array of methodologies, starting with OLS models with country fixed effects, progressing through various matching methods like greedy and optimal matching, as well as several weighing methods like Inverse Probability Weighting and Entropy balancing, and culminating in advanced machine learning techniques like Augmented Inverse Propensity Weighting. Our findings indicate a general decline in positive attitudes towards immigrants and Muslims post-Charlie Hebdo, with notable variances across countries like France, Germany, Ireland, Czech Republic, and Finland. This conclusion remains robust across different methodological applications (RDD and Multilevel Linear Regression with Random Intercept) and time-window adjustments (we varied the time window from 10 days to 2 months with 10 days increment). Pre-existing time trends are also ruled out through falsification tests.

Second, in addition to data on attitudes collected immediately before and after the incident, we also analyze news articles published online in 5 European nations on the issue in the month following the incident. These news data enable us to assess the role of the media in shaping public opinion against or in favor of immigrants. Utilizing two different natural language processing methods, a large language model and structural topic modeling, we analyze these news articles, uncovering a correlation between media sentiment and public opinion. Notably, Germany and France emerge as countries with predominantly positive media portrayal of immigrants and Muslims consistent with the improvement in public attitudes towards these groups observed in both countries. This contrasts with Ireland, Czech Republic, and Finland where both attitudes towards immigrants and Muslims and media sentiments towards these groups are more negative. To the best of our knowledge, this is one of the few studies exploring how the framing of terror incidents around immigration and cultural out-groups, like Muslims, can specifically play a role in shaping attitudes towards these groups.

2 Data and Methods

2.1 Data

To test the impact of the Charlie Hebdo terrorist attack on attitudes towards immigrants and Muslims, we utilized data from the seventh round of the European Social Survey (ESS). The ESS is a cross-national survey that reviews attitudes, beliefs, and behavior patterns in European countries every two years based on probability samples gathered through face-to-face interviews. The seventh round of the ESS includes several questions that allow measurement of attitudes towards immigrants and Muslims, and it includes a large number of respondents interviewed immediately before the attacks took place. This makes it possible to assess whether the event shifted attitudes towards immigration from a causal angle by adopting the UESD framework. The terror attacks can be seen as a natural experiment where the attacks are an exogenous shock to attitudes toward immigration and immigrants.

For our analyses, we rely on ±30 days' bandwidth in the main specification because almost all media coverage on the issue was published within one month after the event. Beyond that period, there was almost no press coverage on the issue (see figure 10 in appendices). Nonetheless, in the robustness tests, we discuss results for smaller and larger bandwidths. We kept only respondents from countries with participants interviewed both before and after the attack, resulting in a sample size of 6,041 respondents from 12 European countries, with 2,886 respondents interviewed before the attack and 3,155 after the attack.

The identification of valid causal estimates based on the comparison of respondents interviewed before and after the day of the event hinges on several identifying assumptions. The first is excludability: any difference between respondents interviewed before and after the event shall be the sole consequence of the event. That is, the timing of the interview should affect the outcome variable only through the event. For this assumption to hold, it is important that no other, unrelated events take place around the date of the attack that could affect the outcome. In our case, we have a great advantage with the Charlie Hebdo shooting because no major terrorist attack happened in the months leading to the attack in Europe. We adopt a short time span so that the subjects had very limited time to change some of their attitudinal characteristics due to any other exogenous event than the attack. We also followed the literature Muñoz et al. (2020) and applied the following robustness checks: examination of pre-existing time trends through falsification tests (placebo specifications for one and two months before the actual date) as well as with the RDD that includes a running variable. We also assessed the effect of the shooting on other outcomes for which there should be no effect (attitude toward Jewish and Gypsy populations).

The second key assumption is temporal ignorability: this assumption requires that the timing of the interviews during the fieldwork period occurred by chance. In other words, every individual must have the same, a priori, probability of being interviewed before or after the exogenous treatment. In this way, the treatment assignment is randomized. However, the timing of the interview in the ESS is not under the researchers' control and is not random. The fieldwork often follows a geographical pattern for efficiency reasons. Any correlation between subject location and time of the interview will lead to a violation of the ignorability assumption since the location can be correlated with many other respondents' characteristics that might bias the findings. This potential source of bias can be resolved by controlling for covariates that have been shown to influence how the fieldwork was organized.

2.1.1 Outcomes, Treatment and Control variables

The primary objective of this study is to assess the impact of the Charlie Hebdo attack on public attitudes towards immigrants and Muslims. To this end, we have selected two pertinent variables from the European Social Survey (ESS). The first dependent variable, measuring attitudes towards immigrants, is derived from a survey question asking respondents their views on allowing immigrants from poorer countries outside Europe into their country. The response options range from:
"*Allow none to come (1)*", "*Allow a few (2)*", "*Allow some (3)*" or "*Allow many (4)*".[1]
Our second dependent variable mirrors the first, albeit focusing on attitudes towards Muslims. Respondents were asked to express their views on the extent to which they believe their country should allow Muslims to immigrate. The response categories for this question are identical to those used for the variable on immigrants, ensuring consistency in the measurement of attitudes towards these two groups.

Figure 1 presents the mean scores of the outcome variables for respondents interviewed

[1] We recoded the attitude variables so that a higher value represented a more positive view on immigrants.

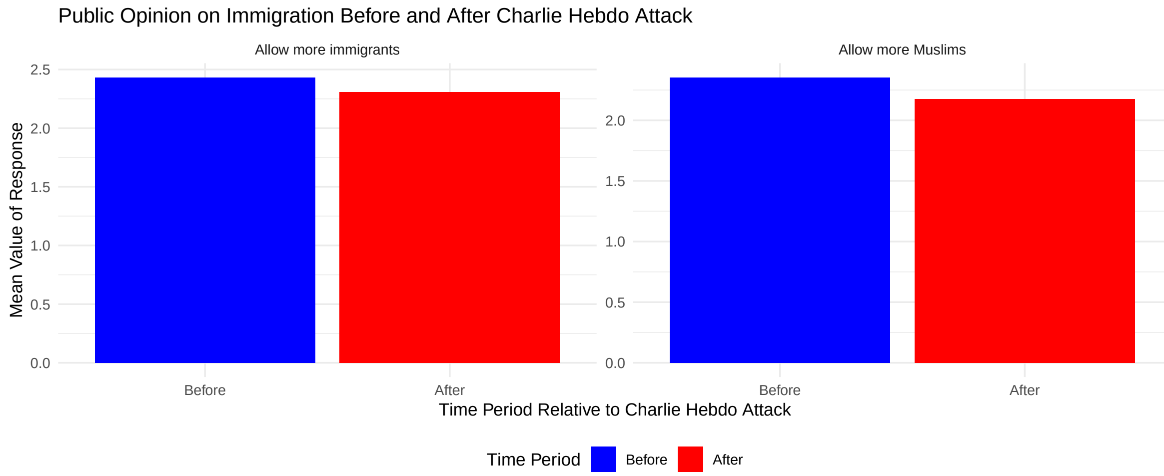

Figure 1: Means of outcome variables Before and After Charlie Hebdo Attack

before and after the Charlie Hebdo attack. An initial observation reveals that respondents interviewed post-attack exhibit more negative attitudes towards both immigrants and Muslims. However, we cannot interpret this difference in means causally unless we assess whether the key identifying assumptions of the UESD are met.

The treatment variable in our study is a binary indicator that segregates the respondents into two distinct groups: those interviewed before the attack date (07/01/2015) and those interviewed after. We deleted all respondents from the analyses who were interviewed on the day of the event, as we cannot be sure whether those respondents had been exposed to the event or not.

In our main specification, we incorporate a range of socio-demographic variables that are relevant for assessing the ignorability assumption. These include gender (with male as the reference category), age, education level (categorized into seven levels), income (segmented into deciles), living area (divided into five categories), self-reported political orientation (on a scale from 0, left-leaning, to 10, right-leaning), religiosity (with non-religious respondents as the reference group), and country of residence. The decision to use these specific variables was based on recommendations for this particular design Muñoz et al. (2020) but also on previous studies. As we progress to more advanced methodologies for estimating the causal effect that better handle high-dimensional datasets, we will introduce additional controls.

Table 1 presents the results of an imbalance analysis, which assesses whether the respondents in the pre- and post-intervention groups are sufficiently comparable regarding our included exogenous covariates. The last column provides the absolute value of the standardized difference in means; it is often recommended that the SMD should not exceed 0.25, which is the case for all our variables. Nonetheless, the smaller these differences are, the more accurate our causal effect estimate will be. We therefore discuss below methods for controlling for these imbalances to increase confidence that the assignment to pre- or post-intervention groups is essentially exogenous.

2.2 Methods

To address the ignorability assumption, most studies in the literature employ either matching/weighting techniques to improve covariate balance between treatment and control groups or use controls within a regression framework. To enhance our analysis and mitigate the limitations inherent in both approaches, we employ both methods in this study.

2.2.1 Regression-Based Methods

We begin our analysis with a Regression-Based Method (RBM) to estimate the average treatment effect (ATE) as the difference between the mean outcomes of treated and control groups, conditional on covariates. The ATE is given by:

$$\tau_{ATE}^{Reg} = E[f(1, X) - f(0, X)]$$

Table 1: Control Variables Stratified by Treatment Status

	Before	After	SMD
n	2886	3155	
Gender = Male (%)	1386 (48.2)	1502 (47.8)	0.0018
Age (mean (SD))	47.02 (18.06)	47.55 (17.30)	0.024
Education (%)			
Education Level I	207 (7.2)	153 (4.9)	0.0286
Education Level II	440 (15.3)	413 (13.2)	0.0165
Education Level IIIb	605 (21.0)	846 (27.0)	0.0617
Education Level IIIa	550 (19.1)	594 (19.0)	0.0082
Education Level IV	426 (14.8)	541 (17.3)	0.0263
Education Level V1	278 (9.7)	220 (7.0)	0.0364
Education Level V2	367 (12.8)	365 (11.6)	0.0148
Other	3 (0.1)	2 (0.1)	0.0001
Income (mean (SD))	5.39 (2.77)	5.17 (2.69)	0.0710
Political Orientation (mean (SD))	5.02 (2.17)	4.89 (2.03)	0.0772
Living Area (%)			
Big city	715 (25.0)	636 (20.3)	0.0497
Suburbs of big City	282 (9.9)	313 (10.0)	0.0029
Small City	829 (29.0)	1189 (37.9)	0.0928
Village	845 (29.6)	752 (24.0)	0.0509
Farm	185 (6.5)	247 (7.9)	0.0049
Religiosity = No (%)	1673 (58.2)	1750 (55.7)	0.0222
Country (%)			
Austria	38 (1.3)	230 (7.3)	0.0568
Belgium	255 (8.8)	64 (2.0)	0.0816
Switzerland	63 (2.2)	62 (2.0)	0.0050
Czech Republic	754 (26.1)	1182 (37.5)	0.1349
Germany	294 (10.2)	550 (17.4)	0.0976
Denmark	149 (5.2)	23 (0.7)	0.0483
Finland	132 (4.6)	162 (5.1)	0.0063
France	446 (15.5)	267 (8.5)	0.0843
Ireland	238 (8.2)	533 (16.9)	0.0795
Netherlands	326 (11.3)	36 (1.1)	0.1138
Sweden	51 (1.8)	11 (0.3)	0.0157
Slovenia	140 (4.9)	35 (1.1)	0.0264

Here, $f(1, X) = E[Y \mid A_n = 1, X_n = x]$ represents the mean outcome Y for the treated group (A = 1), and $f(0, X) = E[Y \mid A_n = 0, X_n = x]$ the mean outcome Y for the control group (A = 0). This estimation is feasible because treatment assignment can be considered random once we control for the covariates X. Formally, this implies:

$$\{Y_i^{(1)}, Y_i^{(0)}\} \perp A_i \mid X_i$$

The RBM that we employ is the Country Fixed Effects (FE) estimator, a method widely accepted across disciplines for estimating causal effects from non-experimental data, particularly effective when time-variant confounders are minimal. This is especially pertinent given the short timeframe (1 month) of our study. The country fixed effects approach has a shortcoming in our application: it does not take into account the dependence of observations nested within countries. In the robustness tests, we replace the country fixed effect estimator with the Multilevel Linear Regression (Random Intercept) to account for the hierarchical structure of the data. Another downside general to all RBM is that they assume linearity in the functional form of $f(1, X)$ and $f(0, X)$. Matching or weighting methods do not make such assumptions, making them more reliable for estimating the causal effect. We therefore complement the FE approach with matching and weighting methods.

2.2.2 Matching Methods

Matching methods create comparable groups of treated and control based on observed characteristics. The most popular matching methods are all based on Propensity-Score Matching (PSM). PSM facilitates balanced comparisons between treated and untreated subjects using their propensity scores. A propensity score is the probability that a respondent belongs to the treatment group conditional on the covariates X.

$e(X) = E[A = 1 \mid X]$. In our study, we employ several PSM techniques to estimate the ATE, including nearest neighbor (greedy matching), Coarsened Exact Matching (CEM), and optimal matching. All methods are applied without replacement to ensure unique pairings.

"*Greedy matching*" pairs subjects based on the closest propensity scores. Although computationally efficient, it can result in suboptimal pairings. Figure 11 in our appendices assesses the plausibility of the ignorability assumption by analyzing balance on pre-treatment covariates between treatment and control groups before and after adjustment through greedy matching. The figure plots the difference in means in standard deviation units of all sociodemographic characteristics of respondents as well as their place of residence. Although the algorithm has achieved considerable improvements in Standardized Mean Differences (SMD), there is still room for improvement for certain covariates like country of residence. To enhance matching precision, especially for these critical variables, "*Coarsened Exact Matching (CEM)*" is used in combination with greedy matching, striking a balance between precision and the practicality of sample size. CEM coarsens continuous covariates into broader categories to allow for exact matching within these groups. This combined matching approach significantly improved balance on the chosen variable for exact matching (country of residence), as evidenced by the SMDs presented in figure 12 in our appendices, showing a reduction to zero for these variables. However, improvement is still limited on other covariates. We therefore use "*Optimal matching*" to minimize the average within-pair difference in propensity scores across all pairs, rather than seeking the closest match per individual. This method, suited for our study's manageable sample size, significantly improved the balance across our sample as shown in figure 13 in our appendices by the SMD reductions below 0.05 for all covariates, including prioritized country of residence and political orientation, enhancing the credibility of our causal inferences.

2.2.3 Weighting Methods

The most popular of the weighting methods is the Inverse Probability of Treatment Weighting (IPTW) (See (Chesnaye et al., 2022) for application to observational data). IPTW balances the characteristics of treated and control groups in observational studies by assigning weights based on their propensity scores. Treated subjects are weighted by the inverse of their probability of receiving treatment, and control subjects by the inverse of their probability of not receiving treatment. This weighting simulates the conditions of a randomized controlled trial, adding robustness to our causal analysis by minimizing bias due to confounding variables. The ATE is given by:

$$\hat{\tau}_{IPW} = \frac{1}{N} \sum_{n=1}^{N} \left(\frac{A_n Y_n}{e(X_n)} - \frac{(1 - A_n) Y_n}{1 - e(X_n)} \right)$$

Another weighting method for causal inference that has gained popularity in recent years is Entropy Balancing Hainmueller (2012). This method adjusts the weights of the treatment and control groups such that the weighted distributions of the covariates in both groups match precisely. It reweights the sample by minimizing the entropy distance subject to covariate balance constraints, ensuring that the weighted means, variances, and higher-order moments of the covariates are equivalent across the treatment and control groups. The ATE is given by:

$$\hat{\tau}_{EB} = \frac{1}{N} \sum_{n=1}^{N} w_n Y_n A_n - \frac{1}{N} \sum_{n=1}^{N} w_n Y_n (1 - A_n)$$

where w_n are the weights derived to achieve covariate balance.

2.2.4 Machine Learning Methods

Machine learning (ML) methods are becoming increasingly popular for causal inference due to their flexibility and ability to handle complex data structures, as highlighted by (Baiardi and Naghi, 2021). The ML method we employ is the Augmented Inverse Propensity-Weighted (AIPW) estimator (Robins et al., 1994). The AIPW estimator involves a two-step process: First, we fit a propensity score model using machine learning techniques. Unlike traditional methods, these techniques do not assume linearity in the functional forms of $f(1, X)$ and $f(0, X)$, allowing for a more flexible and accurate modeling of the treatment assignment mechanism. Second, we estimate the treatment effects by weighting outcomes based on these propensity scores. The AIPW estimator is notable for its "double robustness," meaning that the estimator remains reliable as long as either the propensity score model or the outcome model is correctly specified. This robustness is particularly valuable, as it ensures consistency even if one of the models is misspecified. The ATE is estimated as:

$$\hat{\tau}_{ATE}^{AIPW} = \frac{1}{N} \sum_{n=1}^{N} \left(\underbrace{f(1, X_n) - f(0, X_n)}_{\text{Regression Adjustment}} + \left(\underbrace{\frac{A_n}{e(X_n)}(Y_n - f(1, X_n))}_{\text{Augmentation}} - \underbrace{\frac{1 - A_n}{1 - e(X_n)}(Y_n - f(0, X_n))}_{\text{Augmentation}} \right) \right)$$

In this formula, the first term $f(1, X_n) - f(0, X_n)$ represents the regression adjustment, capturing the difference in predicted outcomes for treated and control units. The second term involves augmentation components that adjust for any discrepancies in the propensity score model, thus ensuring the estimator's robustness.

Because ML methods are better capable of handling high-dimensional data sets, to further bolster our analysis against potential omitted variable bias, we expanded our control variables to include factors such as respondents' TV watching time, interest

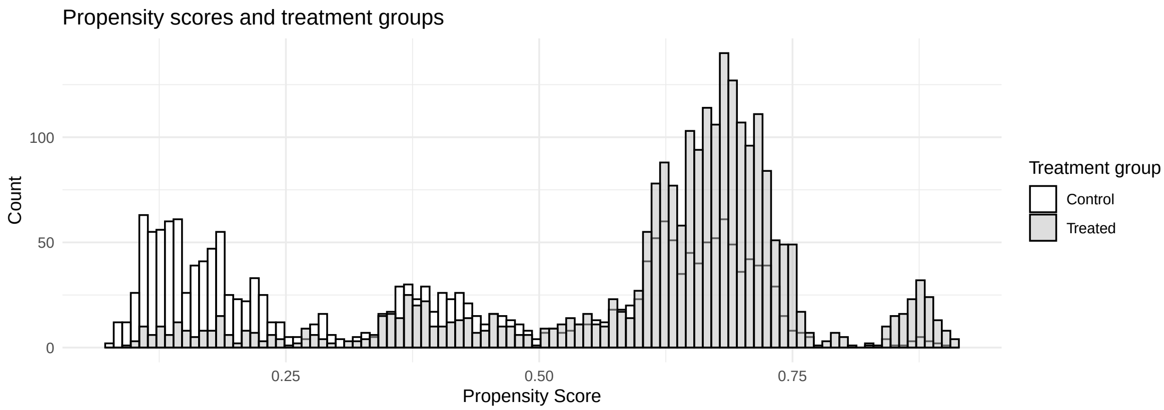

Figure 2: Propensity scores and treatment groups

in politics, satisfaction with life, happiness level, experiences of discrimination, and current employment status. We also applied piecewise polynomials for all continuous variables, increasing our total number of covariates to 55.

One key assumption for all matching and weighting methods that we have not discussed yet is positivity. This assumption states that every subject must have a non-zero probability of receiving treatment Zhu et al. (2021). We verified this assumption by examining the overlap in propensity score distributions, ensuring all scores fell within a typical range of 0.05 to 0.95, as illustrated in figure 2. We now proceed to the results section, where we present the estimated Average Treatment Effects (ATE) for each method, providing a detailed analysis of the societal impact of the Charlie Hebdo shooting on attitudes towards immigrants and Muslims.

3 Results

Figure 3 presents the point estimate of the Average Treatment Effects (ATE) of the Charlie Hebdo shooting on attitudes towards immigrants and Muslims, with 95% confidence intervals drawing on the ±30 days bandwidth. The figure shows a remarkably uniform picture as all seven methods indicate a statistically significant decrease in atti-

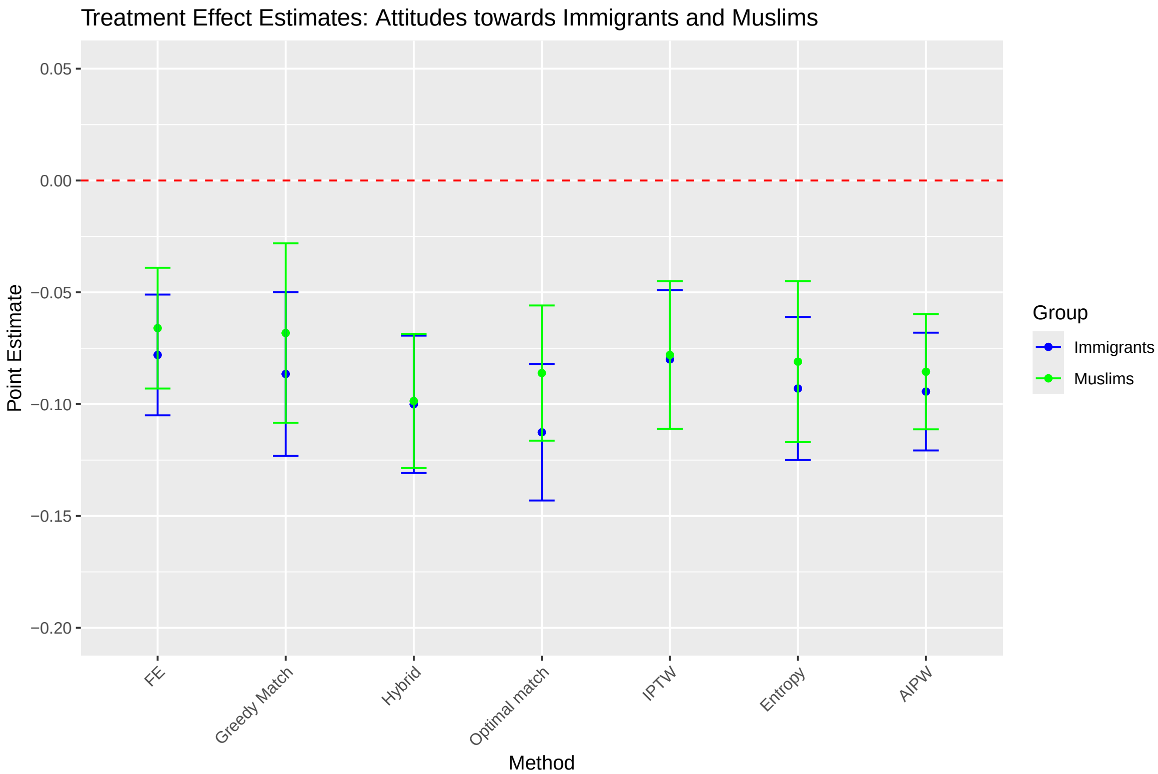

Figure 3: Effect of Terror on Attitudes towards Immigrants and Muslims
Note:

- This figure represents the Average Treatment Effects (ATE) of the Charlie Hebdo shooting on attitudes towards immigrants and Muslims. Attitudes towards immigrants and Muslims are measured from a survey question about allowing immigrants/Muslims in the country. Responses ranged from "*Allow none to come (1)*" to "*Allow many (4)*". Treatment is a binary indicator based on interviews conducted before or after the shooting (07/01/2015). Fixed effects estimator (FE) used 4,479 observations. All matching methods point estimate derived from a paired t-test after matching. Greedy Match: 2,278 matched out of 2,278 treated; Hybrid method combines Exact and Greedy matching: 1,359 matched out of 2,278 treated; Optimal Match: 1,359 matched out of 2,278 treated; IPTW (Inverse Probability Weighting): Utilized 4,413 unweighted observations. Entropy Balancing Utilized 4,413 unweighted observations. AIPW (Augmented Inverse Propensity-Weighted): Estimation involved 4,385 observations.

tudes, suggesting that those interviewed after the attacks had more negative attitudes towards immigrants and Muslims compared to those interviewed before. This finding corroborates results from previous research from both the United States and Europe that terrorism causes a negative shift in attitudes towards outgroups. The magnitude of these effects varies between 8% to 12% of the standard deviation pretreatment distribution for attitudes towards immigrants and from 6% to 10% for attitudes towards Muslims. Although the scale of these effects is modest relative to the standard deviation, they are more significant than the gender gap. This uniformity in the direction of the effects suggests a discernible average shift in public sentiment post-Charlie Hebdo in Europe, albeit not a drastic one.

Next, we test whether the shift in attitude happened immediately after the Charlie Hebdo attacks. For this purpose, we specify a model that includes a running variable ranging from -30 to 30 (the fieldwork days before and after the attacks), with 0 corresponding to January 8th, 2015. This variable is interacted with the treatment indicator. In this interactive model, the constitutive term for the "treatment group" variable corresponds to the effect of the terrorist attacks when the running variable "days" equals 0 (i.e., on January 8th, the first day immediately after the attacks). The interactive term, in turn, indicates whether or not the effects of the attacks changed (weakened or strengthened) as time went by after the attacks. This design is very similar to the Regression Discontinuity Design (RDD) as assignment to the treatment group varies discontinuously with the day of the interview. Specifically, the treatment effect is identified by comparing outcomes on either side of a threshold.[2]

[2]The equation that is estimated

$$Y_i = a + \beta_1 Treatment + \beta_2 Days + \beta_3 Treatment_i \times Days + \delta Controls_i + \epsilon_i \tag{1}$$

The set of dependent variables—contained in Y measures attitudes toward immigrants and Muslims. The main independent variable (Treatment) is equal to 1 for respondents answering the survey after the attack occurred (0 otherwise). The vector of control variables (Control) includes respondents' gender, age, education level, income, living area, self-reported political orientation, religiosity, and country of residence.

Table 2: Effects of Charlie Hebdo attacks on Attitudes

	Dependent variable:	
	Immigrants	Muslims
	(1)	(2)
Treatment	-0.141^{**}	-0.113^{**}
	(0.056)	(0.055)
Treatment:Days_Since_Attack	0.005	0.008^{***}
	(0.003)	(0.003)
Constant	2.857^{***}	2.735^{***}
	(0.109)	(0.109)
Observations	4,479	4,453
R^2	0.186	0.356

The "treatment" coefficient in the interactive model as shown in table 2 reveals that attitudes toward immigrants and Muslims shifted negatively immediately after the attacks on January 8th, 2015. Moreover, the coefficient of the interactive term (Treatment*Days) of this model is positive and significant for Muslims but not for immigrants, indicating that the effects of the attacks did not significantly change thereafter for immigrants but the negative trend in attitudes towards Muslims is less steep as it diminishes over time. Using the same design, we also tested whether the attack caused the same effect for attitudes toward unrelated groups such as Jewish and Gypsy populations and found no evidence of such effect (results available upon request).

One issue highlighted in research on terrorist attacks and public opinion is the potentially short-term impact of such events, whereby effects may last for merely days. Mancosu and Ferrin Pereira (2021), for example, find that negative attitudes against immigrants after the 2017 Manchester bombing lasted only for a week. To assess the duration of the effect beyond one month, we considered comparing attitudes of those interviewed in the 30 days leading to the attack to the attitudes of those interviewed from the 31st day to the 60th day following the attack. Unfortunately, too few respondents were interviewed in the post-attack period and most of those interviewed were from a single country - Austria. Data limitations mean that we are unable to effectively assess the long-term effect of the attack. The best we can do here is to investigate to what extent the absolute size of the effect varied by the length of the time window, which we discuss in the next section.

Before discussing further robustness checks, we briefly examine the estimated coefficients for the control variables in the fixed effects (FE) model. The controls display the expected signs: women report more negative attitudes compared to men; older respondents also exhibit more negative attitudes. In contrast, respondents with higher levels of educational attainment report significantly more positive attitudes. Political alignment plays a role as well, with those on the right of the political spectrum reporting more negative attitudes. Additionally, respondents residing in rural areas show more negative attitudes. Geographic variation is also evident: respondents living in Belgium, the Czech Republic, Finland, and Ireland report below-average attitudes. Those in Switzerland, Germany, France, and Sweden report above-average attitudes, while respondents in Denmark, the Netherlands, and Slovenia report average attitudes.

3.1 Robustness checks

To assess the robustness of our results to the choice of event window, we varied the bandwidths, considering -10 to 10, -20 to 20, -40 to 40, -50 to 50, and -60 to 60 days. For each bandwidth, we ensured that survey respondents were balanced around the event date after controlling for confounders. Figures 14 and 15 in the appendices confirm that our conclusions about the negative impact of the Charlie Hebdo shooting on attitudes toward immigrants and Muslims hold across different time windows and methods (FE, Greedy Match, Hybrid Match, Optimal Match, IPW, Entropy Balancing, or AIPW). The precise number of days or the method used does not change our conclusion.

Furthermore, as we increase the time window, the magnitude of the impact diminishes. This pattern reinforces the credibility of our conclusion that the attitude shift is due to the shooting and suggests that the impact is short-lived, with attitudes likely returning to their pre-shooting levels a few months after the event. We limit our analysis to two months post-event due to data availability and to avoid issues with causal inference, as unrelated events, such as the massive arrival of migrants in Europe from March 2015, could bias the findings.

To assess the validity of the exclusion restriction, which implies that the timing of the survey only affects the outcome of interest through exposure to the event, we performed two falsification tests. These tests examine whether the observed effects are uniquely linked to the Charlie Hebdo event or if similar patterns could occur in the absence of such a dramatic incident. If our findings are caused by a declining trend in attitude, we would find a statistically significant treatment effect when assuming that the attacks happened one month before its actual date on December 7th, 2014 (Placebo 1) or two months before its actual date (Placebo 2) on November 7th, 2014. The results of these tests, as shown in 16 and 17 in the appendices, demonstrate that our main findings are not driven by preexisting time trends unrelated to the event of interest. All seven methods indicate no treatment effect or a slight improvement in attitudes towards immigrants, in contrast with the decline observed on the actual attack date. These findings, particularly the consistency across all methods, further validate our primary results that the shifts in attitudes are directly associated with the Charlie Hebdo incident, rather than a general decline in sentiment towards immigrants and Muslims in Europe.

As mentioned earlier, an important shortcoming in the country fixed effect estimator that we use in our initial analysis is that it does not take into account the dependence of the observation nested within countries. We therefore replace it with the Multilevel Linear Regression (Random Intercept) to account for the hierarchical structure of the data. The results (see table 3 in the appendices) remain the same for both attitudes towards immigrants and Muslims. We also perform leave-one-out analyses to ensure that the main results are not driven by one single country. The results (see table 4 and 5 in the appendices) hold when each country is left out one at a time.

3.2 Heterogeneous Findings

So far, we have focused on the average impact of the attack on all 12 European nations, ignoring that we are dealing with largely differentiated contexts and environments that can moderate or exacerbate the average negative effect observed. The heterogeneous nature of terrorism's impact is evident in the varied findings of previous studies. For instance, (Legewie, 2013) observed significant cross-national variations in the anti-immigrant sentiments triggered by the Bali attack with pronounced effects found in countries like Portugal, Poland, and Finland, while no impact is found in others including Belgium, Switzerland, Sweden, the Netherlands, Norway, and Great Britain. (Savelkoul et al., 2022) reported mixed responses within Europe to the Charlie Hebdo attack, with notable differences in public resistance towards Muslims across countries such as Ireland, the Czech Republic, France, Austria, Finland, and Germany.

Given this backdrop of varying national responses, our study aims to delve deeper by estimating the impact of the Charlie Hebdo attack on attitudes towards immigrants and Muslims for each country within our sample that has a sufficient number of respondents both before and after the attack. Based on the data presented in Table 1, we have identified five countries (the Czech Republic, Germany, Finland, France, and Ireland) that meet the criterion of having at least 100 respondents interviewed both before and after the attack. We summarize below the key findings:

Czech Republic: The findings for the Czech Republic reveal a consistent trend of declining positive attitudes towards both immigrants and Muslims following the Charlie Hebdo attack (See figure 18 in the appendices).

Finland: For Finland, there is a clear indication of a decline in positive attitudes towards immigrants, as shown by all the methods applied (see figure 19 in the appendices). A similar trend of a negative shift is observed for attitudes towards Muslims, with three of the methods reporting a minor decrease in positive attitudes while the other four methods indicate no significant change.

Ireland: The analysis for Ireland presents a picture similar to Finland, with five methods suggesting a minor to moderate decline in positive attitudes while the other two methods indicate no significant change in attitude toward immigrants. In contrast, the negative impact on attitudes towards Muslims is more pronounced and consistent across methodologies (see figure 20 in the appendices).

Germany: The analysis of Germany's response to the Charlie Hebdo attack offers an intriguing perspective on how such events can influence public attitudes in varying national contexts. Across all methodologies, the data indicates no significant shift in attitudes towards immigrants post-Charlie Hebdo attack. The results for attitudes towards Muslims in Germany also point to either no significant change (two methods) or a slight improvement (five methods) post-attack. Overall, the results for Germany suggest a notable improvement in public sentiment towards Muslims and no significant change towards immigrants in the aftermath of the Charlie Hebdo attack (see figure 21 in the appendices).

France: As we turn to France, the focal point of the Charlie Hebdo attack, we inherently anticipate the strongest reactions here due to the attack's local occurrence. Typically, in the context of terrorist attacks, especially one as significant as Charlie Hebdo, one might expect a deterioration in attitudes towards both immigrants and

Muslims, assuming that the proximity and direct impact of the event would intensify public sentiment. Surprisingly, the results for France show a different trend than initially expected. Five of the methods employed suggest an improvement in attitudes towards both immigrants and Muslims following the Charlie Hebdo attack. The other two methods indicate no significant change (see figure 22 in the appendices).

These country-specific results highlight the diverse ways in which societies react to terrorist attacks, influenced most likely by local media narratives, historical context, and societal values. The next section will discuss the implications of these findings and explore potential explanations for the variations observed.

3.3 Discussion

Reflecting on the overall negative impact observed in the sample when all countries are combined and the contrasting country-specific impacts—ranging from negative shifts in the Czech Republic, Finland, and Ireland to improvements in France and Germany—we must examine underlying theories and alternative explanations. Two main theoretical frameworks typically used to explain reactions to terrorist attacks are "*System Justification Theory (SJT)*" (Jost and Banaji, 1994) and '*Intergroup Conflict Theory (ICT)*" (of Oklahoma. Institute of Group Relations and Sherif, 1961). These theories explain negative reactions to terrorist attacks through perceived threats and competition over resources. Recent work has combined explanations from SJT and ICT to better understand the driving forces behind prejudice toward immigrants following terrorism.

However, applying SJT to our findings does not offer a straightforward explanation, especially for France, where one might expect a heightened sense of threat and consequently more negative attitudes, yet we observe the opposite. One possible explanation for the SJT mechanism not being activated is that fear of being targeted in future attacks was probably not relevant in this case since the attack was specifically directed toward the Charlie Hebdo journal and not the general public. How people perceive their own risk of being a victim of a terror attack and the motive of the terrorists seem to play an important role in determining how people will respond to an attack Jakobsson and Blom (2014). Moreover, statistics on the Muslim population in 2015 (higher for France and Germany) and unemployment rates across the five countries do not consistently align with predictions from intergroup conflict theory.

A theoretical argument that could help explain this puzzling finding is based on the contact hypothesis Allport et al. (1954), which states that increased personal interaction between two groups improves attitudes (for an overview, see Pettigrew and Tropp (2006)). Personal interaction between natives and Muslim immigrants in Germany and France may have helped reduce pre-existing stereotypes and thereby stabilize anti-immigration attitudes in these countries (Enos (2014).

In short, the arguments based on SJT and ICT suggest that contextual factors moderate the extent to which terrorism might result in anti-immigrant attitudes. The magnitude of the public response depends on various factors, including people's economic position and their resulting vulnerability to changes in the labor market, or the extent to which they deem cultural homogeneity within their country an important policy goal. Another important contextual factor that is often overlooked is the media environment surrounding the incident. This study argues that media coverage of immigrants and Muslims following terrorism can be an important mediating factor.

4 Role of the media

It is widely acknowledged that information relayed through mass media influences the formation of anti-immigration attitudes. However, comparative studies on how media coverage of terror incidents affects public opinion toward specific outgroups are limited.

To date, research has primarily focused on identifying the causal link between terrorism and attitudes toward outgroups, without exploring the specific channels through which this shift occurs, aside from the size and increase of the outgroup population and unemployment levels as mediating factors. This study argues that in addition to the societal and economic context, news coverage of the attack itself must be considered as a contextual factor.

To understand how news coverage of a terror incident affects public opinion about immigrants and Muslims, it is important to assess whether the media environment surrounding the incident encourages or discourages assigning responsibility to these groups. Mainstream media coverage influences readers' assessments of the gravity of the act and the reasons behind it. When mass media highlight immigration and Muslims in their coverage of the incident, the audience is more likely to view these groups as relevant to the act. This is particularly significant for topics like terrorism, which the general public finds difficult to understand and has little direct experience with. Consequently, the media often serve as the primary information source, leading to strong fluctuations in how incidents are interpreted and responsibility is assigned.

Media influence on public opinion is traditionally understood through agenda-setting theory, which posits that frequent reporting on an issue increases its perceived importance Maxwell (2004). However, understanding media influence requires examining not only the salience of issues but also how they are framed. Media frames are thematic structures that shape audience interpretation, pushing them to view issues within a specific context Erhard et al. (2022). Distinguishing between agenda-setting and framing is crucial for accurately attributing media effects on attitudes, especially in sensitive domains like immigration and ethnic relations post-terrorism.

In this context, we first focus on the agenda-setting effects on attitudes toward immigrants and Muslims. We explore the relationship between the visibility of immigrants and Muslims in news about the Charlie Hebdo shooting and shifts in public opinion about these groups. We hypothesize that countries where the media frequently link immigrants and Muslims with the incident will exhibit the most negative shifts in public opinion.

However, analyzing the visibility of immigrants and Muslims in Charlie Hebdo news articles is insufficient to understand their impact on public opinion. We also assess how the content, specifically the stance of the news toward immigrants and Muslims, could explain differences in attitudes across the five countries. We focus on how the media influence public opinion by assigning responsibility for the incident to these groups or by determining which political claims and comments about immigrants and Muslims are relevant to the discussion. We hypothesize that countries with a higher share of unsupportive Charlie Hebdo news articles about immigrants and Muslims will show the most negative shifts in public opinion.

These expectations are tested through descriptive analysis using a corpus of 19,320 news articles published online up to a month after the Charlie Hebdo attack, containing the keyword "Charlie Hebdo." This corpus includes a wide array of local and national perspectives, predominantly from France (16,284) and Germany (2,671), where no negative shift in public opinion was observed, and Ireland (216), the Czech Republic (131), and Finland (18), where a negative shift was noted. These countries differ significantly in the size of their outgroup populations, allowing us to rule out outgroup size as an explanatory factor for the differences in reactions. For the analysis, we assume that all survey respondents were exposed to news about the incident either directly or through interpersonal communication. We use two Natural Language Processing (NLP) methods for analysis: Large Language Models (LLMs) and Structural Topic Modeling (STM).

4.1 Large Language Models (LLMs)

The emergence of Large Language Models (LLMs) like ChatGPT has revolutionized the field of Natural Language Processing (NLP), offering unprecedented opportunities for understanding and analyzing human-produced text. LLMs, with their vast knowledge bases and contextual understanding capabilities, have become powerful tools in various NLP tasks. One such task is stance detection, a nuanced form of sentiment analysis focused on identifying the position or perspective expressed in text towards a specific subject or entity. Unlike general sentiment analysis, which classifies text as positive, negative, or neutral, stance detection delves into understanding the explicit or implied stance towards particular targets, often crucial in analyzing public opinions on controversial topics Küçük and Can (2020).

Recent studies have showcased the impressive performance of LLMs across a range of NLP tasks (Cruickshank and Ng, 2023). For instance (Zhang et al., 2022) evaluate ChatGPT's performance on detecting political stance and highlight ChatGPT's proficiency in political stance detection, often outperforming traditional models. (Zhang et al., 2023) provide a comprehensive investigation into the capabilities of LLMs in performing various tasks, confirming that LLMs can significantly surpass standard language models, especially in scenarios requiring a nuanced understanding of text. Nonetheless, despite their impressive capabilities, LLMs are not without challenges. Variability in performance across different versions Aiyappa et al. (2023), the tendency to produce plausible but incorrect answers (hallucinations) Cruickshank and Ng (2023), and limitations in handling complex tasks requiring deep understanding Zhang et al. (2023) are among the issues researchers and practitioners face. With these limitations in mind, the next section will delve into our strategy for employing these models in analyzing the media coverage of the Charlie Hebdo attack. We will discuss the prompting strategies designed to optimize the quality of outputs and prevent undesirable behaviors like hallucination.

4.1.1 Prompting Strategy

As we delve into the sophisticated world of Large Language Models (LLMs) such as GPT-3.5 Turbo, it becomes crucial to master the art of prompt engineering—the technique of crafting effective prompts to elicit accurate and relevant responses from these models Schmidt et al. (2023). A few key techniques in prompt engineering that are particularly relevant for LLMs include: "*Few-Shot Prompting:*" Providing the model with a few examples of the desired output to guide its responses White et al. (2023); "*Chain-of-Thought Reasoning:*" Encouraging the model to "think aloud," detailing its reasoning process step-by-step, which often enhances the quality and reliability of its responses Wei et al. (2022); "*Preventing Hallucinations:*" Introducing strategies to minimize the risk of the model generating plausible but incorrect or unverifiable information Huo et al. (2023). In light of these techniques, and over several iterations and adjustments, we designed a prompt that consistently yields the desired outputs while containing only essential components to identify and analyze the stance of news articles towards immigrants and Muslims. The prompt includes:

- A clear task definition: Asking the model to identify whether Muslims or Immigrants are discussed in the article and to discern the presence of supportive or unsupportive opinions about them.
- Desired output format: Specifying a straightforward 'Yes' or 'No' response format to facilitate the classification of the articles' sentiments.
- Evidence requirement: Asking the model to provide evidence from the article for each answer to minimize hallucinations and ensure reliability.

Table 6 in the appendices showcases an example of the employed prompt and the corresponding output obtained from GPT-3.5 Turbo.

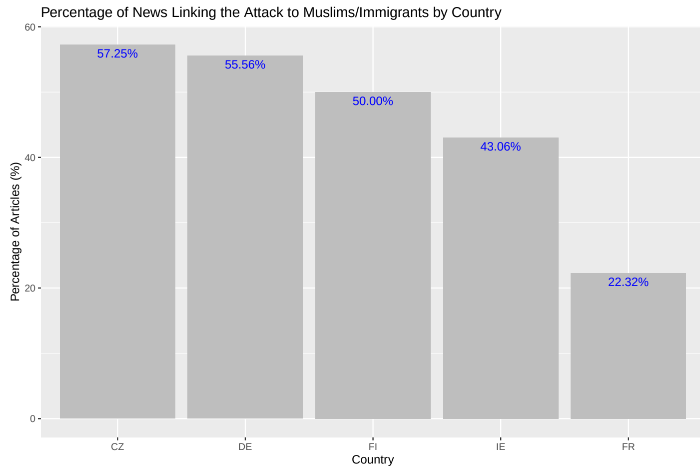

Figure 4: Percentage of News Linking the Attack to Muslims/Immigrants by Country

4.1.2 Result of the LLM

In our quest to understand media influence on public attitudes post-Charlie Hebdo, we have utilized the GPT-3.5 Turbo, a popular Large Language Model developed by OpenAI. To manage this extensive data and ensure accuracy in model responses, we undertook two essential preprocessing steps:

Sample Reduction for France: Given the overwhelming number of articles from France, we randomly selected 2,500 articles for a more manageable and representative analysis. This selection was meticulously done to retain the distribution of news sources. This step was crucial for ensuring the feasibility of manually verifying model-generated evidence and categorizing the opinions expressed in the articles.

Multilingual Corpus Management: To address the multilingual nature of our data, we employed Microsoft machine translation APIs to translate non-English articles into English, the language in which LLMs, particularly GPT-3.5 Turbo, show superior performance. This step ensures consistency in analysis and broadens the accessibility of our findings to an English-speaking audience.

Our analysis began by exploring the extent to which news articles link the Charlie Hebdo attack to Muslims or immigrants. Figure 4 presents the percentage of news articles associating the attack with these groups in each country. Interpreting these percentages within the framework of agenda-setting theory provides a nuanced understanding of media's impact. The theory posits that a higher frequency of media coverage, or salience, should correlate with increased public concern about the topic. Intriguingly, our findings align with this theory for France, where lower media salience correlates with an improved public attitude post-attack. However, the results for Germany present an unexpected deviation. Despite a high salience of immigrants and Muslims in the media post-Charlie Hebdo attack, reflecting a supposedly heightened public concern, attitudes towards these groups improved. This divergence in Germany, contrasted with the alignment in France, suggests that while agenda-setting theory can provide some insights, it does not fully encapsulate the complex dynamics at play. This discrepancy further emphasizes the necessity to explore beyond mere salience to the specific stance and framing of news articles to understand their influence on public attitudes.

Moving forward, we will focus on the second question posed to the LLM regarding the presence of supportive and unsupportive opinions about immigrants and Muslims in the news articles. The distribution of news articles with supportive and unsupportive opinions about immigrants or Muslims post-Charlie Hebdo, as shown in figure 5, presents a much clearer picture.

Germany: Clearly stands out with more supportive opinions than unsupportive opinions toward immigrants and Muslims. This is consistent with the improvement in public attitudes towards these groups in Germany. This finding suggests that more positive coverage of outgroups following terror can lead to more positive public perceptions.

France: Interestingly, France exhibits the lowest percentages of both supportive and unsupportive news articles, with as many supportive as unsupportive, offering balanced

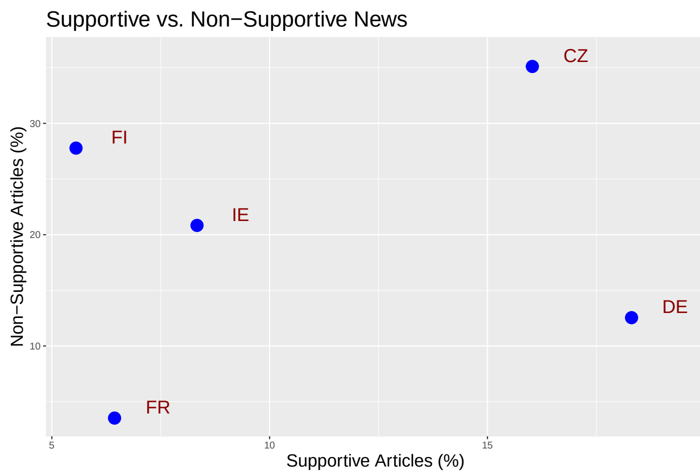

Figure 5: Supportive vs. Non-Supportive News Percentages by Country

coverage of immigrants and Muslims following the incident. This is consistent with the improvement or stability in public attitudes towards immigrants and Muslims in France. The result from France reinforces the idea that balanced coverage of minorities following terror can also lead to more positive public perceptions.

Czech Republic Finland and Ireland: These three countries have a higher share of unsupportive opinions than supportive ones, aligning with the more negative shifts in public attitudes observed in these countries post-attack. The Czech Republic, for instance, exhibits the highest percentages of news articles with non-supportive opinions, closely followed by Finland and Ireland. Overall, the results support the idea that a negative media environment towards Muslim immigrants following a terrorist attack can increase hostility towards these groups.

As we delve deeper into the media's influence, our next step involves categorizing the various themes present in the evidence provided by the model. This approach was crucial for ensuring accuracy in our analysis; each piece of evidence cited by the model to support the identification of supportive or unsupportive opinions was manually verified to confirm its presence in the actual text. This verification process helped us minimize any potential hallucinations. During this verification, we also classified the evidence

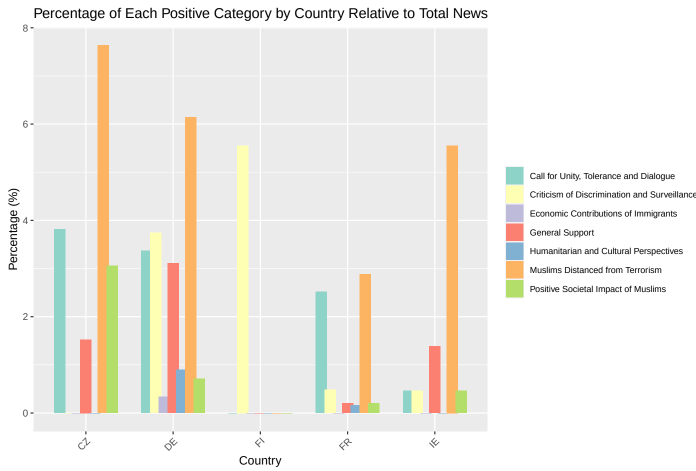

Figure 6: Percentage of Each Positive Category by Country Relative to Total News

into broader categories of supportive and unsupportive opinions. This classification allowed us to analyze if certain arguments were more prevalent in specific countries. For instance, Figure 6 organizes supportive opinions into themes such as calls for unity, criticism of discrimination, the positive societal impacts of Muslims, and narratives distancing Muslims from terrorism by country. This categorization by country helps us better understand the nuances of media narratives in each country and their potential impact on public sentiment.

Muslims Distanced from Terrorism: This theme captures articles where Muslims or their leaders explicitly condemn terrorism and distinguish their faith from extremist acts. Example: "The rector of the Givors Mosque strongly condemns the attack on the satirical newspaper and calls it an inadmissible act." This theme is particularly present in the Czech Republic, Germany, and Ireland.

Call for Unity, Tolerance, and Dialogue: This category includes articles advocating for solidarity, understanding, and open conversation among different communities. Articles in this category usually call for a dialogue between religious leaders or highlight community leaders advocating for peace and unity. Example: "The article calls for dialogue between priests, rabbis, and imams to define common values." This category is mostly present in the Czech Republic, Germany, and France.

Criticism of Discrimination and Surveillance: Articles in this category criticize unfair treatment, discrimination, or excessive surveillance targeting Muslims or immigrants. Example: "Chancellor Merkel's statement protects Muslims against general suspicion and affirms that the vast majority of Muslims in Germany are constitutional citizens." This category is dominated by Finland, followed by Germany.

Positive Societal Impact of Muslims: Articles in this category usually describe heroes issued from immigration or who are Muslims. Example: "It describes Bathily's efforts to save the customers and his willingness to help the police with information." This category is mainly present in the Czech Republic.

Humanitarian and Cultural Perspectives: This theme captures articles that show support for multiculturalism or discuss the humanitarian obligation of Europeans. Example: "The article highlights the need to address Germany's humanitarian obligation to people who flee their home countries due to war and persecution." This category is present only in Germany and France and represents less than 1% of each country's news articles about Charlie Hebdo.

Economic Contributions of Immigrants: This category includes articles that remind Europeans that they need immigrants out of their own economic interest. Example: "The Siemens CEO promotes more immigration and tolerance in Germany, saying that Germany needs more openness out of its own economic interest. We must be attractive to top talent." This category is also present only in Germany and represents less than 1% of Germany's news articles about Charlie Hebdo.

General Support: This category encompasses articles expressing sympathy or support for Muslims and immigrants in a general sense. Example: "Integration Commissioner Aydan Özoguz calls for support for Muslims in Germany and emphasizes the need for

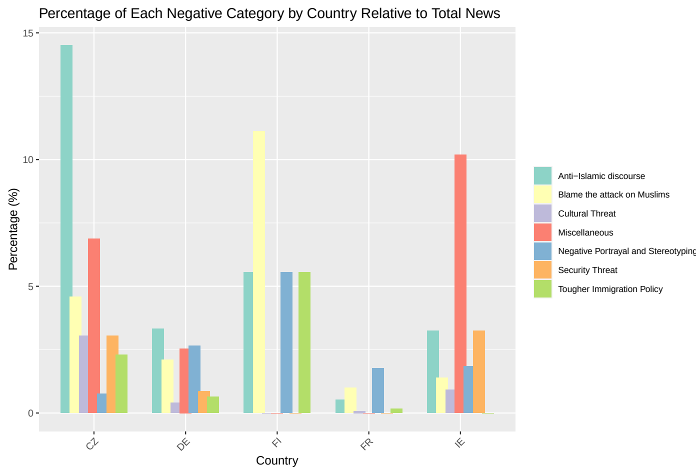

Figure 7: Percentage of Each Negative Category by Country Relative to Total News

unity." This category is mostly present in Germany, followed by the Czech Republic.

In conclusion, Germany clearly stands out with a significant percentage of articles featuring a variety of supportive themes, especially highlighting that the terrorists do not represent Muslims' faith, warning against discriminating against Muslims, and calling for unity, tolerance, and dialogue. This aligns with the observed improvement in public attitudes towards these groups in Germany. The Czech Republic also shows a notable focus on distancing Muslims from terrorism and on unity, tolerance, and dialogue, even though attitudes did not improve in the Czech Republic as the country also featured the highest share of unsupportive opinions. France presents a mix of supportive categories but at lower percentages compared to Germany and the Czech Republic. Interestingly, despite these lower percentages, public attitudes in France improved post-attack, mostly because France also featured the lowest share of unsupportive opinions. Ireland and Finland, with lower percentages of supportive opinions, reflect the more negative shifts observed in public attitudes post-attack, indicating a potential correlation between less supportive media coverage and more adverse public sentiment.

Similarly, figure 7 categorizes non-supportive opinions into several themes. We discuss below the most common categories with examples.

Anti-Islamic Discourse: Articles in this category contain critical or hostile views towards Islam as a religion or generalize Muslims negatively. Example: "The article presents multiple opinions that are critical of Islam and suggests there is a problem with the faith itself leading to terrorism." This category is mainly present in the Czech Republic.

Negative Portrayal and Stereotyping: This theme includes articles portraying Muslims or immigrants in a stereotypical or negative light, often associating them with various societal issues. Example: "The article mentions that not a minority, but a significant proportion of Muslims who refer to the Koran hold extreme attitudes." This theme is mainly present in Finland.

Blame the Attack on Muslims: Here, the articles specifically attribute the cause or responsibility for terrorism directly to the Muslim community or faith. Example: "Opinions suggesting Muslims need to ask themselves why so many terrorists invoke the Koran." This category is dominated by Finland, followed by the Czech Republic.

Security Threat: Articles categorized here discuss Muslims or immigrants as a potential security risk or threat. Example: "Concerns about young Muslims leaving for jihad in Syria or Iraq." This category is mainly present in the Czech Republic and Ireland.

Tougher Immigration Policy: Articles categorized here discuss the need for a tougher immigration policy as a response to the attack. Example: "The Prime Minister speaks out against Islamism, and there is mention of the need for deportation of rejected asylum seekers." This category is mainly present in Finland and the Czech Republic.

Cultural Threat: This category includes articles that criticize multiculturalism or warn against the Islamization of the West. Example: "Petr Fiala, who warns of the dangers of radical Islam and suggests that multiculturalism and political correctness have failed." This category is mainly present in the Czech Republic.

Miscellaneous: This broad category encompasses anything that could not be classified within other categories. Example: "The article implies that the blasphemy law is being used by Muslim terrorists to justify their atrocities." This category is mostly present in Ireland and the Czech Republic.

As we conclude our analysis of the results from the Large Language Model (LLM), the intricate relationship between media portrayals and public attitudes becomes more evident. The detailed categorization of supportive and non-supportive opinions provided by the LLM offers a granular understanding of how immigrants and Muslims are represented in the media following the Charlie Hebdo attack. These insights highlight the important role of media in shaping public attitudes and underscore the need for deeper analysis to fully understand these dynamics. To verify and further explore these findings, we will now employ Structural Topic Modeling (STM) as an alternative natural language processing (NLP) method. This approach will allow us to validate the insights gained from the LLM analysis and examine the robustness of our findings, particularly focusing on whether Germany still appears as the country with the highest percentage of supportive news, and the Czech Republic, Finland, and Ireland as the countries with the highest percentage of unsupportive news, aligning with the shifts in public attitudes observed in these countries.

4.2 Structural Topic Modeling (STM)

To assess the validity and robustness of the results obtained using LLMs, we next apply Structural Topic Modeling (STM) (Roberts et al., 2013) to classify and analyze themes in news coverage produced by the incident in France, Germany, the Czech Republic, Finland, and Ireland. We focus specifically on articles that have been identified by the LLM as discussing both the shooting and immigration as well as Muslims, aiming to discern the stance—supportive or unsupportive—of these news articles toward immigrants and Muslims.

STM, a derivative of topic models, facilitates the discovery of latent themes across large text corpora by examining word co-occurrence within documents. Unlike traditional topic models, which assign each word in a document to a single topic (single-membership models like Latent Dirichlet Allocation), STM (Roberts et al., 2014) allows documents to exhibit multiple topics (mixed-membership). STM's distinctive feature is its ability to incorporate document metadata (e.g., publication country, publication year) into the analysis, enabling the model to adjust topic prevalence and content based on these attributes.

4.2.1 Model Specification Choices

Covariates: The first choice in the STM model specification is to choose covariates to incorporate in the model. The only covariate included in our model is the country of publication, assumed to influence the prevalence of topics but not the content. Including such a covariate allows us, for instance, to determine how frequently specific themes appear in articles from different countries.

Number of Topics k: The second choice is the number of topics (k) to be discovered. Selecting the appropriate k depends on the desired granularity and the specific research questions. For this study, we trialed models with k ranging from 8 to 14, evaluating each model for coherence (documents within a topic discuss the same theme) and exclusivity (documents discussing similar themes are grouped under the same topic). A k value of 10 was chosen as optimal, providing a clear distinction between topics without sacrificing coherence.

4.2.2 Results of STM

We analyzed the discovered topics by identifying exemplar documents, those with the highest probability of belonging to a topic, thus being the most representative of the idea being discussed. Each topic was assessed for its supportiveness toward immigrants and Muslims based on the dominant tone in its exemplar documents, five of which were selected for each topic and provided to three different LLMs (Perplexity.ai, Claude.ai, and ChatGPT) [3] to assess the supportiveness of the document toward immigrants and Muslims. The overall tone of a document was determined by majority classification from the LLMs. Similarly, a topic is classified as Supportive/Unsupportive if the vast majority, at least 4 out of 5 of its exemplar documents, are classified that way.

Based on the classifications provided in Table 7 in the appendices, three topics clearly stand out as very supportive of immigrants and Muslims, with at least 4 out of 5 of their exemplar documents classified as supportive toward these groups. The supportive topics include: **Topic 1:** *Interfaith Solidarity and Condemnation of Violence.* The exemplar documents consistently promote a narrative of unity, peace, and condemnation of violence, involving leaders from multiple faiths who collectively reject terrorism and advocate for interfaith solidarity and understanding. **Topic 3:** *Expressions of National Unity and Democratic Values.* The exemplar documents focus on defending democratic principles and national unity in response to terrorism, emphasizing an inclusive approach that transcends religious and cultural differences. **Topic 8:** *Community Responses to Anti-Islam Demonstrations.* The exemplar documents portray a consistent theme of widespread public demonstrations supporting diversity and inclusivity, often overshadowing anti-Islam demonstrations.

[3] These LLMs were chosen because they were the most popular at the time of the analysis in July 2024. Gemini was not included in the analysis as it refuses to answer any question that could be linked to politics. The prompt used: "The news article below was extracted in the month following the Charlie Hebdo shooting, I would like you the assess whether the articles express opinions that could be considered supportive or unsupportive of Immigrants/Muslims."

Regarding unsupportive topics, none can be clearly classified as fully unsupportive. However, some topics can be classified as Neutral and leaning slightly toward unsupportive. These are topics with 3 out of 5 of their exemplar documents classified as Neutral and the remaining 2 as unsupportive. Two such topics are: **Topic 4:** *Controversial Discourse in the Wake of the Paris Attacks.* The exemplar documents report several controversial statements regarding Muslims, including comments by President Miloš Zeman about Muslims' inability to adapt genetically, without endorsing or condemning these views. **Topic 5:** *Societal Reflections and Cultural Discourse in Contemporary Literature.* This topic engages with complex and sometimes controversial reflections on modern society, using literature (e.g., Michel Houellebecq's novel "Submission") as a mirror to societal fears and aspirations. While the articles maintain a generally neutral stance toward Muslims, the novel itself is controversial, stirring debates about societal fears of Islamization.

To summarize, the STM analysis has uncovered a diverse range of themes in media coverage regarding immigrants and Muslims following the Charlie Hebdo attack. These themes include highly supportive topics focused on collective interfaith efforts to condemn violence and promote unity (Topic 1), the defense of democratic values and unity in the face of terrorism (Topic 3), and strong community support for diversity, opposing xenophobic rhetoric (Topic 8). On the other hand, neutral to slightly unsupportive topics include divisive rhetoric (Topic 4) and discussions of Michel Houellebecq's controversial novel "Submission" (Topic 5). Additionally, neutral topics (Topics 2, 6, 7, 9, and 10) offer nuanced perspectives that are not clearly supportive or unsupportive.

To explore the distribution of these themes across Germany (DE), France (FR), Ireland (IE), the Czech Republic (CZ), and Finland (FI), we employed linear regression to estimate the proportion of documents from each country that discuss supportive or unsupportive topics, as identified in the STM model. This approach incorporates measurement uncertainty from the STM model using the method of composition, as

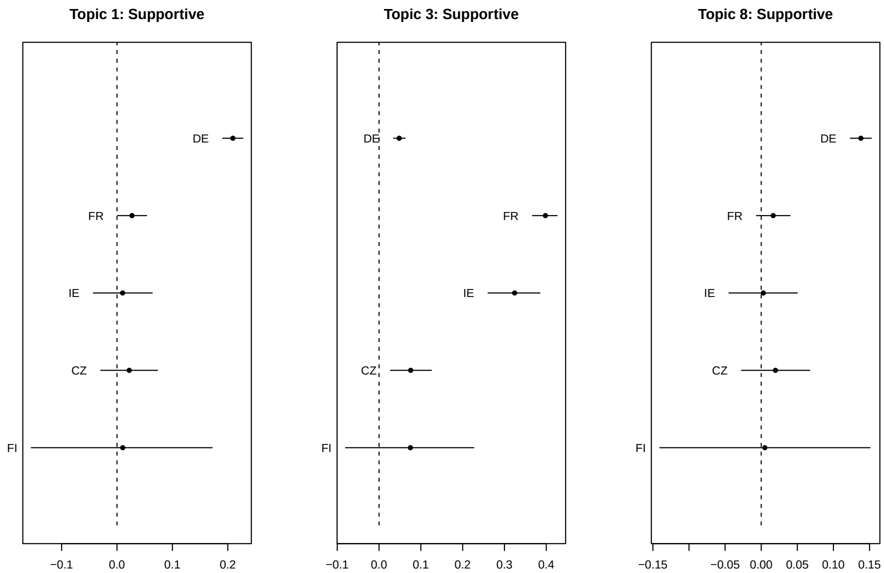

Figure 8: Estimate of Supportive Topics Proportion across Countries

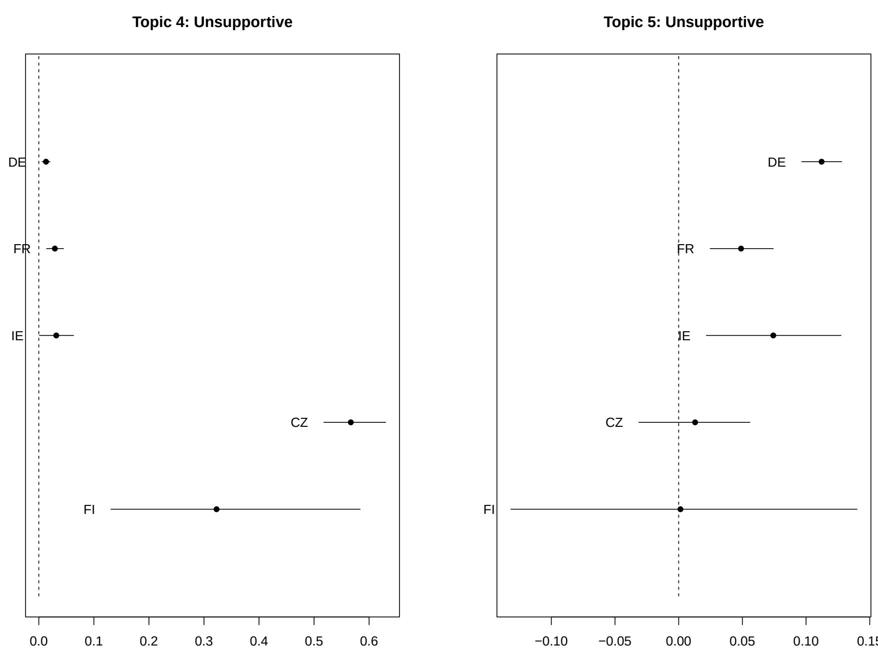

Figure 9: Estimate of Unsupportive Topics Proportion across Countries

detailed by (Roberts et al., 2018).

Figure 8 displays the estimates of the proportion of supportive topics (Topics 1, 3, and 8) across the five countries, including 90% confidence intervals. Topics 1 and 8 are notably prominent in Germany, suggesting a strong narrative of interfaith solidarity and resistance to xenophobia. This aligns with the LLM findings, which identified Germany as having the highest percentage of supportive media coverage of immigrants and Muslims, correlating with improved public attitudes toward these groups. Topic 3 appears most frequently in France, indicating a robust emphasis on republican values and democratic ideals after the attack, consistent with improved public attitudes toward immigrants and Muslims there. Topic 3 is also significant in Ireland; however, no improvement in public attitudes toward immigrants and Muslims is observed, likely due to the presence of Topic 5 (illustrated in Figure 9), which is neutral but slightly leaning toward unsupportive. The coexistence of supportive and unsupportive topics in Ireland could explain the lack of observed improvement in public attitudes.

Figure 9 illustrates the estimates of the proportion of unsupportive topics (Topics 4 and 5) across the five countries, complete with 90% confidence intervals. Topic 4 is predominantly high in the Czech Republic and also notable in Finland, indicating a national discourse that is critically unsupportive of Muslims. This correlates with the negative shift in public attitudes in these countries and aligns with the LLM results, which show high percentages of non-supportive media coverage. Topic 5 appears in almost all countries, including Germany and France. These findings confirm the alignment between the STM and LLM analyses, offering a coherent picture of how media coverage influences public attitudes in different national contexts.

5 Conclusion

In this article, we exploited the timing of the seventh round of the European Social Survey (ESS) and the Charlie Hebdo shooting to study the causal link between the attack and attitudes toward immigrants and Muslims. Our study combines observational data from ESS, collected immediately before and after the Charlie Hebdo attack, with a natural experiment design to assign respondents to treatment (post-attack) and control (pre-attack) groups. We employed several causal inference methods, including OLS models with country fixed effects, various matching and weighting techniques, and a machine learning approach.

Our main findings are as follows. First, we identified a negative average impact of the attack on European attitudes toward immigrants and Muslims. This finding is robust across different methodological and time window adjustments. Second, our regression discontinuity design provides clear evidence that the shift in attitude occurred immediately after the attack on January 8, 2015, and is specific to immigrants and Muslims, with no shift in attitudes detected toward Jews or Gypsies. Third, the further away respondents were interviewed from the date of the attack, the smaller the magnitude of the impact, suggesting that the effects of the attack dissipate within a few months.

While the average reaction was more negative, there were significant variations across countries. The Czech Republic, Finland, and Ireland exhibited negative shifts, while France and Germany showed improvements. This raises the question of why reactions to the same terrorist attack differ so significantly from one country to another.

Traditional theoretical frameworks, such as System Justification Theory (SJT) and Intergroup Conflict Theory (ICT), suggest that perceptions of safety threats or competition over scarce resources could explain negative shifts in attitudes. However, these theories alone do not account for the nuanced differences observed across nations in our study. To delve deeper into these attitudinal shifts, we examined the role of the media. Using news articles published about the Charlie Hebdo shooting in the month following the incident in five European countries, we employed two natural language processing (NLP) methods: Large Language Models (LLM) and Structural Topic Modeling (STM). We tested the hypothesis that both the visibility of outgroups and the stance of the news toward them influence public reactions to the incident.

Our descriptive results indicate that anti-Muslim sentiment is not solely a direct artifact of terrorist attacks; rather, it is significantly influenced by the media coverage that follows these events. While terrorist attacks create a context for potential shifts in attitudes, the nature of media coverage—supportive or unsupportive—toward immigrants and Muslims plays a critical role in moderating or exacerbating the effect of these attacks on public attitudes. Specifically, our findings suggest that supportive media coverage can mitigate the negative impact of terrorism on attitudes toward Muslims, while unsupportive or divisive media coverage can amplify negative sentiments.

For example, our results revealed that Germany displayed a substantial proportion of supportive media coverage, focusing on collective interfaith efforts to condemn violence and support for diversity, which aligned with more positive public attitudes toward immigrants and Muslims. In contrast, France exhibited improved public attitudes despite featuring both lower levels of supportive and unsupportive media content. This suggests that the media coverage in France effectively emphasized the country's republican values and democratic ideals in response to the attack. Conversely, the Czech Republic, Finland, and Ireland, which presented higher levels of unsupportive media content focusing on divisive rhetoric, experienced a correspondingly negative shift in public attitudes. This pattern indicates that unsupportive content typically exacerbates negative perceptions towards immigrants and Muslims, while supportive and nuanced media approaches mitigate such effects.

This study contributes significantly to our understanding of the media's influence on public opinion in the aftermath of terrorism, particularly regarding directly involved outgroups. While predicting the impact of specific terror incidents on attitudes towards immigrants is difficult, we have shown that considering the political climate in which the attacks occurred offers a starting point for understanding the conditions under which terrorist attacks lead to shifts in attitude toward outgroups. We encourage further studies in understanding under which conditions terrorism leads to outgroup hostility. Our findings also demonstrate the importance of knowing about the actual contents of the news to understand how it might affect readers. Additionally, our study showcases the utility of LLMs for quantifying opinions across countries and demonstrates how STM can analyze topic prevalence across different national contexts. Combining these methods provides deeper insights into the complex dynamics of media influence.

Our analysis does have limitations. First, we cannot establish a causal link between media coverage and attitude shifts, as only five countries were analyzed. Future research could expand this scope to include more countries or different contexts, potentially using variations in media exposure to establish causality. Additionally, while we focused on direct measures of attitude towards immigrants, future studies might explore indirect measures, such as voting patterns for anti-immigrant parties, to assess whether terrorist attacks influence election outcomes. Another limitation is our inability to investigate the long-term effect of the attack on attitudes due to the lack of data beyond two months after the attack. Understanding the exact duration of the effect is crucial from a policy perspective. Complementing short-term reactions with long-term consequences would provide a more comprehensive picture of how terrorist events impact citizens.

References

Rachith Aiyappa, Jisun An, Haewoon Kwak, and Yong-Yeol Ahn. Can we trust the evaluation on chatgpt? *arXiv preprint arXiv:2303.12767*, 2023.

Gordon Willard Allport, Kenneth Clark, and Thomas Pettigrew. The nature of prejudice. 1954.

Henrik Andersen and Jochen Mayerl. Attitudes towards muslims and fear of terrorism. *Ethnic and Racial Studies*, 41(15):2634–2655, 2018.

Anna Baiardi and Andrea A Naghi. The value added of machine learning to causal inference: Evidence from revisited studies. *arXiv preprint arXiv:2101.00878*, 2021.

Hajo G Boomgaarden and Claes H De Vreese. Dramatic real-world events and public opinion dynamics: Media coverage and its impact on public reactions to an assassination. *International Journal of Public Opinion Research*, 19(3):354–366, 2007.

Bjorn Burscher, Joost van Spanje, and Claes H de Vreese. Owning the issues of crime and immigration: The relation between immigration and crime news and anti-immigrant voting in 11 countries. *Electoral studies*, 38:59–69, 2015.

Bruno Castanho Silva. The (non) impact of the 2015 paris terrorist attacks on political attitudes. *Personality and social psychology bulletin*, 44(6):838–850, 2018.

Nicholas C Chesnaye, Vianda S Stel, Giovanni Tripepi, Friedo W Dekker, Edouard L Fu, Carmine Zoccali, and Kitty J Jager. An introduction to inverse probability of treatment weighting in observational research. *Clinical Kidney Journal*, 15(1):14–20, 2022.

Mathieu Couttenier, Sophie Hatte, Mathias Thoenig, and Stephanos Vlachos. Anti-muslim voting and media coverage of immigrant crimes. *The Review of Economics and Statistics*, pages 1–33, 2021.

Iain J Cruickshank and Lynnette Hui Xian Ng. Use of large language models for stance classification. *arXiv preprint arXiv:2309.13734*, 2023.

Enny Das, Brad J Bushman, Marieke D Bezemer, Peter Kerkhof, and Ivar E Vermeulen. How terrorism news reports increase prejudice against outgroups: A terror management account. *Journal of Experimental Social Psychology*, 45(3):453–459, 2009.

David De Coninck. Fear of terrorism and attitudes toward refugees: An empirical test of group threat theory. *Crime & Delinquency*, 68(4):550–571, 2022.

Milena Djourelova. Media persuasion through slanted language: Evidence from the coverage of immigration. 2020.

Agustin Echebarria-Echabe and Emilia Fernández-Guede. Effects of terrorism on attitudes and ideological orientation. *European Journal of Social Psychology*, 36(2): 259–265, 2006.

Ryan D Enos. Causal effect of intergroup contact on exclusionary attitudes. *Proceedings of the National Academy of Sciences*, 111(10):3699–3704, 2014.

Lukas Erhard, Raphael H Heiberger, and Michael Windzio. Diverse effects of mass media on concerns about immigration: New evidence from germany, 2001–2016. *European Sociological Review*, 38(4):629–647, 2022.

Jens Hainmueller. Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies. *Political analysis*, 20 (1):25–46, 2012.

Siqing Huo, Negar Arabzadeh, and Charles LA Clarke. Retrieving supporting evidence for llms generated answers. *arXiv preprint arXiv:2306.13781*, 2023.

Niklas Jakobsson and Svein Blom. Did the 2011 terror attacks in norway change citizens' attitudes toward immigrants? *International Journal of Public Opinion Research*, 26 (4):475–486, 2014.

John T Jost and Mahzarin R Banaji. The role of stereotyping in system-justification and the production of false consciousness. *British journal of social psychology*, 33(1): 1–27, 1994.

Sekou Keita, Thomas Renault, and Jérôme Valette. The usual suspects: Offender origin, media reporting and natives' attitudes towards immigration. *Documents de travail du Centre d'Économie de la Sorbonne*, 2021.

Dilek Küçük and Fazli Can. Stance detection: A survey. *ACM Computing Surveys (CSUR)*, 53(1):1–37, 2020.

Joscha Legewie. Terrorist events and attitudes toward immigrants: A natural experiment. *American journal of sociology*, 118(5):1199–1245, 2013.

Moreno Mancosu and Monica Ferrin Pereira. Terrorist attacks, stereotyping, and attitudes toward immigrants: The case of the manchester bombing. *Social science quarterly*, 102(1):420–432, 2021.

McCombs Maxwell. Setting the agenda: the mass media and public opinion, 2004.

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

University of Oklahoma. Institute of Group Relations and Muzafer Sherif. *Intergroup conflict and cooperation: The Robbers Cave experiment*, volume 10. University Book Exchange Norman, OK, 1961.

Thomas F Pettigrew and Linda R Tropp. A meta-analytic test of intergroup contact theory. *Journal of personality and social psychology*, 90(5):751, 2006.

Faisal Rabby and William M Rodgers iii. The impact of 9/11 and the london bombings on the employment and earnings of uk muslims. 2010.

Margaret Roberts, Brandon Stewart, Dustin Tingley, Kenneth Benoit, Maintainer Brandon Stewart, LinkingTo Rcpp, et al. Package 'stm'. *R Package Version*, 1(3):3, 2018.

Margaret E Roberts, Brandon M Stewart, Dustin Tingley, Edoardo M Airoldi, et al. The structural topic model and applied social science. In *Advances in neural information processing systems workshop on topic models: computation, application, and evaluation*, volume 4, pages 1–20. Harrahs and Harveys, Lake Tahoe, 2013.

Margaret E Roberts, Brandon M Stewart, Dustin Tingley, Christopher Lucas, Jetson Leder-Luis, Shana Kushner Gadarian, Bethany Albertson, and David G Rand. Structural topic models for open-ended survey responses. *American journal of political science*, 58(4):1064–1082, 2014.

James M Robins, Andrea Rotnitzky, and Lue Ping Zhao. Estimation of regression coefficients when some regressors are not always observed. *Journal of the American statistical Association*, 89(427):846–866, 1994.

Frank Rickard Sandell, Christofer Edling, and Jens Rydgren. Terrorism, belief formation, and residential integration: population dynamics in the aftermath of the 2004 madrid terror bombings. 2016.

Michael Savelkoul, Manfred te Grotenhuis, and Peer Scheepers. Has the terrorist attack on charlie hebdo fuelled resistance towards muslim immigrants in europe? results from a natural experiment in six european countries. *Acta Sociologica*, 65(4):357–373, 2022.

Douglas C Schmidt, Jesse Spencer-Smith, Quchen Fu, and Jules White. Cataloging prompt patterns to enhance the discipline of prompt engineering. *URL: https://www. dre. vanderbilt. edu/˜ schmidt/PDF/ADA_Europe_Position_Paper. pdf [accessed 2023-09-25]*, 2023.

Paul M Sniderman, Louk Hagendoorn, and Markus Prior. Predisposing factors and situational triggers: Exclusionary reactions to immigrant minorities. *American political science review*, 98(1):35–49, 2004.

Jasper Van Assche and Kim Dierckx. Attitudes towards outgroups before and after terror attacks. *Terrorism and political violence*, 33(7):1530–1545, 2021.

Julie Van de Vyver, Diane M Houston, Dominic Abrams, and Milica Vasiljevic. Boosting belligerence: How the july 7, 2005, london bombings affected liberals' moral foundations and prejudice. *Psychological science*, 27(2):169–177, 2016.

Marijn Van Klingeren, Hajo G Boomgaarden, Rens Vliegenthart, and Claes H De Vreese. Real world is not enough: The media as an additional source of negative attitudes toward immigration, comparing denmark and the netherlands. *European Sociological Review*, 31(3):268–283, 2015.

Rens Vliegenthart and Hajo G Boomgaarden. Real-world indicators and the coverage of immigration and the integration of minorities in dutch newspapers. *European Journal of Communication*, 22(3):293–314, 2007.

Christian Von Sikorski, Desirée Schmuck, Jörg Matthes, and Alice Binder. "muslims are not terrorists": Islamic state coverage, journalistic differentiation between terrorism and islam, fear reactions, and attitudes toward muslims. In *Media, Terrorism and Society*, pages 91–114. Routledge, 2020.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837, 2022.

Jules White, Quchen Fu, Sam Hays, Michael Sandborn, Carlos Olea, Henry Gilbert, Ashraf Elnashar, Jesse Spencer-Smith, and Douglas C Schmidt. A prompt pattern catalog to enhance prompt engineering with chatgpt. *arXiv preprint arXiv:2302.11382*, 2023.

Bowen Zhang, Daijun Ding, and Liwen Jing. How would stance detection techniques evolve after the launch of chatgpt? *arXiv preprint arXiv:2212.14548*, 2022.

Wenxuan Zhang, Yue Deng, Bing Liu, Sinno Jialin Pan, and Lidong Bing. Sentiment analysis in the era of large language models: A reality check. *arXiv preprint arXiv:2305.15005*, 2023.

Yaqian Zhu, Rebecca A Hubbard, Jessica Chubak, Jason Roy, and Nandita Mitra. Core concepts in pharmacoepidemiology: Violations of the positivity assumption in the causal analysis of observational data: Consequences and statistical approaches. *Pharmacoepidemiology and drug safety*, 30(11):1471–1485, 2021.

Appendix A Media coverage: Charlie-Hebdo

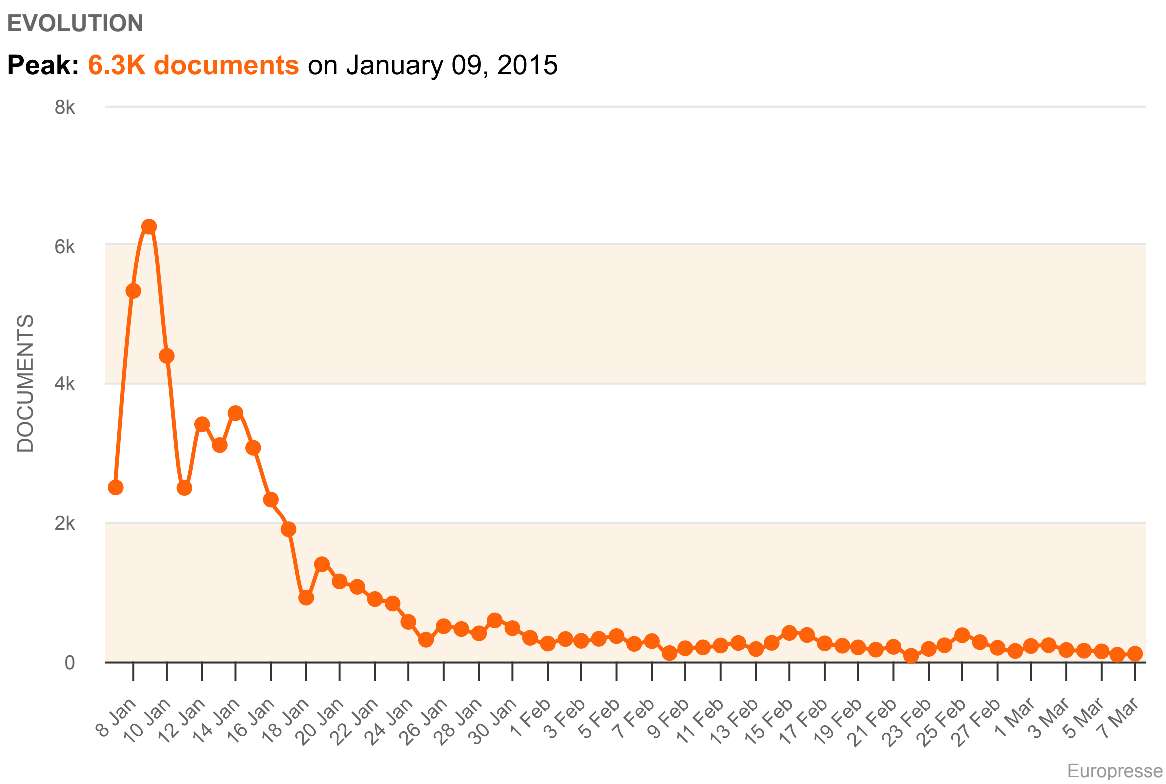

Figure 10: Media Coverage: Charlie-Hebdo

Appendix B Difference in SMD after adjusting

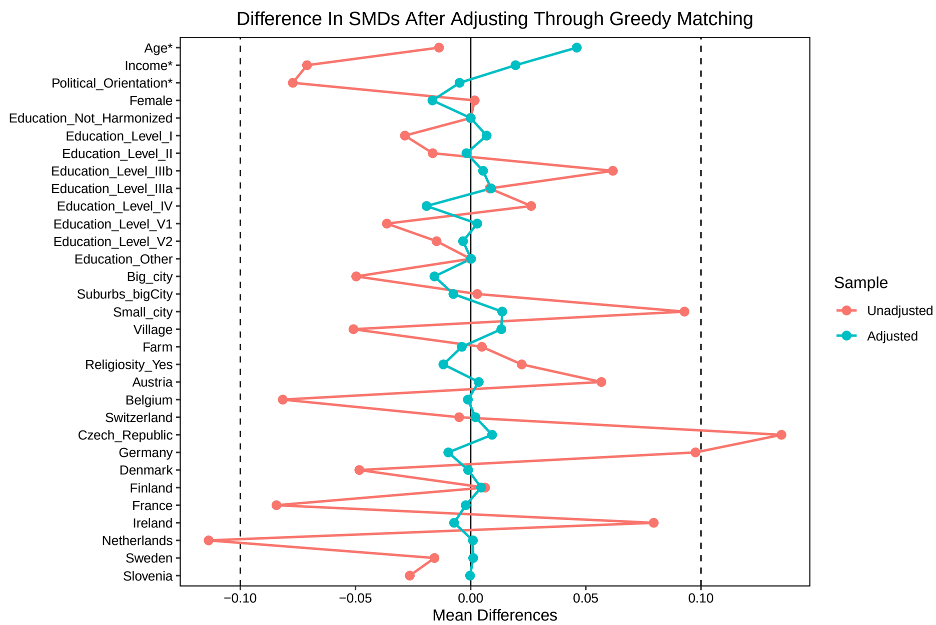

Figure 11: Difference In SMDs After Adjusting Through Greedy Matching

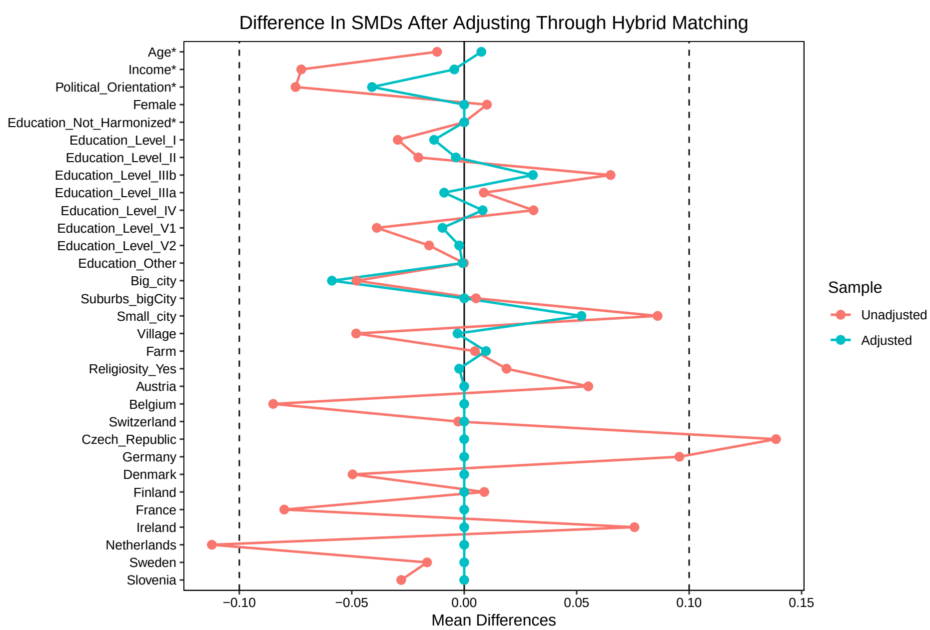

Figure 12: Difference In SMDs After Adjusting Through Hybrid Matching

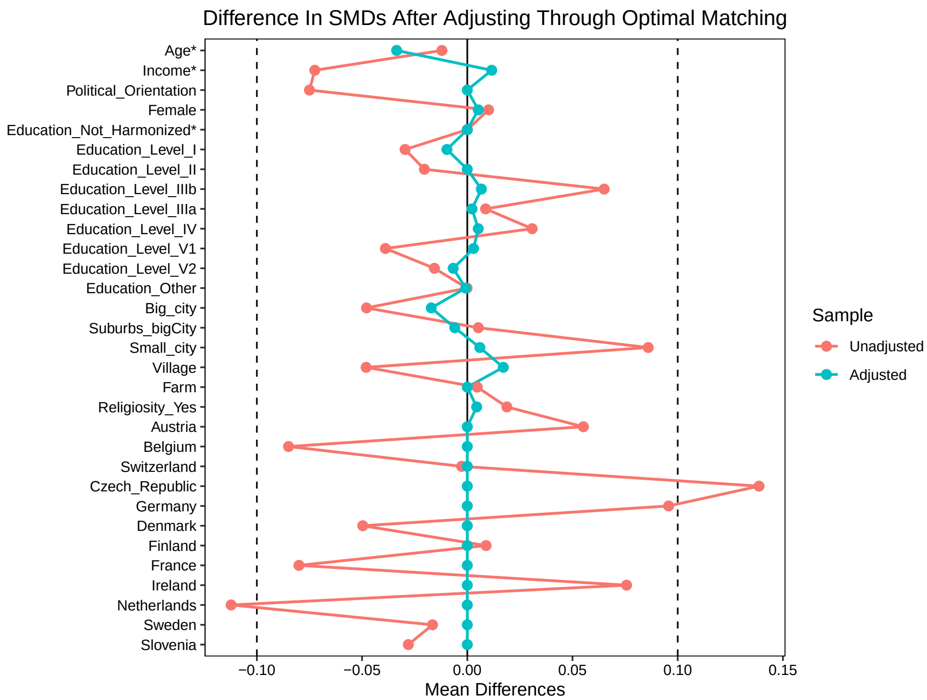

Figure 13: Difference In SMDs After Adjusting Through Optimal Matching

Appendix C Robustness Checks

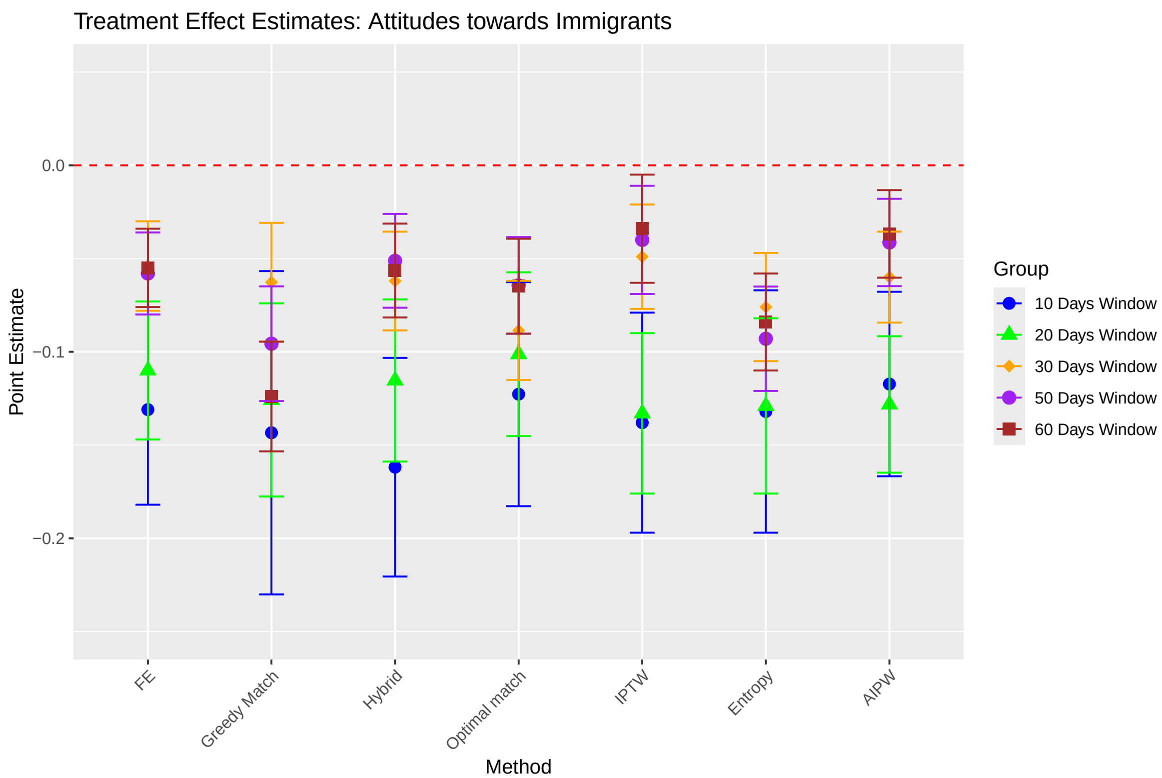

Figure 14: Attitudes towards Immigrants: Treatment Effect

Note:

- This figure represents the Average Treatment Effects (ATE) of the Charlie Hebdo shooting on attitudes towards immigrants for different bandwidth choices.

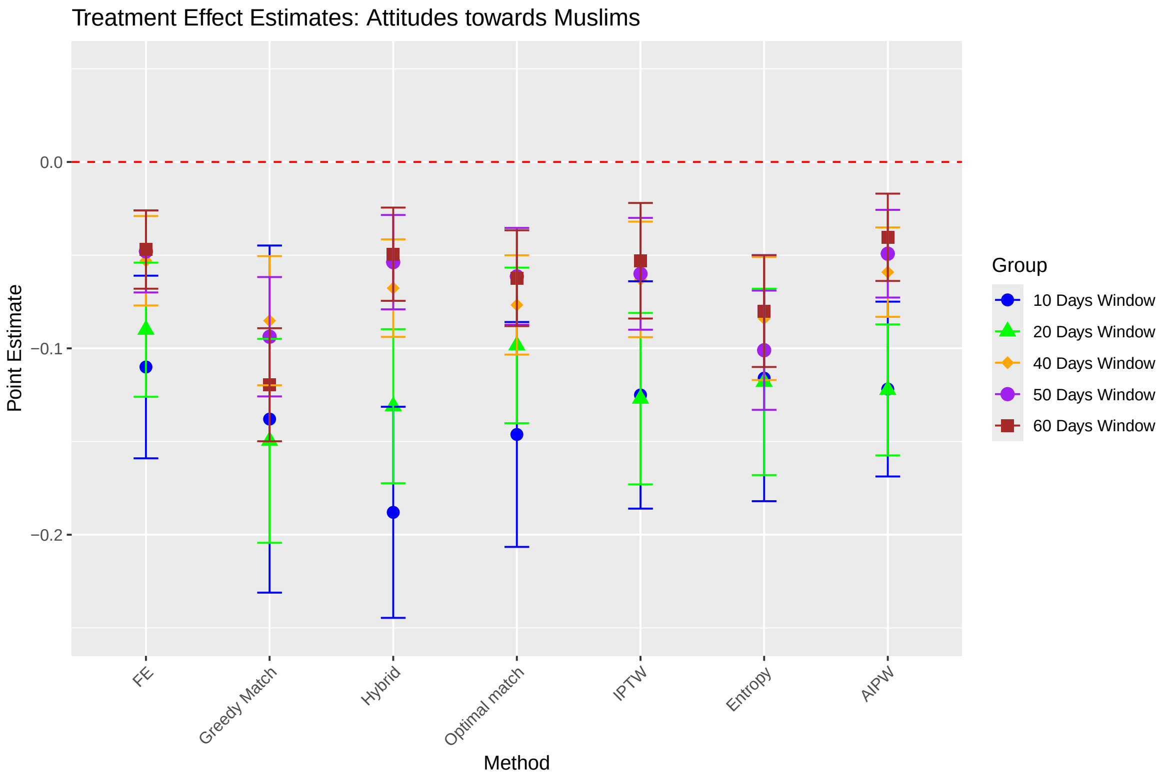

Figure 15: Attitudes towards Immigrants: Treatment Effect

Note:

- This figure represents the Average Treatment Effects (ATE) of the Charlie Hebdo shooting on attitudes towards immigrants for different bandwidth choices.

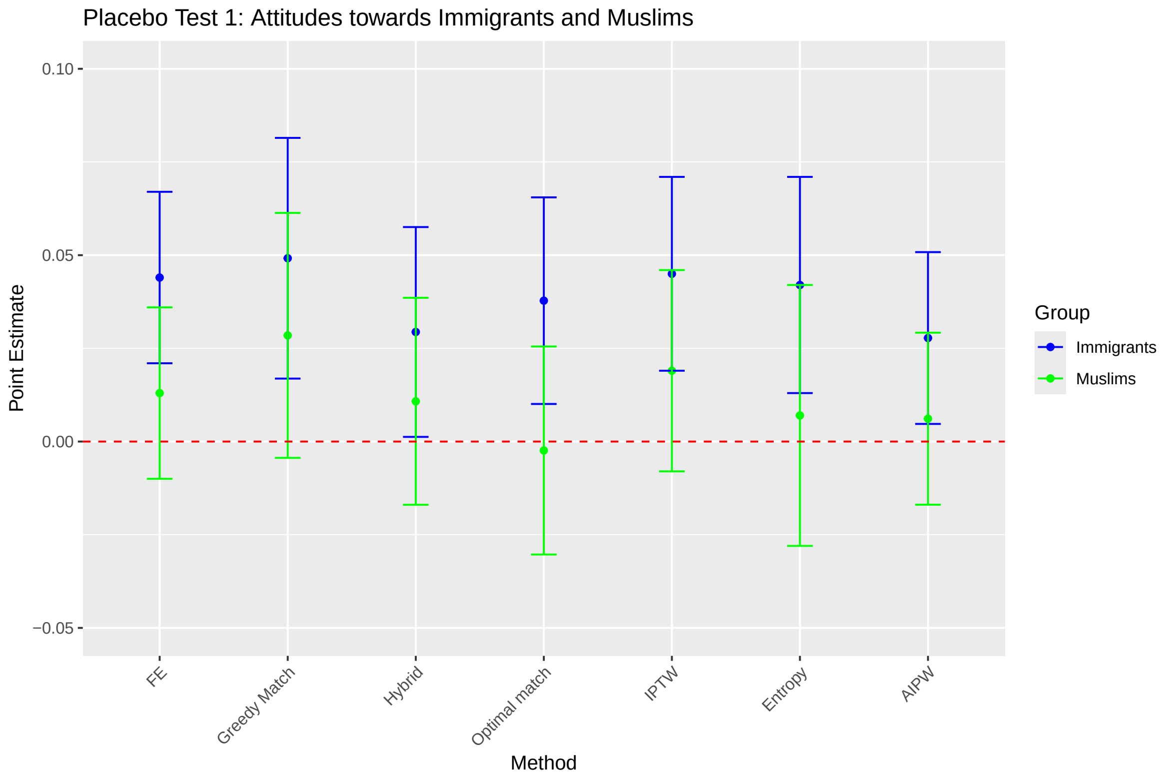

Figure 16: Attitudes towards Immigrants and Muslims: Placebo 1 (Fake attack date, 2014-12-07)

Note:

- This figure represents the Average Treatment Effects (ATE) of the Charlie Hebdo shooting on attitudes towards immigrants and Muslims when we falsely assume that the attack happened one month before the actual date. Fixed effects estimator (FE) used 6,283 observations. All matching methods point estimate derived from a paired t-test after matching. Greedy Match: 2,063 matched out of 2,063 treated; Hybrid method combines Exact and Greedy matching: 1,667 matched out of 2,063 treated; Optimal Match: 1,667 matched out of 2,063 treated; IPTW (Inverse Probability Weighting): Utilized 6,193 unweighted observations. Entropy Balancing: Utilized 6,193 unweighted observations. AIPW (Augmented Inverse Propensity-Weighted): Estimation involved 6,167 observations.

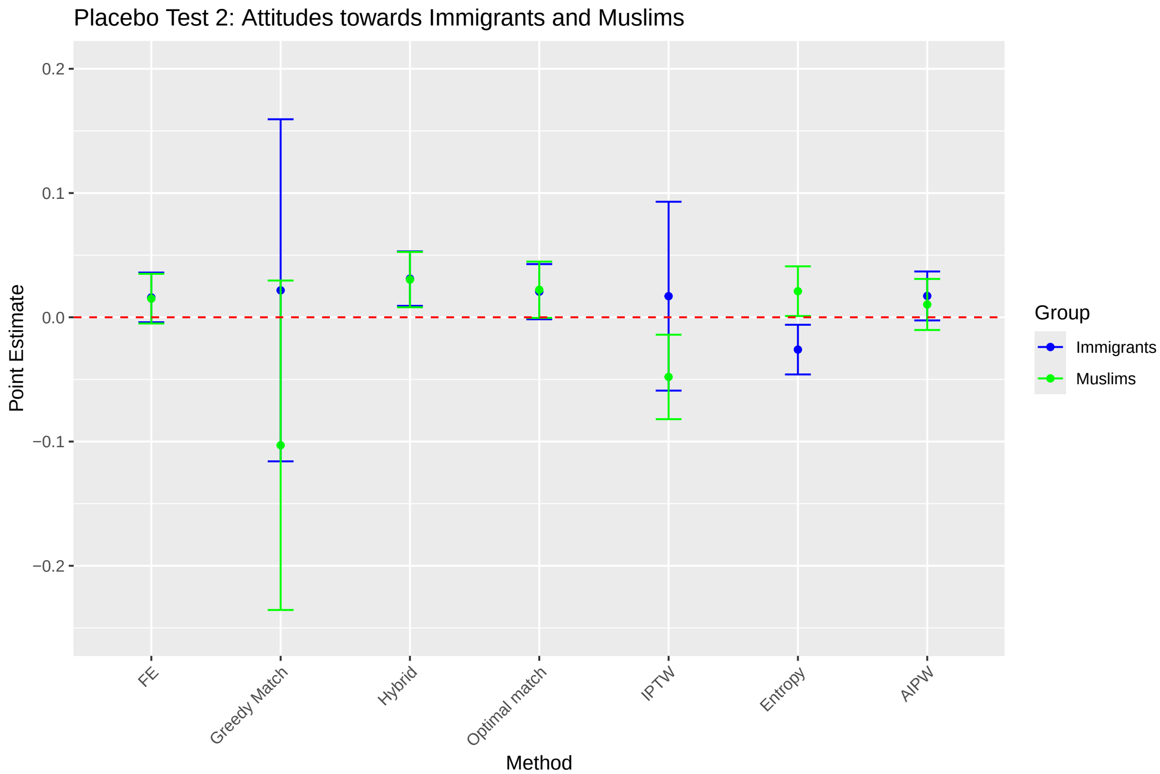

Figure 17: Attitudes towards Immigrants and Muslims: Placebo 2 (Fake attack date, 2014-11-07)

Note:

- This figure represents the Average Treatment Effects (ATE) of the Charlie Hebdo shooting on attitudes towards immigrants and Muslims when we falsely assume that the attack happened two months before the actual date. Fixed effects estimator (FE) used 8,366 observations. All matching methods point estimate derived from a paired t-test after matching. Greedy Match: 4,092 matched out of 4,092 treated; Hybrid method combines Exact and Greedy matching: 2,571 matched out of 4,092 treated; Optimal Match: 2,571 matched out of 4,092 treated; IPTW (Inverse Probability Weighting): Utilized 8,256 unweighted observations. Entropy Balancing Utilized 8,256 unweighted observations. AIPW (Augmented Inverse Propensity-Weighted): Estimation involved 8,230 observations.

Table 3: Multilevel Linear Regression With Random Intercepts

	Immigrants		Muslims	
Predictors	**Estimates**	**CI**	**Estimates**	**CI**
(Intercept)	2.80***	2.59 – 3.02	2.73***	2.46 – 3.00
Treatment	-0.08**	-0.13 – -0.03	-0.07*	-0.12 – -0.02
Male	-0.02	-0.07 – 0.03	0.06*	0.01 – 0.11
Age	-0.01***	-0.01 – -0.00	-0.01***	-0.01 – -0.00
Education_Level_II	0.09	-0.03 – 0.21	0.15*	0.02 – 0.27
Education_Level_IIIb	0.16**	0.04 – 0.28	0.20***	0.08 – 0.32
Education_Level_IIIa	0.26***	0.14 – 0.39	0.27***	0.15 – 0.40
Education_Level_IV	0.32***	0.20 – 0.45	0.37***	0.25 – 0.50
Education_Level_V1	0.52***	0.39 – 0.66	0.60***	0.47 – 0.74
Education_Level_V2	0.59***	0.46 – 0.72	0.63***	0.50 – 0.76
Education_Other	-0.26	-1.06 – 0.53	0.60	-0.19 – 1.39
Income	0.01	-0.00 – 0.02	0.01	-0.00 – 0.02
Political_Orientation	-0.06***	-0.07 – -0.04	-0.04***	-0.05 – -0.03
Suburbs_bigCity	0.06	-0.03 – 0.15	-0.00	-0.09 – 0.09
Small_city	-0.07*	-0.14 – -0.01	-0.14***	-0.20 – -0.07
Village	-0.03	-0.10 – 0.04	-0.09**	-0.16 – -0.02
Farm	-0.04	-0.15 – 0.07	-0.16**	-0.27 – -0.05
Religious_No	0.02	-0.04 – 0.07	0.01	-0.04 – 0.06
Random Effects				
σ^2	0.64		0.63	
τ_{00}	0.06 (Country)		0.13 (Country)	
ICC	0.09		0.18	
N	12 (Country)		12 (Country)	
Observations	4479		4453	
Marginal R^2 / Conditional R^2	0.088 / 0.165		0.087 / 0.247	

Note: * p<0.05 ** p<0.01 *** p<0.001

Table 4: Leave One Country Out Analysis: Attitude toward Immigrants

	Dependent variable:											
	Attitude Toward Immigrants											
	(CZ)	(FR)	(DE)	(BE)	(NL)	(CH)	(IE)	(AT)	(FI)	(DK)	(SE)	(SI)
Treatment	-0.076^{**}	-0.110^{***}	-0.078^{***}	-0.073^{***}	-0.073^{***}	-0.069^{**}	-0.071^{**}	-0.076^{***}	-0.072^{***}	-0.086^{***}	-0.076^{***}	-0.082^{***}
	(0.032)	(0.029)	(0.029)	(0.027)	(0.027)	(0.027)	(0.028)	(0.027)	(0.028)	(0.027)	(0.027)	(0.027)
Constant	3.060^{***}	2.865^{***}	2.869^{***}	2.858^{***}	2.898^{***}	2.858^{***}	2.838^{***}	2.720^{***}	2.855^{***}	2.875^{***}	2.873^{***}	2.880^{***}
	(0.109)	(0.111)	(0.107)	(0.104)	(0.105)	(0.103)	(0.108)	(0.098)	(0.104)	(0.103)	(0.103)	(0.103)
Observations	3,274	3,856	3,765	4,208	4,176	4,381	3,949	4,294	4,223	4,339	4,428	4,376
R^2	0.185	0.184	0.161	0.191	0.193	0.183	0.196	0.185	0.187	0.186	0.178	0.186
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Note: *p<0.1; **p<0.05; ***p<0.01

Between Parentheses is the code of the country left out

Table 5: Leave One Country Out Analysis: Attitude toward Muslims

	Dependent variable:											
	Attitude Toward Muslims											
	(CZ)	(FR)	(DE)	(BE)	(NL)	(CH)	(IE)	(AT)	(FI)	(DK)	(SE)	(SI)
Treatment	-0.059^{*}	-0.099^{***}	-0.094^{***}	-0.055^{**}	-0.062^{**}	-0.065^{**}	-0.044	-0.062^{**}	-0.067^{**}	-0.069^{**}	-0.064^{**}	-0.065^{**}
	(0.032)	(0.029)	(0.029)	(0.027)	(0.027)	(0.027)	(0.028)	(0.027)	(0.027)	(0.027)	(0.027)	(0.027)
Constant	2.964^{***}	2.771^{***}	2.833^{***}	2.803^{***}	2.856^{***}	2.811^{***}	2.727^{***}	2.459^{***}	2.762^{***}	2.796^{***}	2.792^{***}	2.797^{***}
	(0.111)	(0.111)	(0.107)	(0.105)	(0.105)	(0.103)	(0.107)	(0.097)	(0.103)	(0.103)	(0.102)	(0.102)
Observations	3,242	3,840	3,746	4,182	4,150	4,360	3,930	4,270	4,198	4,313	4,401	4,351
R^2	0.214	0.356	0.326	0.363	0.367	0.354	0.386	0.360	0.365	0.354	0.352	0.356
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Note: *p<0.1; **p<0.05; ***p<0.01

Between Parentheses is the code of the country left out

Appendix D Heterogeneous Findings

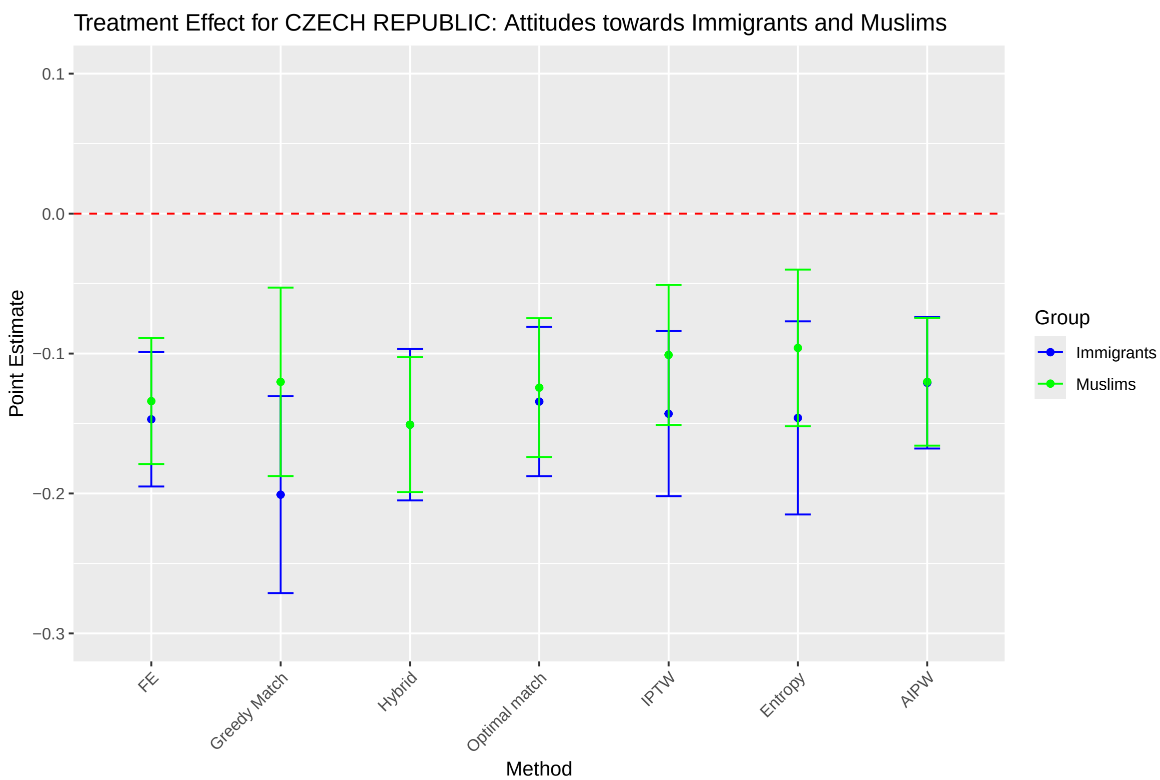

Figure 18: Attitudes towards Immigrants and Muslims: Treatment Effect Estimates for Czech Republic

Note:

- This figure represents the Average Treatment Effects (ATE) of the Charlie Hebdo shooting on attitudes towards immigrants and Muslims in CZECH REPUBLIC. Fixed effects estimator (FE) used 1,205 observations. All matching methods point estimate derived from a paired t-test after matching. Greedy Match: 763 matched out of 763 treated; Hybrid method combines Exact and Greedy matching: 411 matched out of 763 treated; Optimal Match: 402 matched out of 763 treated; IPTW (Inverse Probability Weighting): Utilized 1,190 unweighted observations. Entropy Balancing: Utilized 1,190 unweighted observations. AIPW (Augmented Inverse Propensity-Weighted): Estimation involved 1,172 observations.

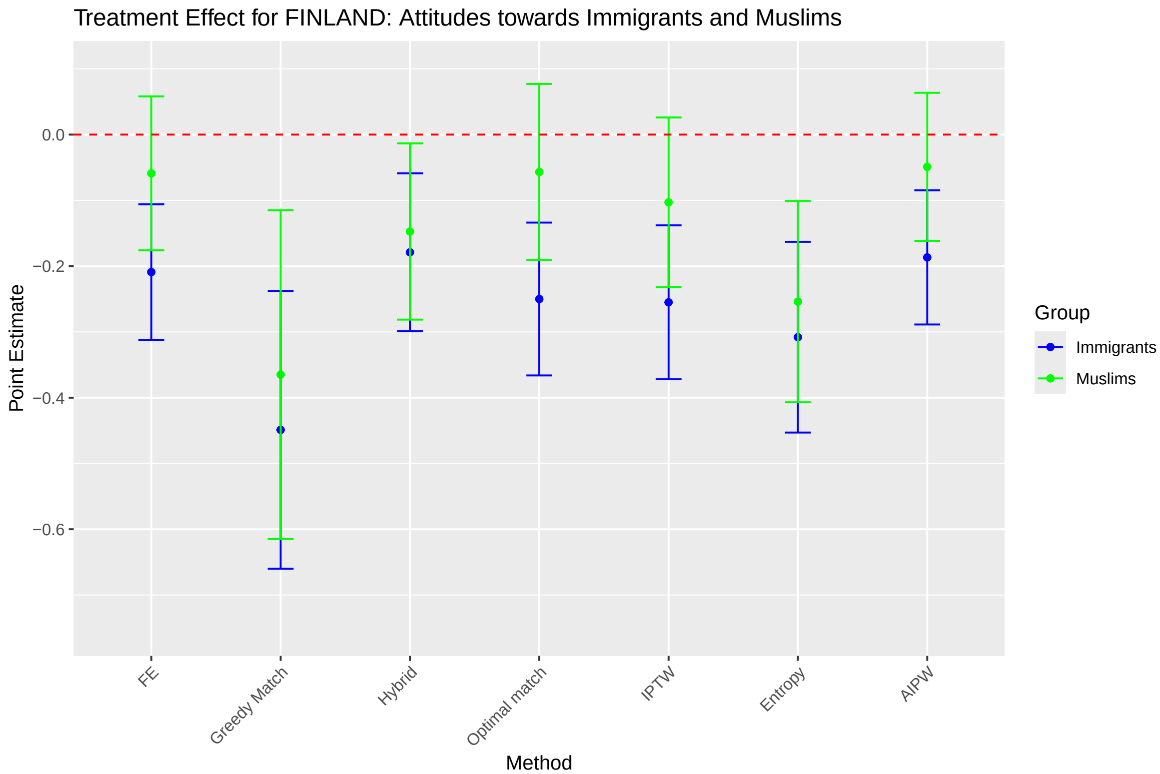

Figure 19: Attitudes towards Immigrants and Muslims: Treatment Effect Estimates for Finland

Note:

- This figure represents the Average Treatment Effects (ATE) of the Charlie Hebdo shooting on attitudes towards immigrants in Finland. Fixed effects estimator (FE) used 256 observations. All matching methods point estimate derived from a paired t-test after matching. Greedy Match: 137 matched out of 137 treated; Hybrid method combines Exact and Greedy matching: 95 matched out of 137 treated; Optimal Match: 88 matched out of 137 treated; IPTW (Inverse Probability Weighting): Utilized 252 unweighted observations. Entropy Balancing Utilized 252 unweighted observations. AIPW (Augmented Inverse Propensity-Weighted): Estimation involved 252 observations.

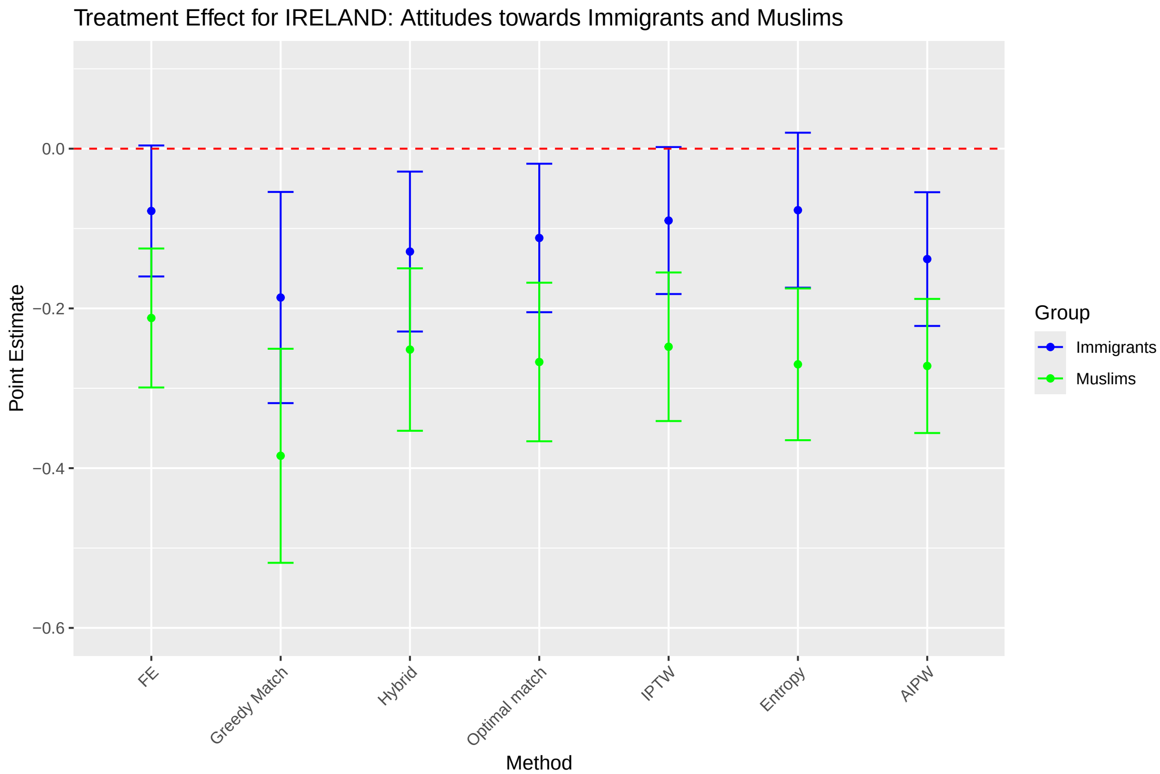

Figure 20: Attitudes towards Immigrants and Muslims: Treatment Effect Estimates for Ireland

Note:

- This figure represents the Average Treatment Effects (ATE) of the Charlie Hebdo shooting on attitudes towards immigrants in Ireland. Fixed effects estimator (FE) used 530 observations. All matching methods point estimate derived from a paired t-test after matching. Greedy Match: 355 matched out of 355 treated; Hybrid method combines Exact and Greedy matching: 163 matched out of 355 treated; Optimal Match: 161 matched out of 355 treated; IPTW (Inverse Probability Weighting): Utilized 518 unweighted observations. Entropy Balancing Utilized 518 unweighted observations. AIPW (Augmented Inverse Propensity-Weighted): Estimation involved 514 observations.

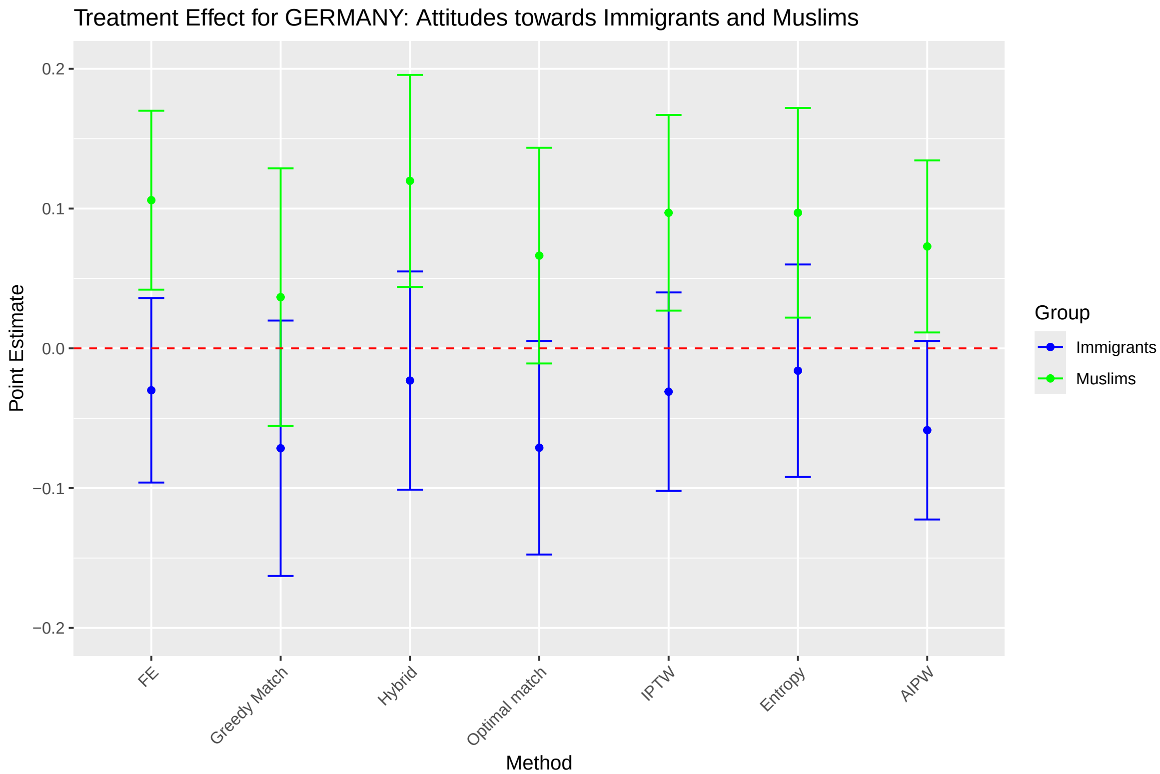

Figure 21: Attitudes towards Immigrants and Muslims: Treatment Effect Estimates for Germany

Note:

- This figure represents the Average Treatment Effects (ATE) of the Charlie Hebdo shooting on attitudes towards immigrants and Muslims in Germany. Fixed effects estimator (FE) used 714 observations. All matching methods point estimate derived from a paired t-test after matching. Greedy Match: 471 matched out of 471 treated; Hybrid method combines Exact and Greedy matching: 217 matched out of 471 treated; Optimal Match: 211 matched out of 471 treated; IPTW (Inverse Probability Weighting): Utilized 704 unweighted observations. Entropy Balancing: Utilized 704 unweighted observations. AIPW (Augmented Inverse Propensity-Weighted): Estimation involved 703 observations.

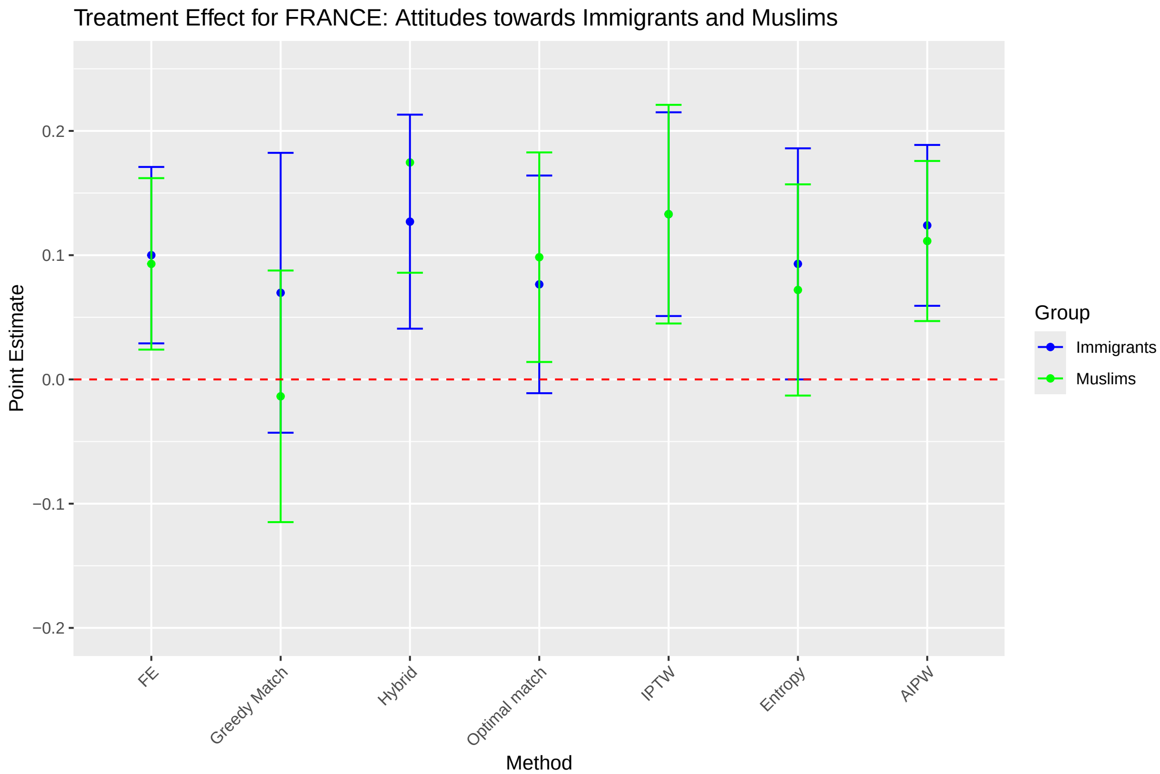

Figure 22: Attitudes towards Immigrants and Muslims: Treatment Effect Estimates for France

Note:

- This figure represents the Average Treatment Effects (ATE) of the Charlie Hebdo shooting on attitudes towards immigrants and Muslims in France. Fixed effects estimator (FE) used 623 observations. All matching methods point estimate derived from a paired t-test after matching. Greedy Match: 223 matched out of 223 treated; Hybrid method combines Exact and Greedy matching: 189 matched out of 223 treated; Optimal Match: 183 matched out of 223 treated; IPTW (Inverse Probability Weighting): Utilized 612 unweighted observations. Entropy Balancing Utilized 612 unweighted observations. AIPW (Augmented Inverse Propensity-Weighted): Estimation involved 610 observations.

Appendix E Prompt used for the LLM

Appendix F Exemplar Documents

Prompt: Please answer the following questions to assess the article's stance towards Islam/Muslims in particular and immigrants in general. You must answer all questions:

1. Does any part of the news article below discuss Islam/Muslims or immigrants? Provide a 'Yes' or 'No' answer and include specific evidence from the article.
2. Is any opinion presented in the article that can be described as supportive of Muslims or immigrants? Provide a 'Yes' or 'No' answer and explain why based on the opinions expressed in the article.
3. Is any opinion presented in the article that can be described as NOT supportive of Muslims or immigrants? Provide a 'Yes' or 'No' answer and explain why based on the opinions expressed in the article.

Article: "*Terror in France: "It is not possible for one to hold a pencil and the other a kalashnikov," says the imam of the mosque in Drancy. French Muslims are indignant at the barbaric act of attackers. They are also afraid of reprisals. DRANCY A tape dammed the road to the brown square mosque building on the outskirts of the northern Paris suburb of Drancy, with three police cars standing in the adjacent car park. The entrance to the building is protected by several armed men who are tasked with preventing possible violent attacks that several mosques across the country faced after Wednesday's assassination of Charlie Hebdo magazine. The mosque is not closed, and for Friday prayer, worshippers come as usual. "The faithful are scared, they are shocked. But they are also resentful of those barbarians brandishing religion. They are afraid of more violence, some mosques were attacked yesterday," says Hassen Chalghoumi, Imam Hassen Chalghoumi, who said violence only breeds more violence. Mr Chalghoumi is preparing for a large demonstration on Sunday, which is said to be attended by a number of leaders of Paris's Muslim community. "Someone who uses guns to kill people has no pranic to do with Islam," he says convincingly. When asked what freedom of speech means to him, he immediately replies: "It is holy to me. One may disagree with some articles, but that's it. It is impossible for one to have a pencil in his hand and the other a Kalashnikov." ...*

Output:

1. Yes, the news article below discusses Muslims and Islam, as well as immigrants in general. The article mentions French Muslims and their concerns in light of the Charlie Hebdo attacks. It also includes information about a mosque in Drancy, and its imam's views on the attacks and Islam in general. The article also mentions the background of a young man who became radicalized and joined ISIS in Syria. Unemployment in the French suburbs is also mentioned, which is often linked to immigrants.
2. Yes, the article includes an opinion that can be described as supportive of Muslims. The imam in Drancy, Hassen Chalghoumi, speaks out against the attackers and affirms that Islam does not condone violence. He also expresses his support for freedom of speech, but believes it should not lead to violence or disrespect.
3. No, the article does not present any opinion that can be described as not supportive of Muslims or immigrants.

Table 6: An example of the prompting strategy and the output from the model.

Table 7

Topic	Document	Perplexity	Chat GPT	Claude	Overall
1	1	Supportive	Supportive	Supportive	Supportive
1	2	Supportive	Supportive	Supportive	Supportive
1	3	Supportive	Supportive	Supportive	Supportive
1	4	Supportive	Supportive	Supportive	Supportive
1	5	Supportive	Supportive	Supportive	Supportive
2	1	Neutral	Neutral	Neutral	Neutral
2	2	Neutral	Neutral	Neutral	Neutral
2	3	Neutral	Neutral	Neutral	Neutral
2	4	Neutral	Neutral	Neutral	Neutral
2	5	Neutral	Neutral	Neutral	Neutral
3	1	Supportive	Supportive	Supportive	Supportive
3	2	Unsupportive	Supportive	Neutral	N/A
3	3	Neutral	Supportive	Supportive	Supportive
3	4	Supportive	Supportive	Supportive	Supportive
3	5	Neutral	Supportive	Supportive	Supportive
4	1	Neutral	Unsupportive	Neutral	Neutral
4	2	Unsupportive	Unsupportive	Neutral	Unsupportive
4	3	Neutral	Supportive	Neutral	Neutral
4	4	Neutral	Supportive	Neutral	Neutral
4	5	Unsupportive	Unsupportive	Neutral	Unsupportive
5	1	Unsupportive	Unsupportive	Unsupportive	Unsupportive
5	2	Neutral	Unsupportive	Neutral	Neutral
5	3	Unsupportive	Unsupportive	Unsupportive	Unsupportive
5	4	Neutral	Supportive	Neutral	Neutral
5	5	Neutral	Supportive	Neutral	Neutral
6	1	Neutral	Supportive	Supportive	Supportive
6	2	Unsupportive	Unsupportive	Neutral	Unsupportive
6	3	Supportive	Supportive	Supportive	Supportive
6	4	Unsupportive	Supportive	Unsupportive	Unsupportive
6	5	Supportive	Supportive	Supportive	Supportive
7	1	Neutral	Neutral	Unsupportive	Neutral
7	2	Neutral	Neutral	Unsupportive	Neutral
7	3	Unsupportive	Neutral	Unsupportive	Unsupportive
7	4	Neutral	Neutral	Unsupportive	Neutral
7	5	Neutral	Neutral	Unsupportive	Neutral
8	1	Supportive	Neutral	Neutral	Neutral
8	2	Supportive	Supportive	Supportive	Supportive
8	3	Supportive	Supportive	Supportive	Supportive
8	4	Supportive	Supportive	Supportive	Supportive
8	5	Supportive	Supportive	Supportive	Supportive
9	1	Neutral	Supportive	Neutral	Neutral
9	2	Neutral	Neutral	Neutral	Neutral
9	3	Neutral	Neutral	Neutral	Neutral
9	4	Supportive	Supportive	Supportive	Supportive
9	5	Neutral	Neutral	Neutral	Neutral
10	1	Neutral	Neutral	Neutral	Neutral
10	2	Neutral	Neutral	Neutral	Neutral
10	3	Neutral	Neutral	Neutral	Neutral
10	4	Neutral	Neutral	Neutral	Neutral
10	5	Neutral	Neutral	Neutral	Neutral