How Big is the Media Multiplier? Evidence from Dyadic News Data

Timothy Besley Thiemo Fetzer Hannes Mueller*

October 29, 2021

Abstract

This paper estimates the size of the *media multiplier*, a model-based measure of how far media coverage magnifies the response to shocks. We combine monthly aggregated and anonymized card activity data from 114 card issuing countries in Turkey, Egypt, Tunisia, Israel and Morocco with a large corpus of news coverage of violent events at these five destinations. To define and quantify the media multiplier we estimate a model in which latent beliefs, shaped by either events or news coverage, drive the behavior of heterogeneous agents. We show here that this parsimonious model does a remarkably good job at picking up the timing and differential responses to the shock in different contexts. According to the model, card activity falls by 53 percent if a country is regarded as dangerous by all agents - more than half of this effect is due to the media multiplier.

JEL Classification: O1, F5, F1, L8

Keywords: media, economic behavior

---

*Corresponding author: Mueller, h.mueller.uni@gmail.com. Besley is based at the London School of Economics, Fetzer at Warwick University and Mueller is based at IAE (CSIC) and the Barcelona GSE. We thank audiences at Paris School of Economics, CAGE/Warwick University, Banco de Espana, CIREQ in Montreal, Barcelona GSE Summer Forum, London School of Economics, CUNEF Madrid. We are grateful to the International Growth Centre (IGC) and to the Mastercard Center for Inclusive Growth for their support. Mueller acknowledges fiancial support from grant PGC2018-096133-B-100 by MCIU/AEI/FEDER, grant 2017 SGR 1359 by AGAUR-GENERALITAT DE CATALUNYA and the Spanish Ministry of Economy and Competitiveness, through the Severo Ochoa Programme for Centres of Excellence in R&D (SEV-2015-0563) and Grant number ECO2015-66883-P. We thank Chris Dann, Olga Starchikova, Bruno Conte Leite, and Azhar Hussain for excellent research assistance. All errors are our own.

1 Introduction

There is strong evidence that media reporting can have an impact on economic behavior.[1] This has been explored for negative events such as crime, violence, traffic accidents or during the pandemic where risk assessments and responses may be driven by media coverage.[2] But while it is well-established that the "media matters" in a range of contexts, the literature has rarely tried to integrate theory and data in quantitative estimation of the magnitude of this effect.

This paper introduces a model of the *media multiplier*, a measure of how far media coverage magnifies the response to shocks. We estimate the model in the context of sudden outbreaks of violence. Our approach is based on the idea that events and news coverage of them change behavior because they change risk perceptions. We show here that a parsimonious model does a remarkably good job at picking up the timing and differential responses to the shock in different contexts. The parameter estimates of the model suggest that, in the context of terrorism and violence, assessments of risk get amplified by a factor of up to 2.5 with consistent, long-lasting negative news coverage. However, the actual effect will depend on how intensive negative reporting is and whether or not it is *drowned out* by other unrelated news coverage.

The paper studies the impact of violent events in 5 travel destinations (Turkey, Egypt, Tunisia, Israel and Morocco) on credit card activity in these places. Our data-set combines monthly aggregated and anonymized card activity data[3] from 114 origin countries at each destination with a corpus of news coverage of the destinations at the different origins. Through downloading and translating more than 450,000 news articles from 57 different origin countries we can track how media outlets in these origin countries cover the same violent events in the 5 destination countries. Supervised machine learning allows us to capture the intensity of media coverage of such events in this large news corpus and this enables us to study the response of travel to the text content of media coverage across 57 different origins.

Our approach exploits a unique combination of dyadic card activity data and dyadic

[1]See Strömberg, 2015 and Prat and Strömberg (2011) for reviews.

[2]In 2016, terrorism in the US caused less than 0.01% of all deaths but was covered by newspapers more than any other cause of death – including the main causes: cancer and heart disease. See Combs and Slovic (1979) and Shen et al. (2018) and data compiled in *Our World in Data*.

[3]Mastercard made the anonymized and aggregated spending and transaction data available subject to robust privacy and data protection controls and in line with their principles guiding the ethical collection, management and use of data. We use data on aggregate spending and number of active cards.

news data, i.e. spending responses and reporting on the same set of events in different origin countries. This allows us to study people who are deciding about travel to a destination country from multiple points of origin. The effect of events at a destination can therefore be separated from the intensity of the news coverage. It is an advance over previous work that looks at the impact of the media which relies exclusively on the *timing* of news (as for example, in Bloom, 2009 and Ramey, 2011) rather than differential coverage across locations.

The kind of card spending data available to us is illustrated in Figure 1 which shows the pattern of card spending around a specific violent incident in which thirty-nine tourists were killed on June 26^{th} 2015 in Sousse, Tunisia. The majority of the victims were UK citizens. The lines in the figure contrast the response of tourism activity on two specific dyads based on the origin country: British (the dotted line) and German (the dashed line), while overall average spending patterns across all dyads is indicated by the solid line.

The figure highlights that overall tourism spending across origins falls immediately in the months following the attack. However, spending of British travelers dropped much more compared to that of Germans. This motivates exploring how differential intensity of news reporting on the event in the two origin countries is responsible for the heterogeneous response. And indeed, robust reduced-form results show that the media matters over and above the violent events themselves. We show that, even when we control for all variation specific to dyads and all time variation at each destination and origin, card activity on a dyad changes distinctly with differential media coverage.

To interpret the responsiveness to news coverage as a media multiplier requires a model of how beliefs are formed. We propose a novel approach in which violent events and news text both evolve with an unobserved categorical state that transitions between *danger* and *safety*. Potential visitors to a country do not observe this state and some agents therefore use news outlets in their home country exclusively even if these are prone to news cycles, sensationalism, restrictions in media freedom or political agendas.[4] Once we model the formation of beliefs, we can estimate the parameters of the updating process and to what degree the news-based beliefs drive the response of card activity to events, i.e. the media multiplier. Our model-based approach suggests that, when a country is perceived as dangerous by all potential visitors, card activity

[4]This is compounded of the use of so-called *availability heuristics*, whereby the ease with which information about a phenomena is recalled, affects an individual's estimate of the likelihood of it recurring. See Tversky and Kahneman (1973). Recent research in psychology is reviewed in Zhu et al. (2020).

falls by 53 percent with more than half of this due to the media multiplier. Drilling down further, we show that news-based beliefs in countries with free media seem to provide a much better explanation for changes in card activity which suggests that travelers rely more on the news coverage in their country of origin.

Having a theoretical model helps to offers both a way of interpreting both heterogeneity in and the time path of media reporting compared to a reduced-form approach. Moreover, it enables us to project the implications of our findings to wider contexts and by considering counterfactual paths. For example, we can use the distinction of coverage of violent events and other background news, which we gained from the supervised machine learning, to show that news shocks due to violent events are less severe when there is more background news. This is because background "drowns out" bad news.

The remainder of the paper is organized as follows. The next section reviews some relevant literature. In section 3 we discuss the data used in some detail. Here, we also discuss the supervised learning method through which we make the text data usable for the subsequent analysis. In section 4, we present reduced form results of the average effect of violence on tourism activity before incorporating the news data. In section 5, we propose a statistical model and fit this on both the news and the tourism activity data. Concluding comments are in section 6.

2 Related Literature

Our modelling approach follows Besley and Mueller (2012) who develop a model of latent beliefs to examine the reaction of house prices to negative events. Here we look at the importance of news coverage in belief formation. We show how both news and events change behavior through their influence on latent beliefs and that this generates a media multiplier. This contributes to the literature that estimates the impact of violence on the economy. For example, Abadie and Gardeazabal (2003) document the sizable negative economic consequences of terrorism in Spain. Amodio and Di Maio (2017) shows how firms bear the direct and indirect costs of violence and political instability. Jha and Shayo (2019) explore how individuals re-evaluate the costs of conflict upon being exposed to financial assets, whose prices may be vulnerable to the economic risks of conflict. Brodeur (2018) shows that successful terror attacks in the US reduce the number of jobs and total earnings in the affected counties. Alfano and Görlach (2019) show that the access to media content amplifies and spreads the impact of terror attacks on schooling choices. In all cases, changing beliefs is the mechanism for changes in behavior but there is no underlying model of how beliefs are formed.

The idea that changes in expectations are the mechanism through which news affect economic activity is a central pillar of work on "news shocks" (Ramey, 2011). Arezki et al. (2017) documents that news reports on resource discoveries have near immediate impact on the current account. Brückner and Pappa (2015) show that the bidding for the Olympic Games leads to a positive news shock. Eggers and Fouirnaies (2014) document that news about the technical declaration of a recession has a contractionary effect on the economy. Similar news shocks have been studied by looking at central bank announcements (see, for example, Glick and Leduc, 2012). The idea of a media multiplier is also relevant in the context of COVID-19 to study how media coverage of death affects economic decisions.[5] Responsiveness to news depends on the intensity of media coverage. To distinguish the impact of an event involving violence from media coverage, we use dyadic data news coverage, giving us heterogeneous news coverage of the same event in different places.

Our model of beliefs requires a way to identify the coverage of violence in news content and a statistical model of how this news content reacts to latent risks. We contribute on both these fronts. First, we deploy supervised machine learning methods to detect the coverage of specific, rare events (violence against tourists and fatal violence) in over 450,000 news articles. This approach allows us to test explicitly, through cross-validation, how good our method is in identifying reporting on violence in the news.[6] This approach stands in contrast to dictionary based methods like in Baker et al. (2016) or Hassan et al. (2019) or the use of topic models in Mueller and Rauh (2017) and Hansen et al. (2018). The performance of our method suggests that this way of producing data from text, i.e. letting the machine determine which parts of the text are most relevant, might also be useful in other applications. But we also add to this literature by providing a framework, which integrates the resulting text features, event coverage, and background news, into a model of belief updating. Other, more complex generative models of text like topic models could also be integrated into the economics literature in this way.

[5] In fact, the numbers from our model are close the effects of a complete lockdown found by Carvalho et al. (2020). Zhao et al. (2018) theoretically show that media coverage can change the dynamics of a pandemic dramatically.

[6] There is still relatively little work in economics leveraging nonlinear supervised machine learning. We use cross-validation to show that an ensemble of naïve bayes and random forests provides relatively large gains over simpler methods commonly used (see Manela and Moreira, 2017; Becker and Pascali, 2019).

Our Bayesian approach to beliefs specifies how individuals make sense of changes in the number and composition of news. We assume that agents are categorical thinkers, some of which only update their beliefs based on the easily available news in their own country. As a result, beliefs formed at a specific location and economic behavior can react strongly to local news coverage. This is a useful way of modeling beliefs for two reasons. First, there is a wealth of evidence from psychology that people use crude categories when making sense of the world.[7] Second, there is a large literature in psychology focusing on availability heuristics, i.e. the human tendency to judge probabilities by the ease with which information underpinning them can be recalled Tversky and Kahneman (1973). These heuristics could provide an alternative motivation for the strong changes in beliefs we find. [8]

The resulting strong, nonlinear response of beliefs to local coverage is particularly relevant to behavior vis a vis developing countries since international news coverage is rare and tends to follow bad events. If this is poorly understood by those who follow news, it can shape the international image of affected countries, and lead to greater economic isolation. Given the economic and social benefits of openness and economic integration in general (see Frankel and Romer, 2008; Melitz and Trefler, 2012) and tourism in particular (see Faber and Gaubert, 2016), the way media reports on countries could thereby have detrimental consequences for economic development.

The impact of the media on economic and political outcomes has been widely studied.[9] Eisensee and Strömberg (2007), shows that news on droughts affect US disaster relief. Similarly, Durante and Zhuravskaya (2018), suggests that offensives in the Israel Palestine conflict are strategically aligned to minimize news coverage in the US, while Jetter (2017) suggests that Al Qaida activity may be endogenous to preceding US television news coverage. We contribute to this literature by providing dyadic news data from 57 origin and 5 destination countries which we embed in a model of violence, reporting, and beliefs. The *drowning out* effect we find in this way, could motivate the kind of strategic behavior demonstrated by Durante and Zhuravskaya (2018) and Jetter (2017) as the respective actors try to minimize (Israeli armed forces) or maximize

[7] Fiske (1998) provides a review of social psychology literature; Fryer and Jackson (2008) provide a model based on this idea and a discussion of the biases that this induces in decision-making.

[8] See Bordalo et al. (2016, 2018); Handel and Schwartzstein (2018); Azeredo da Silveira and Woodford (2019); Zhu et al. (2020) for overviews and examples.

[9] See, for example, Stromberg (2004) on redistributive spending, Besley and Burgess (2002) on government accountability, Gentzkow (2006); Bursztyn et al. (2017) on voter turnout, Snyder and Strömberg (2010) on citizen knowledge, DellaVigna and Kaplan (2007); Enikolopov et al. (2011); Adena et al. (2015) on voting patterns, Durante et al. (2019) on the proclivity towards populist rhetoric.

(terrorists) the impact of their actions.[10]

Finally, we contribute to an emerging literature on the consequences of violence and disorder on trade and economic integration. For example, Besley et al. (2015) measure the cost of piracy in the Gulf of Aden on shipping costs and hence on trade. They emphasize that violence can increase trade costs, which the trade literature has shown can have a significant impact on trade flows (see, Feyrer, 2019, Donaldson, 2018). However, the mechanism we focus on here is closely related to Burchardi et al. (2019) who show that better information on a country, due to ancestry, is an important driver of foreign investment decisions. Tourism is an increasingly important sector that supports up to 313 million jobs across the globe.[11] We show that tourists with more information react less to idiosyncratic media reporting because they have a broader informational base, and this can be an important factor in the overall effect of violence on economic ties. This is particularly relevant in the MENA region, which is one of the least economically integrated regions (Rouis and Tabor, 2012) and where growing economic ties to Europe are important.

3 Data and Feature Extraction

This paper uses three main data sources: (i) aggregated monthly spending data by origin and destination country, (ii) measures of terrorism and conflict events and (iii) a large corpus of dyad-specific news content. We describe each of these followed by a discussion of the supervised machine learning method that is used to identify news coverage of fatal violence and attacks on tourists.

3.1 Aggregated Spending Data

Mastercard provided us access to an anonymized and aggregated monthly data set, which included the number of transactions and number of active cards based on spending in five different countries (Egypt, Israel, Morocco, Tunisia and Turkey) by the country of origin of the card. Where Mastercard's controls may have resulted in, for example, in a blank monthly-dyad observation, we excluded all of those dyads where we have fewer than 60 months (5 years) of data and origin countries with fewer than 3 out of

[10]We observe in our sample that news coverage sharply responds to violent events, but does not increase in anticipation.

[11]See WTTC (2018). A small literature within economics studies tourism using robust methods. For example, Faber and Gaubert (2016) show that, in Mexico, tourism produced significant local economic gains. Neumayer (2004) uses cross country panel data to show that violence is negatively associated with tourism arrivals.

5 destination dyads. The origin countries in our sample span all continents but trend towards higher income countries and those that are geographically closer to the destination countries.[12] Figure A1 maps all of the origin (card-issuing) countries that we have in the sample.

There are notable differences between origin countries with low volume of cards active per month in tourism spending, such as countries like Haiti and Namibia compared to higher volume tourism spending from countries such as Germany and the United States. The aggregated card data can be a proxy for annual patterns in travel flows: for a small set of countries annual data on travel flows is provided to the United Nations World Tourism Organisation (UNWTO). Appendix Figure A2 highlights that the annually-aggregated card data correlates with the travel flows data very closely. Appendix Table A2 presents further regression evidence highlighting the close fit.

3.2 Data on Violent Events

We use five different data sources, three of which are hand-coded event data while the other two are constructed using automation.

Manually-coded data sources As our core data on terrorism we use the Global Terrorism Database (GTD) which is an open-source database that codes information on terrorist events around the world between 1970 and 2017 based on reports from a variety of media sources. These reflect world-wide rather than country-specific news coverage and the information is verified by the GTD research team to establish the credibility of the information source. The data focus on the type of violent events that are likely to influence the desirability of a destination for potential travelers.

As supplementary human-coded sources of data, we also leverage the Georeferenced Event Dataset (GED) provided by the Uppsala Conflict Data Program (UCDP).[13]

Automated Data We use the Integrated Crisis Early Warning System (ICEWS) database created for the Defense Advanced Research Projects Agency (DARPA) and Office of Naval Research (ONR). This event-level data comprises coded interactions between sociopolitical actors (i.e., cooperative or hostile actions between individuals, groups, sectors, and nation states). Similar to the approach used in Fetzer (2019), events

[12] We also include cards originating in the destination countries themselves. These data can be dropped leaving results unaffected.

[13] Results are also similar when studying the Armed Conflict Location & Event Data Project (ACLED) data. As these are currently only available for the three countries on the African continent we do not include them in the analysis.

are identified in a fully automated way and are extracted from news articles, essentially consisting of triplets based on a subject (a source actor), an event type (indicated by a verb) and an object (a target actor). Geographical-temporal metadata are also extracted and associated with the relevant events. In this paper, we focus on events that have been coded as assaults, which include events such as hijacking, suicide bombings, and assassinations along with data on fights or escalations, which includes the use of military force, fights with artillery and tanks and aerial bombing.

The second automated dataset is the GDELT platform which monitors the world's news media from nearly every corner of every country in print, broadcast, and web formats, in over 100 languages, every moment of the day stretching back to January 1, 1979 to produce data on events.[14] GDELT is more inclusive, yet it may also include more false positives and it also has less stable source material over time and codes the news sources from 2014 onwards only.

Both of these data sources have in common that they aim to identify the "true" set objective violent events. Neither of them provides a measure of the likely salience of an event nor the intensity of news coverage about a violent event across different countries.[15] We next describe how we construct a dyadic dataset of news coverage for 57 of our issuing countries, i.e. for 285 dyads.

3.3 Data on News

The news data variable that we construct is intended to proxy the news coverage that potential traveler have access to in a given country when they decide on their holiday destination. A key concern here is measurement error both because the media landscapes differ across countries and because it is not clear a priori which specific news items are viewed. To obtain dyad-specific variation in news coverage, we develop a large scale corpus for 57 tourist-origin countries. For each traveler origin country, we identify a leading news source for which a digital archive of all articles is available over our sample period. For each of these sources, we then download all articles that relate to each of our five destination countries covering the period from 2009 to 2016. The tourist origin countries for which we have both card data as well as media coverage data are indicated in dark grey in Figure A1. The countries for which we have news data represent, by far, the biggest chunk of the world economy, comprising all G20 nations along with a host of other significant emerging-market economies. Hence, although we

[14] This data has been used by Manacorda and Tesei (2020).

[15] To the best of our knowledge, such a data set does not exist.

cannot say that this is globally representative, the consequences of changes in tourist spending in these countries are likely to be economically important for the destination countries that we study.

The resulting data set contains more than 450,000 individual articles, out of which 307,000 articles were translated into English using *Google* Translate. The translation to English allows us to produce a single consistent classifier to code individual articles.[16]

3.4 Supervised Machine Learning Approach

We use supervised machine-learning to classify individual articles according to whether they report violent incidents or incidents directly involving tourists. We proceed in four steps. First, we use human coding to classify a subset of the data which we use as a training dataset to generate our news indicators. Second, we use supervised machine learning to train a set of classifiers to predict the human classifications in the training set and classify unseen articles. In this step, the availability of training data allows us to check performance of the classifier using cross-validation. Third, we check a subset of the classified articles by hand to generate out-of-sample performance measures and reduce measurement error further. Finally, we aggregate the resulting scores to produce a count of news about violence for each dyad/month or dyad/day. We then express this as a share of all news in the same dyad/month.

Training data set To build the training set, human coders classified a sample of around 30,000 articles (approximately 7% of the data). The coding guidelines consisted of two binary classification questions that were used to construct two separate measures of violence. Specifically, human coders were asked to flag up individual articles with a binary indicator if:

1. *the article indicates that there were fatalities as a result of violence*

2. *the article indicates that tourists were harmed due to a violent event*

The underlying classes are quite unbalanced relative to the population of articles. This can make it difficult for statistical learning methods leveraged for classification purposes to separate the data adequately. To navigate this issue, in drawing our training sample, we follow Japkowicz and Stephen (2002) and oversample articles around days for which the Global Terrorism Database indicated that an event occurred.

[16] Appendix Table B1 presents the main source by country, the origin language and the number of articles included in our database. For a few countries only news wire agency reports were available; our results are robust to dropping these countries from the analysis.

Classification approach In the second step, we train a set of classifiers in Python using the scikit-learn packages developed by Pedregosa et al. (2011). Individual articles are represented using the common bag of words language model so that each document can be expressed as a vector of counts. We use standard stemming procedures and remove stop words. We then produce all trigram word features and exclude terms that appear in less than 100 documents.

We use an ensemble of three classifiers to identify violence. To build the ensemble we made extensive use of cross-validation with our training data to get an impression of the likely out-of-sample performance and to refine what part of the text to focus on, which classifiers to use and how to combine them. All three classifiers are built by looking at the full text and headline. We use a simple Naïve Bayes classifier and two Random Forest classifiers with hyperparameters described in the Appendix D. This produces three different classifiers, indexed by k, which allow us to obtain for each document, denoted by D_i, three estimates of the probability that classifier k contains news coverage of the type that interests us, denoted by $\widehat{P}_k(Y_i = 1|D_i)$, where Y_i is an indicator denoting whether a given document D_i is either covering violent events with fatalities or violent events in which tourists were targeted.

Naïve Bayes methods belong to the class of generative linear classifiers, is known to perform well with textual data and sparse feature sets. Random Forests, on the other hand, are particularly suitable to allow for non-linearities using smaller feature sets. The only difference between our two Random Forests, is that in one of them we first use Singular Value Decomposition (also referred to as Latent Semantic Analysis which has recently been used in Iaria et al. (2018)) to reduce the dimensionality of the feature space from (tens of) thousand of word counts features into a much lower 100-dimensional continuous score representation of individual documents D_i. These individual components are then used as numeric features in the construction of the classification trees using the random forest formulation.

In cross validation on three folds our ensemble reaches an AUC of 0.95 and an average precision of 0.85 for fatal violence and an AUC of 0.97 and average precision of 0.65 for attacks on tourists. These are very good statistics but they come from an evaluation of a balanced dataset and precision falls when we instead evaluate an imbalanced sample. This is an important issue for spotting of violence against tourists as this is a heavily imbalanced class even in the training data.

Classification Ensemble and Validation For classification purposes, we use a soft voting ensemble method, i.e. we average our three different classification scores $\widehat{P}_k(Y_i = 1|D_i)$ according to the function:

$$\mathbb{1}(D_i) = \begin{cases} 1 & \text{if } [\frac{1}{3}\sum_{k=1}^{3} \widehat{P}_k(Y_i = 1 \,|D_i)] > c \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

To chose the cutoff c in (1), we count how often the indicator, $\mathbb{1}(D_i)$ would have been correct for different values of c within the training sample. The Bayes optimal decision rule that maximizes overall accuracy would be to pick a cut-off of $c = 0.5$. Yet, the cross-validation exercise highlights that the class imbalance may result in low precision under this rule. As we are concerned that we get too many false positives resulting in a very noisy monthly measure, we choose a higher cutoff which provides 90 percent precision within our training sample.[17] This cutoff gives us 16,906 news articles with fatal violence and 1,082 news with violence against tourists out of over 450,000 articles.

To reduce measurement error we conducted some ex-post manual coding for the classification of articles indicating violence against tourists. While our results are robust to relying only on the machine-generated output, it is prudent to perform such a manual check and some amount of ex-post refinement. We considered all articles with an ensemble probability indicating violence against tourists above 0.75 along with the top 100 articles ranked by the ensemble estimate from in (1) all origin sources. [18] For this sample we re-code mistakes by hand. In other words, we set $\mathbb{1}(D_i) = 0$ by hand if we find a false positive and set $\mathbb{1}(D_i) = 1$ if we find a false negative. Of $1,082$ observations that were marked positive by the algorithm we recoded 103 to negatives, implying that our method did indeed achieve a precision of over 90 percent out-of-sample. In the almost 5,000 additional news items that were hand-coded we only found an additional 608 positives with a rapidly declining rate so that we suspect that the remaining articles will not contain a lot of actual positives. After hand-coding we therefore have $1,587$ positives in over $450,000$ negatives that feed into our media coverage-based measures of violence against tourists.

In the Online Appendix D we describe the classification approach in greater detail, while appendix Tables B2 and B3 provide some sample headlines of articles coded as

[17]Our results are robust to using alternative cut-offs. In Appendix Table A9 we use both a higher cut-off with 95 percent precision and the Bayes-optimal cut-off of 0.5.

[18]We do the latter to ensure that the model has not only fit to sources that emit a lot of news like news agencies in Russia and China.

covering violence and flagged up as capturing that tourists are targeted by a violent event. In the appendix we also discuss the "mistakes" made by the algorithm and why they are often capturing something indicating risks to tourists. It is therefore no surprise that our results, even in the most demanding specifications, are robust to using only the raw $\mathbb{1}(D_i)$ that come out of our automated procedure and using different cutoffs.

This is also important from a methodological perspective. We have managed to provide a meaningful, fully-automated way to identify fatal violence and violence against tourists even though they only appear in about 4 and 0.4 percent of all articles, i.e. are extremely rare. We did this simply by asking our research assistants to code a subset of the articles - the classifier then automatically extracted the relevant features from the data. Furthermore, our supervised learning approach allowed us to check the error rate explicitly and to reduce it through setting hyperparameters and building of the ensemble. This would not have been possible with an unsupervised or dictionary-based method.

3.5 Patterns in the Reporting Data

We approach news reporting on violent events as standing in a relation to overall reporting on a country. Define $B_{hdt} = \sum_{i \in hdt} \mathbb{1}(D_i)$ as the monthly *count* of news stories in dyad (hd) at date t, either about fatal violence or attacks on tourists, based on equation (1). Then our core variable to represent news coverage in a dyad is

$$n_{hdt-1} = \frac{B_{hdt-1}}{N_{hdt-1}} \tag{2}$$

where N_{hdt-1} is the count of *all* news stories featuring country d reported in our news source for country h at date $t-1$. Thus the variable in equation (2) reflects the news coverage of violence as a *share* of all news. This captures the idea that news coverage of violence affects tourists more when they are important relative to other news. Thus, if bad news stories are swamped by other stories, they will have less impact on tourism spending.

Before turning to the full analysis, we document how the news reporting relates to underlying events beginning with daily data. This provides evidence in support of the underlying common trends assumption which matters in the empirical analysis below where we require that reporting only occurs *after* an attack and not prior to one.

Daily Data To look at patterns of news reporting around known events, we use the GTD daily event data to construct a balanced panel at the dyad level covering two week windows around each event. In total there are 3704 recorded events across the five destinations. Given the 57 countries for which we have media coverage, the balanced daily event-level dataset comprises 6.1 million rows.

This data layout allows us to explore the pattern of news reporting around known events. One concern, following Jetter (2017)'s study of US media coverage, is that news stories might precede (and even encourage) acts of violence within a time window (such as a week). This would show up in our data as increases in reporting intensity *before* GTD events. To investigate this possibility, we estimate the following empirical model:

$$p_{hdt} = \alpha_k + \alpha_{hd} + \alpha_t + \sum_{\tau=-14}^{14} (\beta_\tau \times \text{Timetoevent}_{e,t-\tau}) + \epsilon_{hdt}$$

where e indexes a specific event, h and d indicate the reporting dyad, while t indicates time which is now a daily observation. The above regression controls for event fixed effects, α_k, dyad fixed effects, α_{hd}, and daily fixed effects, α_t. In the case of multiple events in close temporal proximity, we would be double counting the reporting on dyad $\{h, d\}$, and hence we adjust standard errors to allow for two-way clustering at the level of the dyad and the event.

In Figure A4 we plot the point estimates $\hat{\beta}_\tau$, which suggests that there is no anticipatory element in the reporting data. Panels A and B show the measures generated from our method for classifying articles. Specifically, we construct the share of articles per day that are classified as reporting either fatal violence or tourists being attacked. The patterns suggest a sharp increase in the share right after the event date. This dissipates quite quickly with most reporting occurring on the day of the event and for around two days afterwards. It is important to note that this happens despite the fact that the total number of news stories increases slightly, i.e. we find this relative reporting effect despite increased reporting overall.[19]

Monthly Aggregates In our analysis we use aggregates of our news measures to the monthly level, as the dyadic aggregate spend data is only available at the monthly level. Figure A6 reports the mean shares at a monthly frequency for the four countries most affected by violence against tourists in our sample period (Tunisia, Turkey, Israel and Egypt). It depicts the average share across all dyads of monthly events defined by (1) for violence against tourists – the black dashed line with the axis on the right hand side – and (1) for fatal violence on a monthly basis – the black solid line with the axis

[19] In Appendix A.2 we provide some further evidence shedding light on which event characteristics are associated with, on average, more extensive media coverage.

on the left hand side.

Figure A6 shows a lot of variation across time in reporting for all countries. But there is also considerable variation in the intensity of reporting across destinations. Reporting on violence is often a considerable part of reporting on Tunisia. At the time of the Sousse attack, for example, violence against tourists occupied around 40 percent of all news. Reporting in Egypt, Turkey and Israel is more intense for fatal violence than it is for violence against tourists. However, this coverage never occupies more than 10 percent of the news. The most extreme example is Israel where news on violence never exceeds 7 percent of reporting and violence against tourists never more than 3 percent.

4 Reduced-form Evidence

We motivate the idea of a media multiplier by first studying reduced-form evidence on the response of card activity to violent events at the destination countries and dyad-specific news coverage of these events. Underlying this analysis is the hypothesis that violence and news reporting are informative in some way about a characteristic that potential travelers care about when they book or cancel their travel. For now we will not make this hypothesis explicit but instead focus on a reduced form relationship between violent events, news, and card activity.[20]

4.1 Core Findings

We begin by looking at the relationship between news, violence and card activity using the following specification:

$$y_{hdt} = \alpha_{hd} + \alpha_{ht} + \alpha_{dm(t)} + \xi_1 n_{hdt-1} + \xi_2 v_{dt-1} + \varepsilon_{hct} \quad (3)$$

where the origin countries are denoted h, destination countries are denoted d and time is denoted t. The dependent variable y_{hdt} are measures of tourism activity either measured as aggregated spending or as the number of distinct cards active in a given month. The specification in equation (3) also introduces dyad fixed effects, α_{hd}, origin/time fixed effects, α_{ht}, and destination/month of year fixed effects, $\alpha_{dm(t)}$, to account for destination-specific seasonality in demand. In some specifications we will use α_{dt} instead of $\alpha_{dm(t)}$, i.e. we include destination by time fixed effects. These fixed effects capture *all* variation at the destination/time level including news events, which are re-

[20] Appendix Section A.1 presents the reduced form relationship between violent events and tourism activity. This is the type of analysis of the early literature documenting negative economic effects of violence (see Neumayer, 2004) which has ignored the media reporting channel that is the focus of this paper.

ported on with a common level of intensity across origin countries. The relationship between news coverage and spending is then identified through idiosyncratic variation in the intensity of news reporting across origins. The variable v_{dt-1} is a broad set of measures capturing violent events occurring in a destination country d, while the variable n_{hdt-1} is our dyad-specific news variable.[21] We lag both the violence measures as well as the news measure by one month as we expect a lagged response.

Before presenting the results, it is worth stressing that our analysis is only exploiting within-dyad variation which absorbs all factors like distance or cultural factors. In addition, we are including home country by time fixed effects, α_{ht}, de facto absorbing a host of factors that may drive the level of tourism activity that is explained by origin-country level idiosyncrasies (such as holiday periods, which may differ across countries). In this way we are modelling the rate of tourism activity for a given destination among our sample of five countries *relative to* the overall amount of tourism originating in country h. Hence all magnitudes are based on comparing the attractiveness of our destinations compared to the other five destinations in our data rather than other parts of the world. Thus, we are only able to say whether international travel to Turkey decreased after the terror attacks in the country relative to Egypt, Israel, Tunisia and Morocco. Arguably this is a conservative approach, since there could be reputational externalities whereby potential travelers shy away from the entire region due to the turmoil in one of our countries.

Finally, it is important to keep in mind that there are no discernible pre-trends in news coverage but that it responds sharply to terror events as shown in Figure A4. It is clear here that, in our sample, reporting sometimes responds to events but not the other way around. This allows us to assess to what extent dyad-specific news coverage of violent events has econometrically relevant explanatory power for dyad-specific card activity. We show in section 4.3 that idiosyncratic factors, such as the cultural proximity to the victims of an attack, drive a lot of this dyad-specific variation.

4.2 Results

The results from estimating (3) are in Table 1 where the top panel (Panel A) is based on news coverage of violence against tourists while the bottom panel (Panel B) is for reporting on fatal violence in general. In columns (1) through (3), the independent

[21] Appendix Section A.1 documents the reduced-form effect of violence on tourism activity across four different data sets to measure v_{dt-1}. We construct principal components to combine these four data sources to construct as-broad a measure of objective measured violence in a destination.

variable is the (log of) card spending. Column (1) of Panel (A) shows that if the share of stories about tourist violence were to go from zero to one then tourism spending would fall by 0.552 log points or 42 percentage points. This results holds up in column (2) when we add the controls for violent events and the coefficient stays roughly similar. Hence, news coverage of violent events is clearly correlated with card activity over and above the underlying events themselves.

Column (3) in Table 1 is our most demanding specification in which we control for destination $\times$ time fixed effects. This set of fixed effects is collinear with any time varying factors at the destination level, such as destination-specific demand seasonality, macro-economic developments, political developments or violence at the destination. As a result, in this specification, we rely only on the differential intensity in news reporting across different reporting countries. The coefficient on the share of bad news (based on tourist violence) remains highly significant although the coefficient in this saturated specification falls to 0.205 log points. This continues to suggest that a significant part of the overall effect in columns (1) and (2) is driven by pure news reporting. Note, this also provides some evidence for the idea that tourists are not perfectly informed as the same basic risks at a given destination trigger dramatically different responses depending on the news environment, something which we return to below.

Columns (4) through (6) in panel A of Table 1 repeat the specifications in the first three columns but with the log of active cards as the dependent variable. The estimates are similar, which is important as it indicates that the results on spending are not only due to changes at the intensive margin in which tourism spend is less, but rather on the extensive margin as fewer tourists travel to a country following media-coverage of violent events.

Panel B of Table 1 repeats the specifications in Panel A except for measuring reporting on all fatal violence rather than just attacks against tourists. The coefficients are only slightly smaller compared to those presented in panel A. We also find statistically and economically significant relationships throughout although we lose statistical significance in the most saturated specification in column (6). The results suggest that differential intensities in media coverage of violence may have an important independent direct effect on tourism travel and spending, which is particularly relevant given that violence more broadly, not necessarily directed at tourists, is a lot more common.[22]

[22] A horse race combining both of these measures together suggests that we find negative and statistically as well as economically significant coefficients for both news measures, suggesting that our news measures are not simply picking up the same type of news. These results are available upon request.

Taken together, the results in Table 1 suggest that card activity reacts to news activity even when controlling for violence and destination/time fixed effects. To guide the theoretical model in section 5 and to test robustness of the reduced form results, we implemented a set of additional tests.

4.3 Robustness and additional reduced form exercises

Relative Nature of the News Effect Up until now we have always introduced news coverage of violent events relative to overall coverage of a destination country. Appendix Table A5 highlights that this is the empirically most suitable way of measuring news in the reduced form exercise. In section 5, we provide a statistical model of reporting and belief formation which makes sense of this relative reporting effect.

Timing In the main reduced form specification, the news measure enters with a one month lag, n_{hdt-1}. In Figure 2, we explore different leads- and lags. This highlights that there is a strong, immediate negative effect which becomes stronger when studying higher order lags. This is not surprising given that tourists book travel in advance and will not book travel into a place with negative reporting at the time of booking. Importantly, we also find no pre-trends in card activity before news events.

However, Figure 2 also highlights that a simple reduced form regression analysis as in specification 3 is unlikely to capture well the overall dynamic impact of news coverage and as such, will result in a worse fit and poor out-of-sample performance. This is one of the reasons why we put emphasis on section 5 in which we provide a statistical model that we fit to the data, and which will end up significantly outperforming the reduced-form model in terms of matching patterns in the data.

An instrumental variable (IV) approach We also present results from an IV approach in Appendix Table A6. The main idea is to exploit the distribution of the nationalities of casualties in the different violent events as an instrument for media coverage of attacks on tourists. An event becomes much more newsworthy in a country of origin if individuals from that country are affected. However, risks for the specific nationalities have not suddenly increased. The idea is therefore that such an instrument is both relevant in driving media-reporting and plausibly excludable as the distribution of nationalities of casualties is, conditional on fixed effects and the event occurring, random.

We also relax this identifying assumption by focusing exclusively on *media reporting spillovers* where reporting is higher for other origin countries, which are culturally close.[23] Specifically, we study the impact of reporting on a casualty from origin country X by focusing exclusively on media-coverage and spending responses in other origin countries that share a common language and geography with country X. In simple terms, we can identify the causal impact of a German tourist being killed on tourism activity of Swiss and Austrian travelers that is due to the German casualties' impact on media reporting in Austria or Switzerland, which share the German language, while fully ignoring German card activity.

We find a consistent first stage in which casualties increase news reporting in the origin country of the victim and strong reporting spillovers through shared language and contiguity.[24] The instrument provides a Weak IV identification F-statistic above 10 for most of the exercises. In the second stage, we find that point estimates on the reporting share increase relative to our reduced form, that may suggest that the reduced form estimates are affected by attenuation bias. Importantly, results are robust when we rely only on spillovers, by dropping all data from all dyads that ever had a casualty. This lends support to our idea that it is variation in whether a news story is picked up or not in a home country that causes a response in card activity.

Further Robustness We conduct some further robustness checks. The results are robust to dropping each potential tourist origin country in turn; dropping countries for which only wire-service reporting is available leaves the results very similar. Appendix Figure A9 shows that our results are broadly carried by all the different destination countries we study. Appendix Table A9 presents some further robustness. Results are robust to controlling for dyad-specific linear time trends, in addition to destination specific non-linear time trends (columns 1 and 4). Results are also robust to controlling for time-varying exchange rate movements at the dyad level (column 2). Further, results are robust to using alternative cutoffs for the classification of individual articles in the news corpus for which we did not use hand-coding (columns 5-10).

5 Modelling the Impact of News Coverage on Card Activity

In this section we posit a model in which agents form beliefs about a latent state (danger/safety) based on observing violence and news. News and violence only matter

[23]This approach is similar to Persson and Tabellini (2009) and Acemoglu et al. (2019) who instrument for democracy using democracy in neighbouring countries.

[24]The results are presented in Appendix Table A6.

to the extent that agents dislike booking travel to dangerous locations because these locations tend to stay dangerous. Our model assumes that two sets of unobserved beliefs, linked to the latent state, drive card activity. One set of beliefs is formed based on observing violent events, whereas the other set of beliefs does not observe events but needs to rely exclusively on media reports to update on whether a destination is dangerous. The reaction of card activity to these two sets of beliefs is what allows us to define and quantify the media multiplier. We can then use the fitted model to consider policy experiments and show that the model of beliefs provides a better description of card activity data than the reduced-form model.

5.1 The Model

Model Overview Figure 3 presents a graphical overview of the model. The circles represent endogenous variables, the boxes represent parameters while the arrows illustrate dependencies between them. Endogenous variables include latent, i.e. unobserved variables in the clear circles whereas those in the shaded circles are observable. The reduced-form reported above relates the shaded circle on the right of Figure 3 representing card activity, y_{hdt}, to those on the left: violent events, v_{dt}, news coverage of violence, B_{hdt}, and overall news items, N_{hdt}.

There is an underlying categorical state s_{dt} denoting whether a country is safe or dangerous and information affects beliefs about this state. To simplify the analysis we assume a dual belief system denoted by Π_{dt} and π_{hdt} where the former reflects data about violent events, v_{dt} , and the latter reflects news coverage, i.e. bad news, B_{hdt}, and all news, N_{hdt}.

The parameters in the square boxes in Figure 3 govern this process and are estimated following a procedure spelled out below. We start by deriving expressions for the beliefs Π_{dt} and π_{hdt} and then estimate the weight on news-based beliefs, χ, which best explains movements in card activity. This parameter represents the media multiplier. In addition to the media multiplier, we are also able to estimate weights on lagged beliefs, ω_τ, which we interpret as the share of travelers that book or cancel their itinerary τ months before they travel.

Card Spending, Latent States and Beliefs Suppose that destination country d at date t is characterized by a state, s_{dt}, where $s_{dt} = 1$ denotes a *dangerous* destination and $s_{dt} = 0$ denotes a *safe* destination. The underlying empirical model is based on Besley and Mueller (2012) and Besley et al. (2015), who suppose that there is an underlying latent state that can be modeled as a Markov process. In Figure 3 this Markov process is denoted as s_{dt} with a recursive arrow together with persistence parameters p_d and q_d. Parameter p_d captures the persistence of danger at destination d wheras q_d captures the persistence of safety.

The state of danger to our model is central as it determines both the statistical process that drives violence through parameters μ_{sd} and σ_{sd} and the statistical process that drives news reporting through parameters η_s. We assume that all agents, including the agents that base their beliefs on observations of the news, know the persistence parameters of the Markov process. We estimate these persistence parameters from the violence series. The implicit assumption here is that agents who observe only news still care about danger because this implies a change in violence expectation for the future as in Besley and Mueller (2012) and Besley et al. (2015).

At each date t, $\hat{P}_{hdt}$ is the belief that a destination country d is dangerous as perceived by potential travelers residing in country h. Due to the different dates at which people book their travel, spending is determined by a weighted average of past beliefs.[25] Hence, equation (3) is replaced by:

$$y_{hdt} = \alpha_{hd} + \alpha_{ht} + \alpha_{dm} + \xi \sum_{\tau=0}^{-9} \omega_\tau \hat{P}_{hdt-\tau} + \varepsilon_{hct} \tag{4}$$

where ω_τ is the weight on each lagged value, i.e. at date $t-\tau$. Equation (4) also has the same fixed effects as equation (3). In this framework, n_{dht} and v_{ht} affect spending through affecting beliefs $\hat{P}_{hdt}$. Since we do not observe beliefs, we posit that they can be represented by a function $\Gamma(\cdot)$ such that

$$\hat{P}_{hdt} = \Gamma\left(\Psi_{hdt}, \Omega_{dt}\right)$$

where Ψ_{hdt} is the history of news reporting up to date t and Ω_{dt} is the history of violent events up to date t. Note, only Ψ_{hdt} varies at the dyad level. We will specify $\Gamma(\cdot)$ below combining two stylized types of belief formation that rely exclusively on either Ω_{dt} or Ψ_{hdt} that we refer to as "event-based" and "news-based". This will allow us to specify the media multiplier, which captures the extent to which $\hat{P}_{hdt}$ are driven by incorporating Ψ_{hdt} into the information set.

Event-based Beliefs We regard beliefs to be *event-based* if they are based on curated information databases like UCDP, GTD and the GDELT event database that tracks all

[25]Note, we do not distinguish the timing of booking travel and changing earlier bookings but simply focus on the aggregate spending response.

available reporting. Since this information is not confined to any specific country, such beliefs are common across all origin countries regardless of media coverage. This is "as if" the individual forming these beliefs can observe the event history, Ω_{dt}. By construction, such beliefs are not subject to media influence, depending only on violent events thus providing a benchmark against which to calibrate the media multiplier. A measure of such beliefs can be constructed based on data on the history of violent events up to t, Ω_{dt}.

To do so, assume that our measure of violent events is distributed normally i.e., $v_{dt} \sim N(\mu_{sd}, \sigma^2_{sd})$. This allows the mean, μ_{sd}, and the variance, σ^2_{sd} to vary with the state s_{dt}. At each date, there is a destination country-specific transition probability between states where p_d is the probability of transitioning from dangerous to safe and q_d is the probability of transitioning form safe to dangerous. This gives a parameter vector for the model with six elements for each destination country d, summarized as $\theta_d = \{\mu_{0d}, \sigma^2_{0d}, \mu_{1d}, \sigma^2_{1d}, p_d, q_d\}$.

Event-based beliefs correctly assume that the probability that a destination is dangerous at time t is given by

$$\Pi_{dt} = \Pr(s_{dt} = 1 \mid \Omega_{dt}, \hat{\theta}_d), \tag{5}$$

where $\hat{\theta}_d$ is the estimated parameter vector which we assume is known to all types. Event-based beliefs are updated using Bayes rule as new information on violent events is revealed, i.e.

$$\Pi_{dt} = \frac{E_{t-1}\left[\Pi_{dt}\right]}{E_{t-1}\left[\Pi_{dt}\right] + \left[1 - E_{t-1}\left[\Pi_{dt}\right]\right] \gamma\left(v_{dt}\right)}$$

where $\gamma\left(v_{dt}\right) = \frac{\phi(v_{dt}|0)}{\phi(v_{dt}|1)}$ is the likelihood ratio derived from the normal distribution densities and where

$$E_{t-1}\left[\Pi_{dt}\right] = \Pi_{dt-1} \times p_d + (1 - \Pi_{dt-1}) \times (1 - q_d)$$

is the prior from the previous period. For v_{dt}, we use the principal components across the data from different violence sources.[26] Together these make up the elements of the history Ω_{dt}. However, as with any Bayesian approach, the prior history is fully captured by beliefs up to $t-1$.

Note, we assume that no additional information about latent risks is contained in dyad-specific reporting. This means that we can estimate the parameters of the latent state, p_d and q_d, from the events data. We assume that news-based beliefs can also rely

[26]To aggregate the different components into a single number, v_{dt}, we use the point estimates on the first two components from Table A3.

on knowing the resulting parameters p_d and q_d. In Figure 3 this is captured by the arrow from these parameters to both event-based and news-based beliefs.

Under this assumption, the empirical estimates of the parameter vector of the Markov switching model, $\hat{\theta}_d$ can be estimated from the data on violent incidents, v_{dt}, using the EM algorithm (Hamilton, 1990) and are provided in Appendix Table A10. They show strong persistence in the state in four out of five destination countries. The model fits less well for Morocco as the country experiences almost no violence and the differences in μ_{sd} are therefore minimal between what the model picks as the two categorical states.

Figure A7 reports our estimates of Π_{dt} for Egypt, Tunisia, Turkey and Israel. This approach permits a classification of whether a country is deemed to be dangerous or safe based on the level of violence specific to each destination. Thus, unlike equation (12), the effect of a given change in v_{dt} is heterogeneous across different destinations depending the history and persistence of violence. This makes sense; what would be deemed to be violence pointing to a state of danger in, say, Israel (bottom right) is different from what would be considered danger in Tunisia (top right). These differences reflect estimates of μ_{sd}. The probabilities based on equation (5) indicate that Tunisia was the first of our destinations to become dangerous based on the level of violence and this was followed by Egypt and Turkey. Danger in Israel based on this method is less persistent.

Spending by agents with event-based beliefs depend on the beliefs displayed in Figure A7. Deviations from this will allow us to understand how large the media multiplier is. However, in order to be able to quantify the media multiplier we need to estimate a model of how news-based beliefs are updated which is consistent with the Markov Chain model and the estimated parameters in $\hat{\theta}_d$ in Appendix Table A10. We now turn to this.

News-based Beliefs News-based beliefs are assumed to use the same statistical model of danger and safety as event-based beliefs and to rely on the same underlying persistence parameters $\hat{p}_d$ and $\hat{q}_d$. The difference is that news-based beliefs are formed based only on news coverage from an individual's home country h, i.e. Ψ_{hdt}. [27] This means that beliefs are specific to a dyad reflecting how media outlets in a home country choose to feature events that unfold in a destination country.

As shown in Figure 3 we allow news reporting to react to the underlying state so that

[27] This can be interpreted as a form of availability heuristic where news at home is more available and salient.

agents can use reporting to update beliefs. We posit a statistical model to represent the data generating process driving news coverage of a destination d by the media from a traveler home country h. However, in contrast to events, where we could assume a simple normal distribution with state-specific parameters, we now need a statistical model consistent with state-specific variation in both general background news coverage and news coverage of violent events.

Let overall news coverage, N_{hdt}, and violent news coverage, B_{hdt}, be measured as the number of news articles about destination d in origin country h at date t. They are assumed to follow a negative binomial distribution parameterized by η_s capturing the fraction of news articles that report on violence. We use the negative binomial distribution to capture the fact that news reporting of B_{hdt} relative to N_{hdt} has fat tails. The distribution could be justified by a stopping rule for media consumption but we follow it purely for pragmatic reasons. If $\hat{\eta}_1 > \hat{\eta}_0$, there is a higher relative frequency of "violent news" relative to the number of news articles when a country is dangerous so that the extent of violent news coverage can serve as a signal about the state of danger at a destination.

We can write the density function for each of the two underlying latent states s as:

$$f\left(B_{hdt}, N_{hdt} \mid s\right) = \binom{N_{hdt}}{B_{hdt}} \left(\eta_s\right)^{B_{hdt}} \left(1-\eta_s\right)^{\left(N_{hdt}-B_{hdt}\right)}. \tag{6}$$

The model implies that beliefs will increase if B_{hdt} increases relative to N_{hdt}. However, the magnitude of this effect depends on the parameters $\hat{\eta}_1$ and $\hat{\eta}_0$.

We again assume that agents know the parameters of the model but we need to derive estimates for the parameters $\hat{\eta}_1$ and $\hat{\eta}_0$. With known states, $\hat{\eta}_s$ would simply be the frequencies of B_{hdt}/N_{hdt} during periods of danger and safety. However, since the underlying states are latent, our estimates of $\hat{\eta}_s$ is based on the following weighted average:

$$\hat{\eta}_{s,d} = \frac{\sum_{hdt} \Pi_{dt} \left[\frac{B_{hdt}}{N_{hdt}}\right]}{\sum_{hdt} \Pi_{dt}}, \tag{7}$$

where the summation is over all dyads and time periods. Hence, we are weighting the frequencies by the probability of the occurence of each state from the events-based estimates.

We estimate that $\hat{\eta}_1 = 0.022$ and $\hat{\eta}_0 = 0.002$, i.e. the share of violent news in all news is around 2.2% for the dangerous state and 0.2% for the safe state. An important implication of this is that some dyads with reporting on violence against tourists are regarded as safe because a lot of *other* news appears about the destination in the same month.

Using the model, news-based beliefs are updated using a likelihood ratio based on the densities in (6) given by:

$$\lambda\left(B_{hdt}, N_{hdt}\right) = \frac{f(B_{hdt}, N_{hdt}|0)}{f(B_{hdt}, N_{hdt}|1)} \tag{8}$$

which now depends only on news coverage (B_{hdt}, N_{hdt}) in a dyad. The probability that a country is perceived as dangerous according to the naïve beliefs is then given by

$$\pi_{hdt} = \Pr(s_{dt} = 1 \mid \Psi_{hdt}, \hat{p}_d, \hat{q}_d, \hat{\eta}_0, \hat{\eta}_1).$$

which evolves according to the Bayesian recursion:

$$\pi_{hdt} = \frac{E_{t-1}\left[\pi_{hdt}\right]}{E_{t-1}\left[\pi_{hdt}\right] + \left[1 - E_{t-1}\left[\pi_{hdt}\right]\right]\lambda\left(B_{hdt}, N_{hdt}\right)}$$

where, as before, $E_{t-1}\left[\pi_{hdt}\right]$ is derived from the Markov chain governing the underlying categorical state.[28]

Important Characteristics of News-based Beliefs When assessing the model's empirical success, it should be noted that the model is quite specific in the way that it allows news coverage to matter and hence there is nothing built in that will guarantee a better fit to patterns in the data. Thus, what we present below is a non-trivial test of whether the model can capture key features of the data. That said, there are two key features of news-based belief formation that are worth highlighting and bear directly on how the media multiplier is estimated.

First, many of the destinations in our sample are not covered at all for months by some home countries even if they suffer from political violence. Instead, after specific news events, reporting will spike and then calm down immediately and so will news-based beliefs. Such intense news reporting on a single event will leave a strong mark on agents with news-based beliefs. A dyad in which there is no reporting interrupted by some news coverage of violence will jump back and forth between strong beliefs of safety and danger.

Second, observing *non-violent news* reporting about a country leads to updating that a country is safe. Since $\hat{\eta}_1 > \hat{\eta}_0$, news-based beliefs attach a higher probability to a place being dangerous if B_{hdt} increases. However, since $\hat{\eta}_0 > 0$, news coverage of

[28]Appendix Figure A10 reports the posterior distribution of π_{hdt} for months that we classify as relatively dangerous ($\Pi_{dt} > \frac{1}{2}$) and safe ($\Pi_{dt} \leq \frac{1}{2}$) destination/months. Both distributions have full support but the density of π_{hdt} has a much thicker tail during dangerous months.

violence against tourists does not immediately imply updating towards a destination being dangerous. Context matters and the non-linearity of the model implies that there is a natural "tipping point" in equation (8) as a function of B_{hdt} relative to N_{hdt} at which news-based beliefs attach a larger probability to a country being dangerous.

Both of these features allow the model to account for the sharp changes in beliefs in response to news that are needed to explain patterns in the spending data. However, given the model structure, these sharp reactions are not driven by a different model of what danger means but by the way that information is processed to update these beliefs.

5.2 Estimating the Media Multiplier

We use movements in the spending data to estimate the weights on news-based beliefs and event-based beliefs. Since travel is usually booked in advance, we also allow for a lagged effect of beliefs of up to nine months, estimating the lag structure that offers the best fit to the data. Therefore, in total we estimate ten parameters (weights) using the following empirical specification based on (4):

$$y_{hdt} = \alpha_{hd} + \alpha_{ht} + \alpha_{dm} + \xi \sum_{\tau=0}^{-9} \omega_\tau \left(\chi \Pi_{dt-\tau} + (1-\chi)\, \pi_{hdt-\tau} \right) + \varepsilon_{hct}. \tag{9}$$

where χ is the weight on beliefs formed by observing violence. This gives rise to a natural formal representation of the media multiplier as $1/\chi$. In our model, it captures the relative importance of news-based beliefs in amplifying the overall spending response up and above the effect coming from the reaction based on violence alone. Formally $\chi + (1-\chi)$ divided by χ which gives $1/\chi$. The total effect is by $1/\chi$ larger when event-based beliefs get complemented by news-based beliefs.

We interpret (9) as modelling the beliefs of a "representative" traveler. But χ and ω_τ could also reflect heterogeneity in the population in terms of booking behavior or news-based and event-based beliefs among sub-groups of travelers. In that case, the parameter χ would correspond to the share of event-based travelers.

The model-based approach differs from equation (3) in three ways. First, the belief estimates, (Π_{dt}, π_{hdt}), are both heterogeneous across countries and non-linear in their response to violence and news data. Second, they depend on the entire history of violence rather than on one-period lagged values and third, they allow for a lag structure to reflect the timing of travel decisions. At the same time the weights (ω, χ) reflect the importance of different kinds of latent beliefs for overall spending.

Our approach allows us to estimate ξ which captures the effect of latent beliefs on card activity. This is an important feature of the model-based specification as it provides a precise meaning to the estimate. Table 2 explores how well our different estimates of latent beliefs $\{\Pi_{dt}, \pi_{hdt}\}$ explain variations in aggregated tourist spending. In column (1) of Table 2, we report the relationship between aggregated spending and Π_{dt-1}. Formally, this implies in relation to equation (9) that we suppose that $\chi = 1$ and $\omega_{-1} = 1$. On average, spending falls by about 20 percent when beliefs Π_{dt-1} increase from 0 to 1, i.e. when a destination goes from being viewed as completely safe to completely dangerous based on event-based beliefs. This magnitude is in the same ballpark as the reduced-form results reported in Table A3. But recall that Π_{dt-1} moves as in Figure A7 and therefore reacts much more strongly to some changes in violence than others. For example, given the Markov chain estimates in Appendix Table A10, some levels of violence that are associated with safety in Egypt would represent a dangerous episode in Tunisia.[29] In Appendix Table A11 we show that, in this way, our model is able to capture an important part of the heterogeneity in the response to violence so that the response to beliefs is relatively homogeneous across destinations.

Column (2) focuses on responsiveness of spending to news-based beliefs, π_{hdt-1}, i.e. imposing the (artificial) case of $\chi = 0$ and $\omega_{-1} = 1$. In other words, we assume that all travelers form news-based beliefs and book their travel one month in advance. We now get a fall of 0.364 log points or 30 percent if the news-based belief that a destination country is dangerous within a dyad moves from 0 to 1. In column (3) we estimate π_{hdt} based on news reporting on violent events more generally, rather than those targeting tourists. And although, as in the reduced-form results, the impact of news reporting is somewhat smaller in magnitude, it moves in the same direction.

However, $\chi = 0$ is not appealing. It represents a case where the media "tail is wagging the dog", i.e. where there are economic effects from news coverage even if there is no actual threat or violence. Therefore, to estimate the media multiplier we combine news-based and violence-based beliefs in the same regression as we do in column (4). Now we find a strong amplifying effect of news coverage with the coefficient on π_{hdt} being quantitatively large and statistically significant. This is because aggregate spending tends to follow media coverage even when the latter is not closely related to the event-based risks highlighted in Figure A7. We regard this as empirical

[29] This is in line with the idea in Becker and Rubinstein (2011) that there is investment in coping with fear which would lead to long-term adjustments to existing levels of violence.

support for the specific form of an *availability heuristic* captured by our model and it also provides a theoretical interpretation of the media multiplier reflected in π_{hdt}.

Column (5) uses the estimate of $\sum_{\tau=0}^{-9} \hat{\omega}_\tau \hat{P}_{hdt-\tau}$ after fitting the entire model to the spending data. For this, we first find the optimal weights (χ and ω_τ) that give the best fit to the spending data using a grid search over the weights in equation (9) to maximize goodness of fit.[30] We find that $\hat{\omega}_0 = \hat{\omega}_1 = 0.2$ so that 20 percent of spending is driven by contemporaneous beliefs and 20 percent are coming from the first lag. After that, the weight based on the best fit falls (the weight sequence is 0.15, 0.1, 0.1, 0.1, 0.05, 0.05, 0.05). Our estimate of the media multiplier is $1/\hat{\chi} = 2.5$, i.e. the effect of an event is more than doubled if it is accompanied by intense media coverage.

Putting these estimates together, we find that "optimal" weighted average reported in column (5) suggests that if all tourists switched their categorical beliefs that a destination is dangerous from zero to one, then spending would fall by about 53 percent (0.75 log points). Of course, this is somewhat extreme thought experiment since it would require a sequence of negative events and persistent, intense negative reporting. But this is relevant for destinations in which news coverage and casualty data lead to maintained beliefs of danger over several months as is would be the case in ongoing terror campaigns in Afghanistan or Iraq or as it has been the case for most countries in the COVID-19 crisis. We will use the estimated weights and the resulting estimate of $\hat{\xi}$ in column (5) to quantify the separate impact of news on spending below.

In summary, we find that card activity and beliefs are strongly associated, even after controlling for dyad and origin/time fixed effects. In the reduced form section we have also shown that this association shows no clear pre-trends either in reporting or in card activity to the event date. What remains are concerns with the interpretation of our results where our model could be capturing differential risk perceptions or government responses. However, we can show that our model is able to capture a lot of heterogeneity in economic outcomes both through beliefs. For example, we find no evidence for heterogeneity of the spending response to π_{hdt} with respect to distance to a destination or the share of Muslims in an origin country (see Appendix Table A12). The only exception is heterogeneity with respect to the freedom of the local press which we turn to in the next sub-section.

[30] See Appendix B for more details.

5.3 The Media Multiplier and Press Freedom

Our results suggest that up to sixty percent of the economic effect of violent events, many of which affected by terrorism, would not arise in the absence of news coverage as the media multiplier would be absent. The fact that many countries around the world suppress media provides an interesting natural experiment and a further test of the role of news coverage on spending.

In terms of our model, suppression of information through news censorship would affect use of news-based beliefs π_{hdt} by agents implying that π_{hdt} should have less predictive power in countries with censored media. To test this, we collected data from Reporters without Borders (RSF), which score countries according to their press freedom. We use this to split traveler origin countries depending on whether they are above or below the median score in our sample of countries which allows us to test whether the variable π_{hdt} enters heterogeneously across the free and censored countries.

(Table 3)

The results are presented in Table 3. In column (1) we show that the effect of π_{hdt} on spending increases by a factor of two in countries with a free press compared to countries without. Column (2) shows that this finding is not driven by press agencies which we have in the dataset from six countries including Russia and China. Also striking is that this amplification of the effect is *not* found for event-based beliefs, Π_{dt}, as shown in column (3). This supports the view that it is belief formation through the (free) media that is driving our results and is not therefore a general propensity to respond to an event. Columns (4) to (6) present similar findings when we use the total number of accounts as the dependent variable.

These findings further corroborate the idea that news-based beliefs, π_{hdt}, capture the impact of signals transmitted through local news reporting and give rise to an amplification effect of news coverage that we have referred to as the media multiplier. That said, we would caution against arguing that suppression of information is normatively desirable in the face of violent events if the latter are prone to overreaction. Public trust in signals coming from the news media is a useful asset in most circumstances. Also, there are many beneficial aspects of societies with free media that transcend concerns of the kind that we have focused on.

5.4 The Media Multiplier and Background News

A key feature of news-based belief formation is that non-violent or background news coverage, $N_{hdt} - B_{hdt}$, affects beliefs. Figure 4 illustrates the role of these background news by considering the spending response to a violent event (occurring at date 0) which changes event-based beliefs from $\Pi_{d,-1} = 0$ to $\Pi_{d,0} = 1$ for one month. The event-based part from the model estimated in columns (5) of Table 2 then yields the spending response represented by the dashed line in Figure 4. According to this, the immediate effect would be that the 20 percent of tourists which react immediately to danger ($\omega_0 = 0.2$) do not travel to the destination and tourism spending would fall by around four percentage points.[31] The dashed line shows the persistent effect on tourism spending, i.e. the effect of $\Pi_{d,0} = 1$ falls only slowly because most tourists book their travel in advance.

We then show the *additional* impact of the media multiplier via a change in news-based beliefs in a scenario in which the violent event is accompanied by reporting $B_{hd0} = 1$. To illustrate the importance of other background news, $N_{hdt} - B_{hdt}$, we contrast two levels of background reporting $N_{hd0} \in \{0, 100\}$. The black solid line shows the effect of the media multiplier with no other background news, $N_{hdt} = 0$. The result is a spending response of over 8.4 percentage points. In other words, the economic impact of the violent event on spending more than doubles due to the media multiplier. Again, because tourists book their holidays in advance, this effect persists.

However, with $N_{hd0} = 100$, the news about tourist violence is "drowned out" by other news and the impact of negative reporting, $B_{hd0} = 1$, is reduced dramatically. The maximum spending effect is now only 5.6 percent. This is because news-based beliefs update according to the density in equation (6) and reporting of $B_{hd0} = 1$, $N_{hd0} = 100$ is relatively likely coming from a safe environment. The impact of the media multiplier therefore depends not only on the estimated parameters of the model but also on the overall news landscape.

This vividly illustrates the importance of background news in "distracting" or "putting things into perspective" for potential travelers when they rely on news coverage to form their beliefs. The behavior of beliefs in our model is explained by the fact that, to the extent that they are news-based, travelers do not learn from wider news sources, taking local news coverage at face-value. Their beliefs then behave in a way similar to what would happen if tourists were to update using the model of Bordalo et al. (2016) in

[31]The exact calculation is $0.2 \times 0.4 \times 0.53 = 4.2\%$.

which destination countries are stereotyped as dangerous if they are covered by bad news without any other background news coverage.

5.5 The Economic Consequence of the Media Multiplier

To quantify the economic consequences of the media multiplier, we contrast the average effect on tourism activity due to a violent event operating exclusively through event-based beliefs,

$$\hat{\xi} \sum_{\tau=0}^{-9} \hat{\omega}_\tau \hat{\chi} \hat{\Pi}_{dt-\tau}, \tag{10}$$

with the overall effect including news-based beliefs,

$$\hat{\xi} \sum_{\tau=0}^{-9} \hat{\omega}_\tau \left(\hat{\chi} \hat{\Pi}_{dt-\tau} + (1-\hat{\chi})\, \hat{\pi}_{hdt-\tau} \right). \tag{11}$$

Figure 5 Panels A and B show the effect that we would expect if all potential travelers held event-based beliefs as the grey-line contrasted with the overall effect represented by the black line. The left-hand panel in Figure 5 shows that, for Tunisia, a large part of the variability in tourism spending comes from changes in event-based beliefs, irrespective of news coverage. Nevertheless, there is a visible news effect and, in 2015, it alone accounts for a spending decline of about 15 percent. This illustrates the media multiplier at work in Tunisia during that year. In the right-hand panel we show the same thing for Egypt and also find a media multiplier, although not as large as the effect driven by event-based beliefs. This further illustrates the crowding out effect in Figure 4.

The estimates in Figure 5 panels A and B suggest material losses to the economy in all four countries that we study. The World Bank reports that tourism receipts in 2010 were 3.48 Billion USD in Tunisia, 5.6 Billion USD in Israel, 13.63 Billion USD in Egypt and 26.3 billion USD in Turkey. Back of the envelope calculations based on the estimates reported in this section indicate losses between 2011 and 2016 of over 35 billion USD due to violence with more than 10 billion USD being due to the independent effect of negative news reporting.[32] Egypt and Tunisia are, for example, predicted to have recovered from the negative shocks towards the end of the sample period. However, it should be kept in mind that our estimates are conservative in that they rely on variation within the region, so that broader trends away from all countries are captured by the origin/time fixed effects.

Figure 5 Panel C explores how well the model-based approach fits the data. The

[32]For calculations see the Appendix C.

left-hand panel illustrates the predicted effect, averaged over all potential tourists' origin countries, for Tunisia based on (9) and compares it to the average of the residuals in the spending data after having taking out the fixed effects. The model captures both the early decline and recovery at the beginning of the Arab spring. However, the most striking observation is for 2015, where it accurately captures both the decline and recovery in spending.

The model-based prediction compares favorably with that from the reduced-form model, as shown in Panel D of Figure 5. For this comparison, we use the estimated effects, $\xi_1 n_{hdt-1} + \xi_2 v_{hct-1}$ averaged over origin countries based on Table 1. Although some of the broad patterns are visible, the fit to the data is notably inferior to that in the left-hand panel, failing to reproduce the timing and magnitude of the fluctuations in spending. This is particularly noticeable during the Arab spring and getting the timing right in sudden expenditure declines in 2015. Thus, Figure 5 supports the utility of an approach which models patterns of belief formation alongside a reduced-form approach. Moreover, it is an interesting insight that a statistical model where beliefs are at least partially news-based in the way we are modelling them can provide such a good account of spending patterns in the data.

6 Concluding Comments

This paper contributes to our understanding of the power of the media in influencing economic decisions by introducing the idea of a media multiplier - which measures how agents' responses to news reporting amplifies the effect of an event. The theory-based measurement approach is in a specific model of belief formation. Our empirical approach exploits a unique dyadic data set on tourist spending and news coverage of violent events in five destination countries. To implement the approach, we trained a machine learning model to spot news coverage of violence. We find that sixty percent of the weight in terms of updating beliefs comes from country-specific news reporting although the response will depend on the intensity and persistence of reporting.

It would be useful in future to use lab experiments to explore in greater detail how news media coverage plays out in influencing behavior. It would also be interesting to study diversion of spending to other locations by using spending data from a broader range of destination countries. Our focus has been on extreme violent events since the news effect is easier to identify. But other events, such as crime perpetrated on tourists, would be interesting to study, particularly in the Caribbean and Latin America, where anecdotal evidence suggest that this could be important in affecting tourism.

Our results suggest that news cycles around negative events can have adverse economic effects, a good part of which could be driven by media “sensationalism” where gory images and dramatic stories increase attention and sales. However, we would caution against interpreting our results as saying that the existence of a media multiplier justifies curbs on press freedom. This is relevant to recent events surrounding COVID-19, when the media are accused of inducing apparently irrational economic responses. Besley (2020) argue that media suppression could be a factor in under-reporting deaths from COVID-19 and the hope of governments could be that this will lower the economic impact of the pandemic. But, in a wider context, media suppression is dangerous and there are persuasive arguments that media suppression reduces political accountability (see Besley and Prat, 2006). It is also clear from our results that economic agents can adjust and will rely much less on suppressed media when forming beliefs.

The media multiplier is also relevant to debates about the potential for a reversal in globalization. Given that increased international travel has been a significant component of global integration, determinants of travel are important not just in terms of generating traditional economic gains from trade but also in fostering greater cross-cultural understanding. To the extent that security concerns increase the perceived costs of travel, media reports may therefore have an important impact on this aspect of international integration. And the results in this paper suggest that the way media chooses to report these risks has a role to play in this process. They also provide a way of trying to calibrate the size of the impact to gain a quantitative estimate.

This will also be important for international investment, where it is often claimed that perceptions matter. Negative coverage of the prospects for African countries, in particular, create a climate of opinion among corporate boards and shareholders that could affect the allocation of FDI. This is particularly poignant in an era where social media and the potential for fake news is attracting increasing attention in a region which is economically not well-integrated (Rouis and Tabor, 2012). How far news coverage and possible biases have real aggregate economic consequences is ripe for further investigation.

More generally, we believe that the agenda of fitting models to data to explore the channels of influence is an important step forward in studying the economic impact of the media. In particular, having parameters from a model can be useful in thinking through how media coverage could play out in other contexts.

References

Abadie, A. and J. Gardeazabal (2003). The economic costs of conflict: A case study of the Basque Country. *American Economic Review 93*(1), 113–132.

Acemoglu, D., S. Naidu, P. Restrepo, and J. A. Robinson (2019). Democracy Does Cause Growth. *Journal of Political Economy 127*(1), 47–100.

Adena, M., R. Enikolopov, M. Petrova, V. Santarosa, and E. Zhuravskaya (2015). Radio and the Rise of The Nazis in Prewar Germany. *The Quarterly Journal of Economics 130*(4), 1885–1939.

Alfano, M. and J.-S. Görlach (2019). Terrorism, Education, and the Role of Expectations: Evidence from al-Shabaab Attacks in Kenya.

Amodio, F. and M. Di Maio (2017). Making Do With What You Have: Conflict, Input Misallocation and Firm Performance. *Economic Journal*, 1–54.

Arezki, R., V. A. Ramey, and L. Sheng (2017). News Shocks in Open Economies: Evidence from Giant Oil Discoveries. *The Quarterly Journal of Economics 132*(1), 103–155.

Azeredo da Silveira, R. and M. Woodford (2019). Noisy Memory and Over-Reaction to News. *AEA Papers and Proceedings 109*, 557–561.

Baker, S. R., N. Bloom, and S. J. Davis (2016). Measuring Economic Policy Uncertainty. *The Quarterly Journal of Economics 131*(4), 1593–1636.

Becker, G. S. and Y. Rubinstein (2011). Fear and the Response to Terrorism: An Economic Analysis. *CEP Discussion Paper*, 1–90.

Becker, S. O. and L. Pascali (2019). Religion, Division of Labor, and Conflict: Anti-Semitism in Germany over 600 Years. *American Economic Review 109*(5), 1764–1804.

Besley, T. (2020). The Political Economy of Lockdown: Does Free Media Make a Difference? *mimeo* (June).

Besley, T. and R. Burgess (2002). The Political Economy of Government Responsiveness: Theory and Evidence from India. *The Quarterly Journal of Economics 117*(4), 1415–1451.

Besley, T., T. Fetzer, and H. Mueller (2015). The Welfare Cost of Lawlessness: Evidence from Somali Piracy. *Journal of the European Economic Association 13*(2), 203–239.

Besley, T. and H. Mueller (2012). Estimating the Peace Dividend: The Impact of Violence on House Prices in Northern Ireland. *American Economic Review 102*(2), 810–833.

Besley, T. and A. Prat (2006). Handcuffs for the grabbing hand? Media capture and government accountability. *American Economic Review 96*(3), 720–736.

Bloom, N. (2009). The Impact of Uncertainty Shocks. *Econometrica 77*(3), 623–685.

Bordalo, P., K. Coffman, N. Gennaioli, and A. Shleifer (2016). Stereotypes. *The Quarterly Journal of Economics* (October), 1753–1794.

Bordalo, P., N. Gennaioli, Y. Ma, and A. Shleifer (2018). Over-reaction in Macroeconomic Expectations. *NBER Working Paper* (March).

Brodeur, A. (2018). The effect of terrorism on employment and consumer sentiment: Evidence from successful and failed terror attacks. *American Economic Journal: Applied Economics 10*(4), 246–282.

Brückner, M. and E. Pappa (2015). News Shocks in the Data: Olympic Games and Their Macroeconomic Effects. *Journal of Money, Credit and Banking 47*(7), 1339–1367.

Burchardi, K. B., T. Chaney, and T. A. Hassan (2019). Migrants, Ancestors, and Foreign Investments. *The Review of Economic Studies 86*(4), 1448–1486.

Bursztyn, L., D. Cantoni, P. Funk, and N. Yuchtman (2017). Polls, the Press, and Political Participation: The Effects of Anticipated Election Closeness on Voter Turnout. *NBER Working Paper*.

Carvalho, V. M., J. R. Garcia, S. Hansen, T. Rodrigo, and J. V. R. Mora (2020). Tracking the COVID-19 Crisis with High-Resolution Transaction Data. *mimeo*.

Combs, B. and P. Slovic (1979). Newspaper Coverage of Causes of Death. *Journalism Quarterly 56*(4), 837–849.

DellaVigna, S. and E. Kaplan (2007). The fox news effect: Media bias and voting. *Quarterly Journal of Economics 122*(3), 1187–1234.

Donaldson, D. (2018). Railroads of the Raj: Estimating the Impact of Transportation Infrastructure. *American Economic Review 108*(4-5), 899–934.

Durante, R., P. Pinotti, and A. Tesei (2019). The Political Legacy of Entertainment TV. *American Economic Review 109*(7), 2497–2530.

Durante, R. and E. Zhuravskaya (2018). Attack When the World Is Not Watching? US News and the Israeli-Palestinian Conflict. *Journal of Political Economy 126*(3), 1085–1133.

Eggers, A. and A. Fouirnaies (2014). The Economic Impact of Economic News. *mimeo*.

Eisensee, T. and D. Strömberg (2007). News droughts, news floods, and US disaster relief. *Quarterly Journal of Economics 122*(2), 693–728.

Enikolopov, R., M. Petrova, and E. Zhuravskaya (2011). Media and political persuasion: Evidence from Russia. *American Economic Review 101*(7), 3253–3285.

Faber, B. and C. Gaubert (2016). Tourism and Economic Development: Evidence from Mexico's Coastline. *American Economic Review 109*(6), 2245–2293.

Fetzer, T. (2019). Can Workfare Programs Moderate Conflict? Evidence from India. *Journal of the European Economics Association*.

Feyrer, J. (2019). Trade and Income – Exploiting Time Series in Geography. *forthcoming, American Economic Journal: Applied Economics* (603), 1–43.

Fiske, S. T. (1998). Stereotyping, prejudice, and discrimination. In &. G. L. D. T. Gilbert, S. T. Fiske (Ed.), *The handbook of social psychology*, pp. 357–411. New York: McGraw-Hill.

Frankel, J. A. and D. Romer (2008). Does Trade Cause Growth? *American Economic Review 89*(3), 379–399.

Fryer, R. and M. O. Jackson (2008). A categorical model of cognition and biased decision making. *B.E. Journal of Theoretical Economics 8*(1).

Gentzkow, M. (2006). Television and Voter Turnout. *Quarterly Journal of Economics 121*(3), 931–972.

Glick, R. and S. Leduc (2012). Central bank announcements of asset purchases and the impact on global financial and commodity markets. *Journal of International Money and Finance 31*(8), 2078–2101.

Hamilton, J. (1990). Analysis of Time Series Subject to Changes in Regime. *Journal of Econometrics 45*, 39–70.

Handel, B. and J. Schwartzstein (2018). Frictions or mental gaps: What's behind the information we (don't) use and when do we care? *Journal of Economic Perspectives 32*(1), 155–178.

Hansen, S., M. McMahon, and A. Prat (2018). Transparency and deliberation within the FOMC: A computational linguistics approach. *Quarterly Journal of Economics 133*(2), 801–870.

Hassan, T. A., S. Hollander, L. van Lent, and A. Tahoun (2019). Firm-Level Political Risk: Measurement and Effects. *The Quarterly Journal of Economics 134*(4), 2135–2202.

Iaria, A., C. Schwarz, and F. Waldinger (2018). Frontier Knowledge and Scientific Production: Evidence from the Collapse of International Science. *The Quarterly Journal of Economics 133*(2), 927–991.

Japkowicz, N. and S. Stephen (2002). The class imbalance problem: A systematic study1. *Intelligent Data Analysis 6*(5), 429–449.

Jetter, M. (2017). The effect of media attention on terrorism. *Journal of Public Economics 153*, 32–48.

Jha, S. and M. Shayo (2019). Valuing Peace: The Effects of Financial Market Exposure on Votes and Political Attitudes. *Econometrica 87*(5), 1561–1588.

Manacorda, M. and A. Tesei (2020). Liberation Technology: Mobile Phones and Political Mobilization in Africa. *Econometrica 88*(2), 533–567.

Manela, A. and A. Moreira (2017). News implied volatility and disaster concerns. *Journal of Financial Economics 123*(1), 137–162.

Melitz, M. J. and D. Trefler (2012). Gains from trade when firms matter. *Journal of Economic Perspectives 26*(2), 91–118.

Mueller, H. and C. Rauh (2017). Reading Between the Lines: Prediction of Political Violence Using Newspaper Text. *American Political Science Review 2*(October), 1–18.

Neumayer, E. (2004). The Impact of Political Violence on Tourism. *Journal of Conflict Resolution 48*(2), 259–281.

Pedregosa, F., V. Michel, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, J. Vanderplas, D. Cournapeau, G. Varoquaux, A. Gramfort, B. Thirion, O. Grisel, V. Dubourg, A. Passos, M. Brucher, M. Perrot, and E. Duchesnay (2011). Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research 12*, 2825–2830.

Persson, T. and G. Tabellini (2009). Democratic Capital: The Nexus of Political and Economic Change. *American Economic Journal: Macroeconomics 1*(2), 88–126.

Prat, A. and D. Strömberg (2011). The Political Economy of Mass Media. *CEPR Discussion Paper*.

Ramey, V. A. (2011). Identifying government spending shocks: It's all in the timing. *Quarterly Journal of Economics 126*(1), 1–50.

Rouis, M. and S. R. Tabor (2012). *Regional Economic Integration in the Middle East and North Africa*.

Shen, O., H. Al-Jamaly, M. Siemers, and N. Stone (2018). Charting Death: Reality vs Reported Data. *mimeo*, 1–10.

Snyder, J. M. and D. Strömberg (2010). Press coverage and political accountability. *Journal of Political Economy 118*(2), 355–408.

Stromberg, D. (2004). Radio's Impact on Public Spending. *The Quarterly Journal of Economics 119*(1), 189–221.

Strömberg, D. (2015). Media and Politics. *Annual Review of Economics 7*(1), 173–205.

Tversky, A. and D. Kahneman (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology 5*(2), 207–232.

WTTC (2018). Travel & Tourism Power and Performance - September 2018. pp. 1–23.

Zhao, Y., L. Zhang, and S. Yuan (2018). The effect of media coverage on threshold dynamics for a stochastic SIS epidemic model. *Physica A: Statistical Mechanics and its Applications 512*, 248–260.

Zhu, J.-Q., A. N. Sanborn, and N. Chater (2020). The Bayesian sampler: Generic Bayesian inference causes incoherence in human probability judgments. *Psychological Review*, 1–26.

Figure 1: Overall aggregated spending patterns of British- and German-issued cards in Tunisia in the wake of the Sousse attack)

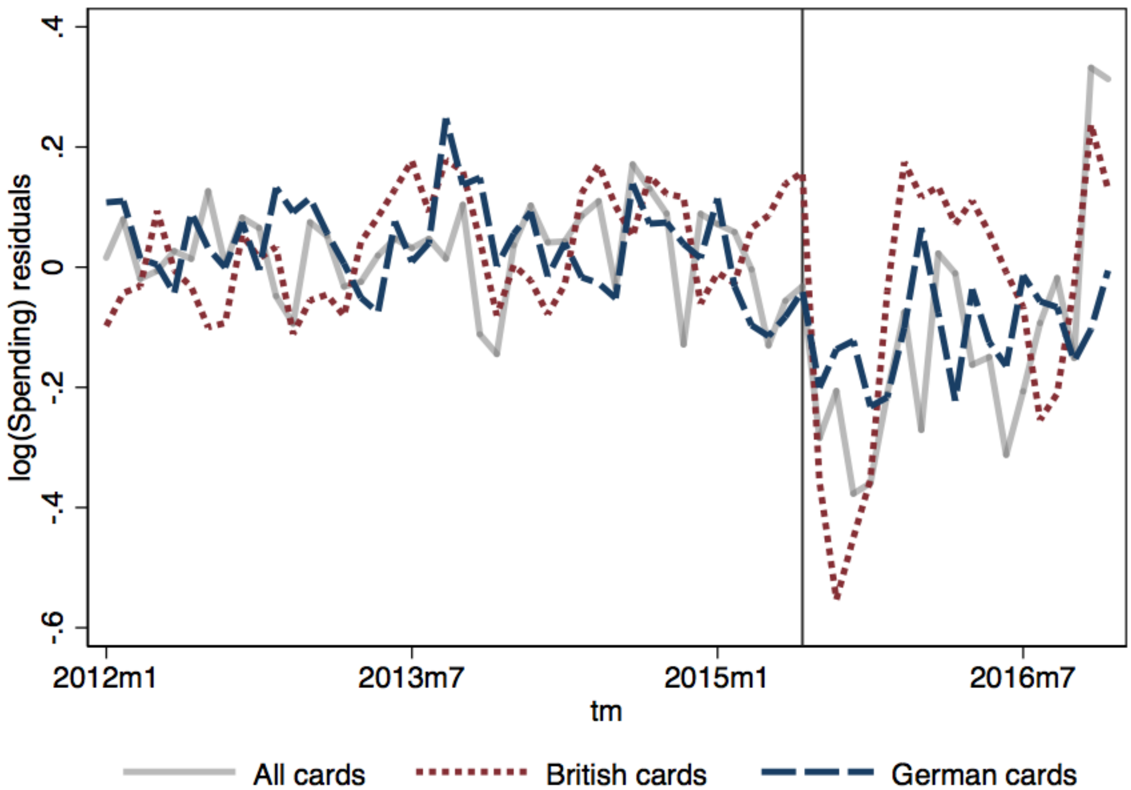

Notes: The solid black line presents residuals of a regression of the log of aggregated card spending removing destination country fixed effects as well as destination-specific seasonality. The other two lines plot residuals of a regression of the log card spend for German- and British-issued cards in Tunisia over time, having removed dyad fixed effects, issuing country by time fixed effects and destination by month fixed effects. The drop in tourism spending is markedly larger for British-issued cards.

Figure 2: Lead- and lagged effect of violence targeted against tourist's on tourism spending

Panel A: News on tourist being targeted

Panel B: News on any fatal violence

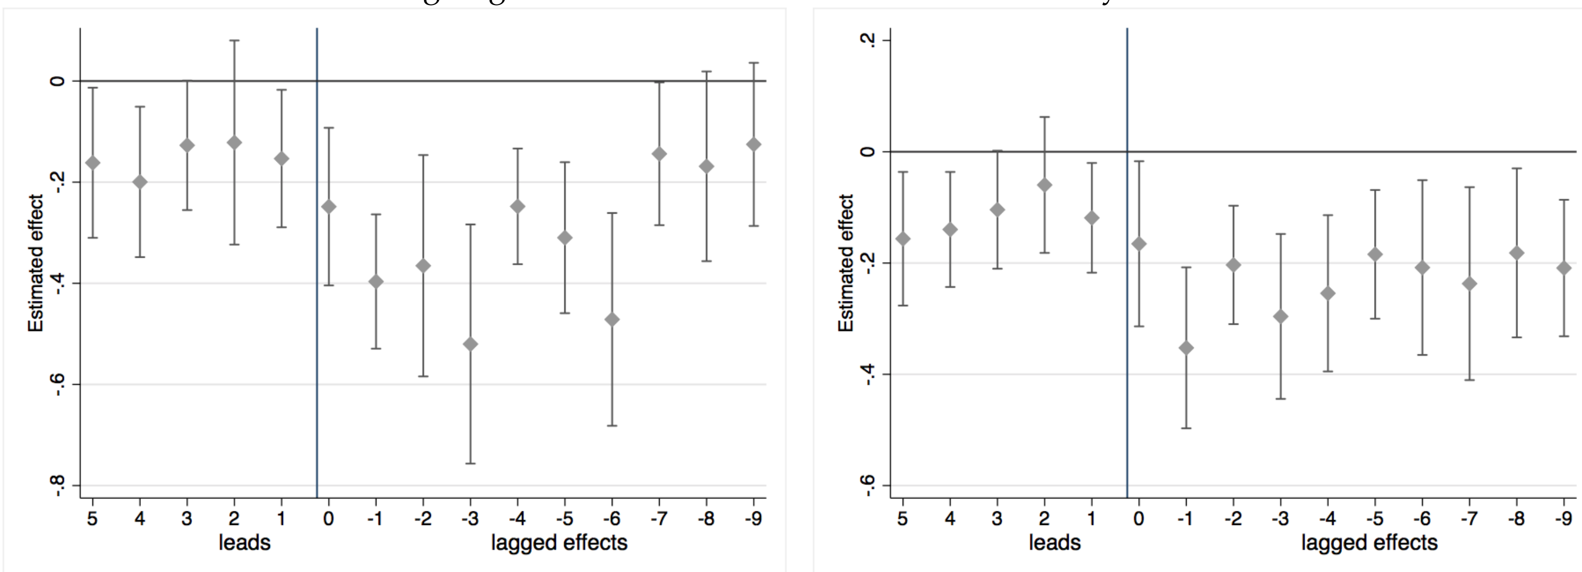

Notes: Figure plots the results from estimating a linear regression controlling for dyad fixed effect, issuing-country by time fixed effect and destination by month seasonality. The coefficients that are reported are the estimates on different leads and lags of the violence reporting measure on the log-value of tourism spending. The news reporting measure in panel A measures the share of articles on a dyad and month that are classified as tourism having been the target of violent events. In panel B the news measure captures share of articles in a dyad and month that that are classified as indicating any violent event involving fatalities. 95% confidence bands obtained from clustering the data at the dyad level are indicated.

Figure 3: Graphical representation of data generating process

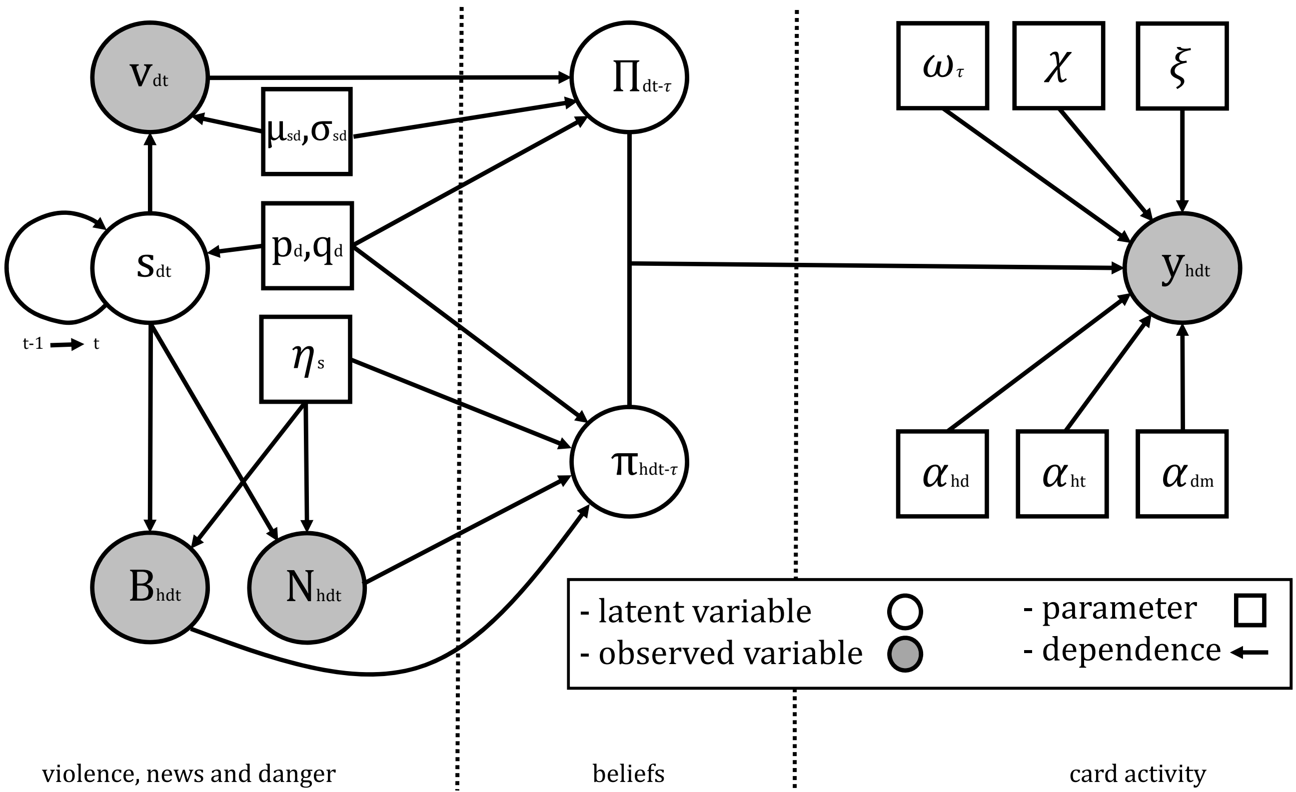

Notes: The figure plots the generative model using an adaptation of plate notation common in machine learning highlighting the observable and unobservable variables, the parameter spaces, and the dependence relationships.

Figure 4: Impact of the Media Multiplier at Different Levels of Background News

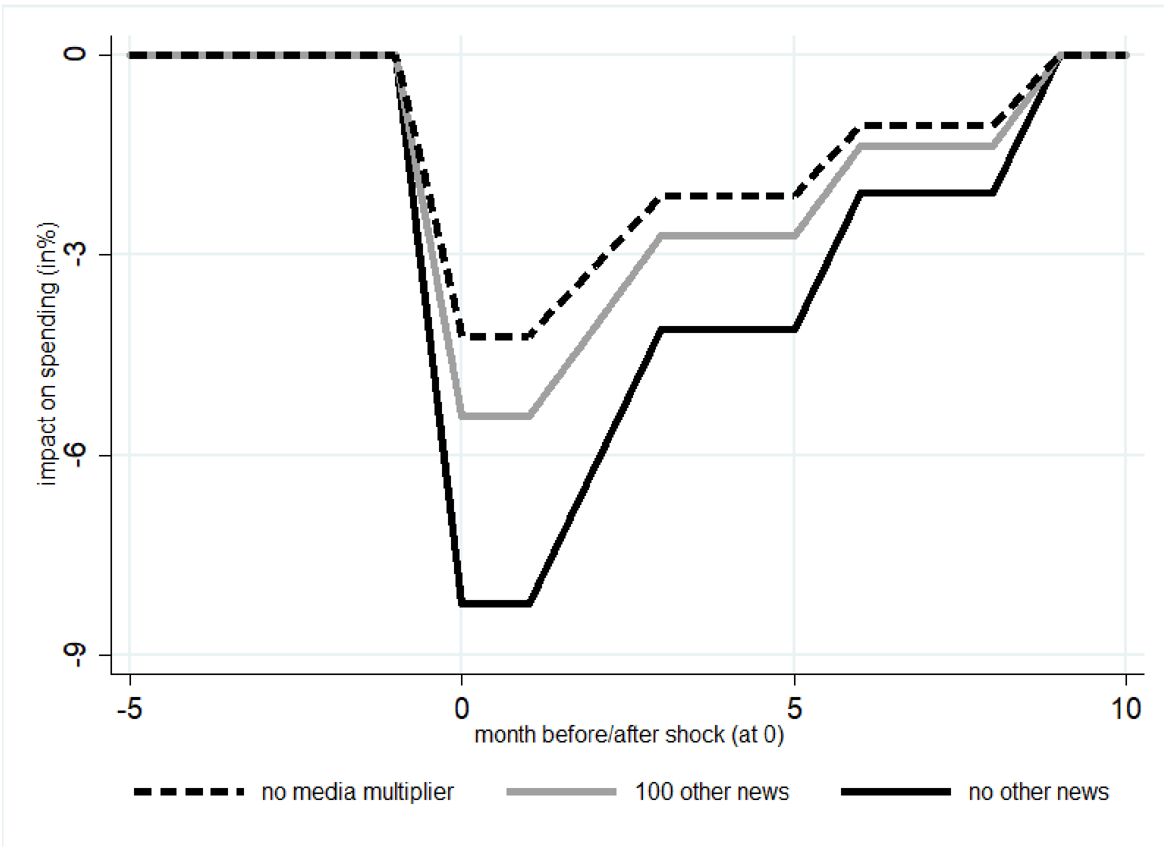

Notes: The figure plots out the impact of a single violent event occurring at month $t = 0$ on tourism spending across subsequent months on one dyad under different scenarios concerning the news environment in the origin country of that dyad. The effect of the violent event without the media multiplier is provided as the dashed line. This is a reference point capturing the impact that is attributable to "Event-based" beliefs. The other lines present the total impact on card spending that incorporates the additional effect of the media multiplier. This effect manifests itself through the impact of news items on "news-based" beliefs. We contrast two information environments: one where there are 100 other background news items each month covering events that are unrelated to violence (solid grey line). The solid black line presents the total effect if there is no other media coverage that could attenuate the effect of the media multiplier. Accordingly, the effect of the event on spending more than doubles.

Figure 5: Model-based Effects in Tunisia and Egypt

Comparison of model-based predictions for Egypt and Tunisia

Panel A: Tunisia

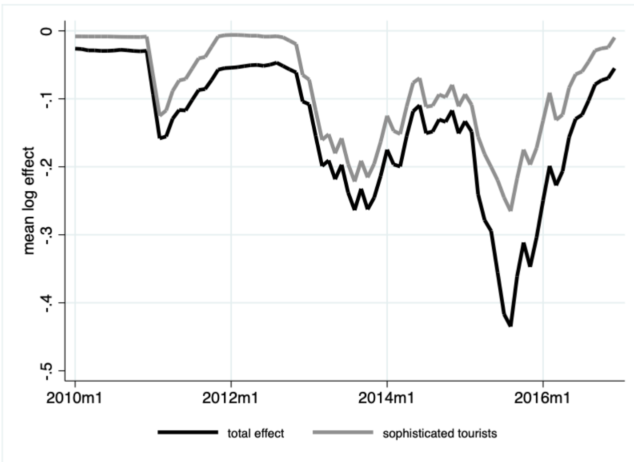

Panel B: Egypt

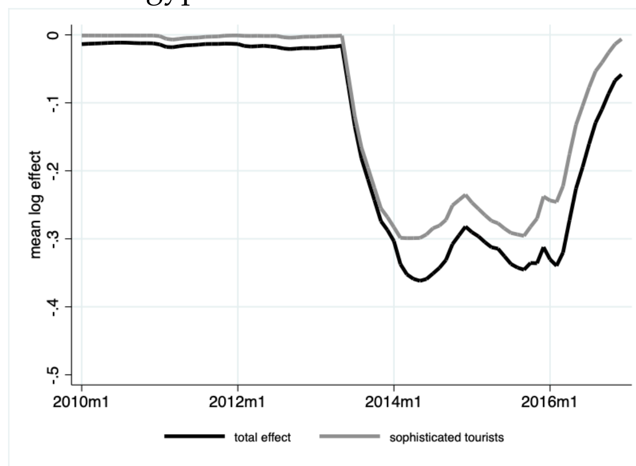

Comparison of model-based predictions versus reduced form for Tunisia

Panel C: Model based

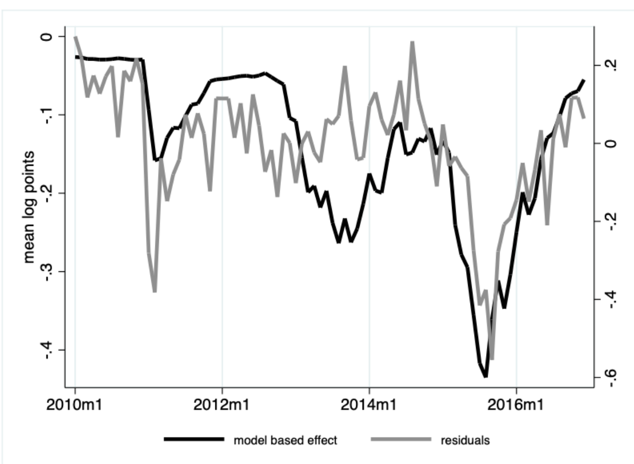

Panel D: Reduced form

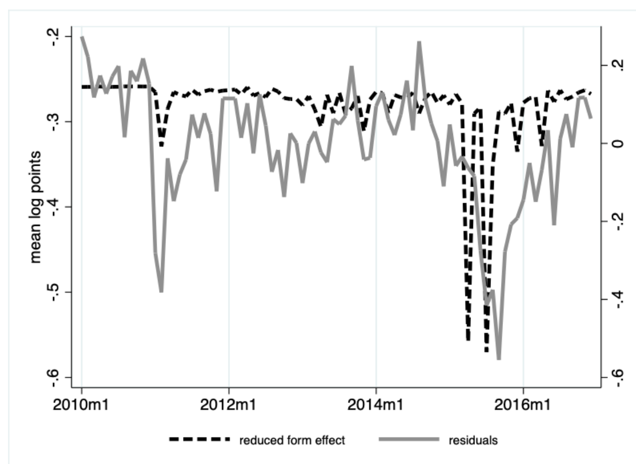

Notes: Figure plots the overall model-based impact of violent news shocks on aggregate card spending for Tunisia (left) and Egypt (right) over time. The solid black line presents the total effect. The grey line presents the effect that is driven solely by "Event-based beliefs". The difference in the two effects is capturing the role of the media multiplier in shaping the economic impact of violent events. Egypt entered a phase of significant and persistent domestic upheaval from mid 2013, while Tunisia saw more erratic violence (see also Figure A7). The figure illustrates the model fit vis-a-vis the fit that would be implied in reduced form exercises commonly studied using the case of Tunisia. The left figure plots the average model-based estimated effect of violent events and their media coverage on tourism spending across the 57 tourist origin countries for which news reporting data is available as a solid line. For comparison, the grey line represents the average of the residuals of the spending data across the 57 tourist origin countries after having removed dyad-, origin-by-time and destination by month fixed effects. The solid line is tracking the grey line quite tightly highlighting that the model fit is approximating closely the patterns in the spending data. The right figure plots these same residuals together with the simple implied average reduced-form effect of lagged violence on spending across the 57 tourist origin countries for which media reporting is available. The simple reduced form approach, as is studied for example in Table 1, does a much poorer job in capturing the variation in the residuals highlighting the value that the model-based approach can bring.

Table 1: Reduced Form Results: Impact of News Reporting and Tourism activity

	(1)	(2)	(3)	(4)	(5)	(6)
	log(Spending)			log(Number of cards)		
Panel A: News on tourist being targeted						
News on tourists targeted (share of all articles)	-0.552***	-0.529***	-0.205**	-0.628***	-0.617***	-0.192***
	(0.092)	(0.098)	(0.090)	(0.068)	(0.074)	(0.066)
Observations	23859	23859	23859	23869	23869	23869
R2	.966	.967	.972	.97	.971	.979
Panel B: News on any fatal violence						
News on violence with fatalities (share of all articles)	-0.458***	-0.337***	-0.196**	-0.364***	-0.269***	-0.061
	(0.092)	(0.098)	(0.093)	(0.060)	(0.062)	(0.054)
Observations	23859	23859	23859	23869	23869	23869
R2	.966	.967	.972	.97	.971	.979
Dyad FE	YES	YES	YES	YES	YES	YES
Origin/Time FE	YES	YES	YES	YES	YES	YES
Dest./Month FE	YES	YES	YES	YES	YES	YES
Event controls	NO	YES	NO	NO	YES	NO
Dest./Time FE	NO	NO	YES	NO	NO	YES

Notes: Table presents regression capturing reduced form effect of dyadic (tourist-origin by destination) specific news coverage on the dyadic log values of card spend in columns (1)-(3) and the number of cards in a month in columns (4)-(6). Panel A uses as news measure the share of articles in a month on a dyad that is classified as capturing tourists being targeted by violent events. In panel B, the news measure captures the share of news in a month on a dyad that is classified as covering violent events. Explanatory variables are lagged by one month to account for the lagged response of tourism to events and news. Event controls are the first four principal components combining the main country-level violence data series studied in Table A3. Robust standard errors clustered at the dyad level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 2: Calibrated Model of Tourism Beliefs and Tourism Activity

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	log(Spending)					log(Number of cards)				
probability of danger (based on violence data) Π_{dt}	-0.208***			-0.190***		-0.203***			-0.185***	
	(0.019)			(0.019)		(0.020)			(0.019)	
probability of danger (tourist news-based) π_{hdt}		-0.364***		-0.273***			-0.367***		-0.278***	
		(0.051)		(0.049)			(0.049)		(0.048)	
probability of danger (fatal news-based) π_{hdt}			-0.239***					-0.213***		
			(0.049)					(0.055)		
weighted probability of danger $\chi\Pi_{dt} + (1-\chi)\pi_{hdt}$					-0.750***					-0.795***
					(0.060)					(0.066)
Observations	23859	23859	23859	23859	23859	23869	23869	23869	23869	23869
R2	.966	.966	.966	.967	.967	.971	.971	.97	.971	.972
Dyad FE	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Origin/Time FE	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Dest./Month FE	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES

Notes: Table presents results from a regression explaining variation in tourist activity measured either as the log value of card spending in columns (1)-(5) and the log number of active cards in columns (6) - (10) on a dyad over time with the probability of a country being in the latent state of being "dangerous". Columns (1) and (6) explore the relationship between the dependent variables the probability of a country being "dangerous" as inferred by "Event-based tourists" Π_{dt}. Columns (2), (3) and (7), (8) explore the relationship between the dependent variables and the probability of a country being "dangerous" as inferred by "news-based tourists" π_{hdt} where the beliefs are either learned through the news reporting on violence targeted against tourists (columns 2 and 7) or through the general news reporting on any violent events with fatalities (columns 3 and 8). Column (5) and (10) explore the weighted average of the two. Please refer to section 5.1 for how we leverage the violence and news reporting data to estimate Π_{dt}, π_{hdt} and χ. Robust standard errors clustered at destination/month level are in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table 3: Role of Freedom of Press in Shaping Impact

	(1)	(2)	(3)	(4)	(5)	(6)
	log(Spending)			log(Number of cards)		
probability of danger (tourist news-based) π_{hdt}	-0.238***	-0.235**	-0.166**	-0.219***	-0.170**	-0.155**
	(0.080)	(0.105)	(0.075)	(0.077)	(0.081)	(0.074)
probability of danger (tourist news-based) π_{hdt} * free press	-0.232**	-0.235*	-0.197**	-0.272***	-0.321***	-0.223**
	(0.099)	(0.120)	(0.097)	(0.095)	(0.098)	(0.094)
probability of danger (based on violence data) Π_{dt}			-0.187***			-0.168***
			(0.027)			(0.028)
probability of danger (based on violence data) Π_{dt} * free press			-0.002			-0.031
			(0.035)			(0.037)
Observations	23859	21340	23859	23869	21349	23869
R2	.966	.967	.967	.971	.973	.971
Dyad FE	YES	YES	YES	YES	YES	YES
Origin/Time FE	YES	YES	YES	YES	YES	YES
Dest./Month FE	YES	YES	YES	YES	YES	YES

Notes: Table presents results from a regression explaining variation in tourist activity measured either as the log value of card spending in columns (1)-(3) and the log number of active cards in columns (4) - (6) on a dyad over time with the probability of a country being in the latent state of being "dangerous" from "event-based tourists" Π_{dt} and those of "news-based tourists" π_{hdt}. Please refer to section 5.1 for how we leverage the violence and news reporting data to estimate these. The table explores whether the impact of violence shocks is heterogenous across tourist origin countries in the extent to which these tourist origin countries have a "free press" measured as whether an origin countries as a press freedom score above average. The results highlight that the impact of violent news shocks (as opposed to violent events) affecting the beliefs are driven by the tourist origin countries that have a free press. Robust standard errors clustered at destination/month level are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Online Appendix and Supplementary Material

"How Big is the Media Multiplier? Evidence from Dyadic News Data"

Appendices A and C for Online Publication; Appendices B and D are Supplementary Material

October 29, 2021

This appendix is subdivided into four sections. Section A presents further robustness checks and additional results as figures or tables that were omitted from the main paper due to space constraints. Section B provides more details on the grid-search used. Section C provides more information on the calculations for the economic impact estimates. Lastly, Section D presents further description, results, and details about the machine-learning approach used to classify the 450,000 news articles.

A Further results and robustness checks

A.1 Violence and Card Activity

In this section we look at the relationship between the (log of) spending by origin country h in destination country d at date t, denoted by y_{hdt}, and actual violent events using the following specification:

$$y_{hdt} = \alpha_{hd} + \alpha_{ht} + \alpha_{dm(t)} + \xi v_{hd,t-1} + \varepsilon_{hdt} \tag{12}$$

where α_{hd} are dyad fixed effects, α_{ht} are origin country/time fixed effects and $\alpha_{dm(t)}$ are destination/month of year fixed effects.

Our core violence measure, denoted by $v_{hc,t-1}$, is lagged by one month to capture the possibility that international travel reacts to past violence. We expect to find that $\xi < 0$ in (12), i.e. violence deters travel activity. We use four sources of data on violent events at the country level in different versions of (12). In order to make the magnitudes from different data sources comparable, we divide the right-hand-side variable measuring violence by its respective standard deviation.

Table A3 reports regressions from the specification in (12) and shows compelling evidence of a negative link between violence and travel activity. Columns (1) through (4) show that there is a significant correlation between all five measures of violence and the level of tourism activity in a country. The size of the coefficients is in the range of 4% to 7.6% decrease in card spending for an increase of violence by one standard deviation.

In column (5) we try to get a more complete impression of the relationship between violence and spending using all available information from the different measures by combining nine different measures using a principal component analysis. We then represent $v_{hd,t-1}$ in (12) with a four dimensional vector comprising the first four principal components, with the results reported in column (5). In line with the results in columns (1) through (4), we find a robust negative relationship between principal components 1, 2 and 4 and spending. In terms of magnitude, spending falls by about 7% with an increase in the first component and by about 4% with the second component. Since this summarizes information from a range of sources, we will use this representation of violence in the analysis that follows.

Columns (6) - (10) show that we obtain similar results when using the log value of number of active cards as the dependent variable. This is important as it indicates that the main spending effect is coming from the *extensive* margin, i.e. the usage of cards from origin countries in the destination countries rather than the average amount spent per card.

Together these results are consistent with the idea that violence may deter potential travelers. Moreover, this is true even when we include, dyad, home country $\times$ time and month effects in the specifications so that the effect of violence is relative to mean dyad spending in a given month. The results are therefore not influenced by macro-trends in the origin country (country in which the cards are issued).

To allay the concern that results based on (12) could be explained by different time trends between times/places, that experience violence and those that do not, we conducted an event-study which studies patterns in aggregated spending and the number of active cards around known violent events. There is no evidence of any anticipatory contraction of spending or reduction in the number of active cards prior to an event taking place. On the contrary, we observe sharp contractions in card spending and the number of active cards with a one month delay only *after* a violent GTD event occurs.[1]

[1]See Appendix Figure A8.

A.2 Event Regression Evidence for News

Complementary to Figure A4 we run the following specification

$$p_{hdt} = \alpha_k + \alpha_{hd} + \alpha_t + \beta \times \text{Post}_{k,t} + \gamma_k \times (\text{Post}_{k,t} \times z_k) + \epsilon_{hdt} \tag{13}$$

where we have defined a dummy variable $Post_{k,t} = 1$ for $\tau = 0, 1, 2$ for up to two days following a violent event. Estimating equation (13) also allows us to explore whether this average effect is heterogenous across a range of event characteristics z_k: the level of casualties, whether American's are among the casualties, and attacks involving tourists.

We present results from specification (13) in Table A4. In columns (1) through (5), the dependent variable is our news coverage covering the share of articles on a day classified as indicating violence with fatalities. In columns (6) to (10), the dependent variable is the share of articles classified as indicating violence against tourists. In columns (1) through (5), we observe that reporting increases sharply in the two days after an event. The increase is larger when there are more casualties (column 1) and if there any American casualties (column 2). Suicide attacks are also more heavily covered (column 3) as well as attacks where tourists are targeted (column 4). Column (5) shows that these all hold up when included together. In columns (6) through (10), we repeat the analysis with the more refined measure that captures the share of articles on a day indicating that tourists were targeted. Here, the most notable observation is column (9), which highlights that, if an event is classified by the GTD as having tourists as targets, the reporting measure increases sharply.

A.3 Event Study Evidence for Spending

To identify the effect of violence, the difference in difference approach relies on there being a common underlying trend in spending between places that experienced violent events and those did not. One way of exploring whether this is plausible is to use an event study approach. This will also give us more insight into the timing of the spending response to violent events.

For this purpose, we define an "event" as a month when casualties in the GTD dataset surpass a given threshold. Across the five destination countries, there is a total of 256 country-by-month windows where an event with at least one casualty occurs (out of a total of the maximum possible 420 country-by-month windows from 2010-2016). For the empirical analysis, we focus on country-month event windows with at least 10 casualties, resulting in a total of 83 event months.

To look at the response in spending, we construct a twelve month window around each of these 83 event months which we denote by index k. We then use the following empirical specification to model the relationship between violent events and tourism activity:

$$y_{khdt} = \alpha_k + \alpha_{hd} + \alpha_{ht} + \sum_{\tau=-6}^{6} (\beta_\tau \times \text{Time to event month}_{k,t-\tau}) + \epsilon_{khd} \tag{14}$$

where, as above, y_{khdt} is the log of tourism spending in an event-month k from home country h in country d at date t. This specification includes event fixed effects α_k, dyad fixed effects α_{hd} and issuing country by time effects α_{ht}. As before, we adjust standard errors two way at the level of the dyad and event.

Estimating (14) permits us to trace out the patterns of aggregate spending around an event month. The results are depicted in Figure A8 for both log of spending and log of active card accounts. In both cases, there is no evidence of any anticipation of the event. Moreover, the observed pattern suggest a sharp contraction in card spending and the number of active cards with effects manifesting a month after an event occurred. Moreover, this occurs with a one month lag as in core specification. That said, it is clear that recovery from an event is quite slow.

In Appendix Table A7 we show that results are robust to dropping each country in turn, highlighting that the results are not an artefact of any of the five destination countries in our sample.

B Grid Search

In the grid search we proceeded as follows. We started from the estimation equation

$$y_{hdt} = \alpha_{hd} + \alpha_{ht} + \alpha_{dm} + \xi \sum_{\tau=0}^{9} \omega_\tau \left(\chi \Pi_{dt-\tau} + (1-\chi)\, \pi_{hdt-\tau}\right) + \varepsilon_{hct}. \tag{15}$$

and used different combinations of weights $\chi, \omega_\tau \in \{0, 0.05, 0.1, 0.15..., 1\}$ to calculate the term

$$\sum_{\tau=0}^{9} \omega_\tau \left(\chi \Pi_{dt-\tau} + (1-\chi)\, \pi_{hdt-\tau}\right).$$

which we then use as a regressor in equation 15. We pick the parameter values that yield the highest within R-squared. From this it should already be clear that assuming two different sets of weights on $\Pi_{dt-\tau}$ and $\pi_{hdt-\tau}$ would lead to an explosion of the complexity of the grid search. We therefore focus on one set of weights.

Note, we did not impose any restrictions on the weights ω_τ. This is remarkable because we get the highest explanatory power with weights that (weakly) fall over time. In particular, we get the weight sequence

$$0.2,\ 0.2,\ 0.15,\ 0.1,\ 0.1,\ 0.1,\ 0.05,\ 0.05,\ 0.05.$$

For χ we get a value of 0.4 which implies that only 40 percent of the agents in our model are estimated to be event-based. However, Event-based tourists will nonetheless drive most spending movements as the shifts in their beliefs are a lot more persistent.

We also ran robustness checks with a different model in which we used the level of violence, i.e. not only the state, as the variable that tourists are interested in. Results are very similar in that model so that we decided not to report it for brevity.

C Calculations of Total Loss

Assume that we have a monthly log spending before the violence that we call y_b. Assume that this takes some value $y_b = x$. The relationship between spending and violence is given as

$$y \approx x + \xi \sum_{\tau=0}^{9} \omega_\tau \left(\chi \Pi_{dt-\tau} + (1-\chi)\, \pi_{hdt-\tau}\right)$$

To compute the dollar value, we use the following transformation:

$$e^{y_b} - e^{y} = e^{x} - e^{x + \xi \sum_{\tau=0}^{9} \omega_\tau (\chi \Pi_{dt-\tau} + (1-\chi) \pi_{hdt-\tau})}.$$

We do this simply by giving every destination country the average treatment value coming out of all origin countries and applying it to the total tourism revenues measured at baseline in 2010.

In the 72 months after 2010 we find the following average losses per month: 0.042 (0.027) billion USD in Tunisia, 0.063 (0.041) billion USD in Israel, 0.163 (0.13) billion USD in Egypt, 0.255 (0.171) billion USD in Turkey. Numbers in brackets indicate the losses from event-based tourists alone. This means a total loss of 37.66 billion USD and 11.09 billion USD from news reporting.

To understand these numbers take the case of Tunisia, which had 3.48 billion USD in tourism receipts in 2010. This implies that $e^{y_b} = 3.48/12$ and the monthly loss is given by $3.48/12 - e^{((ln(3.48/12))+(-0.16))}$, where -0.16 is the average treatment on dyads into Tunisia in the period 2011-2016.

Figure A1: Map of countries included in our estimation samples across the exercises

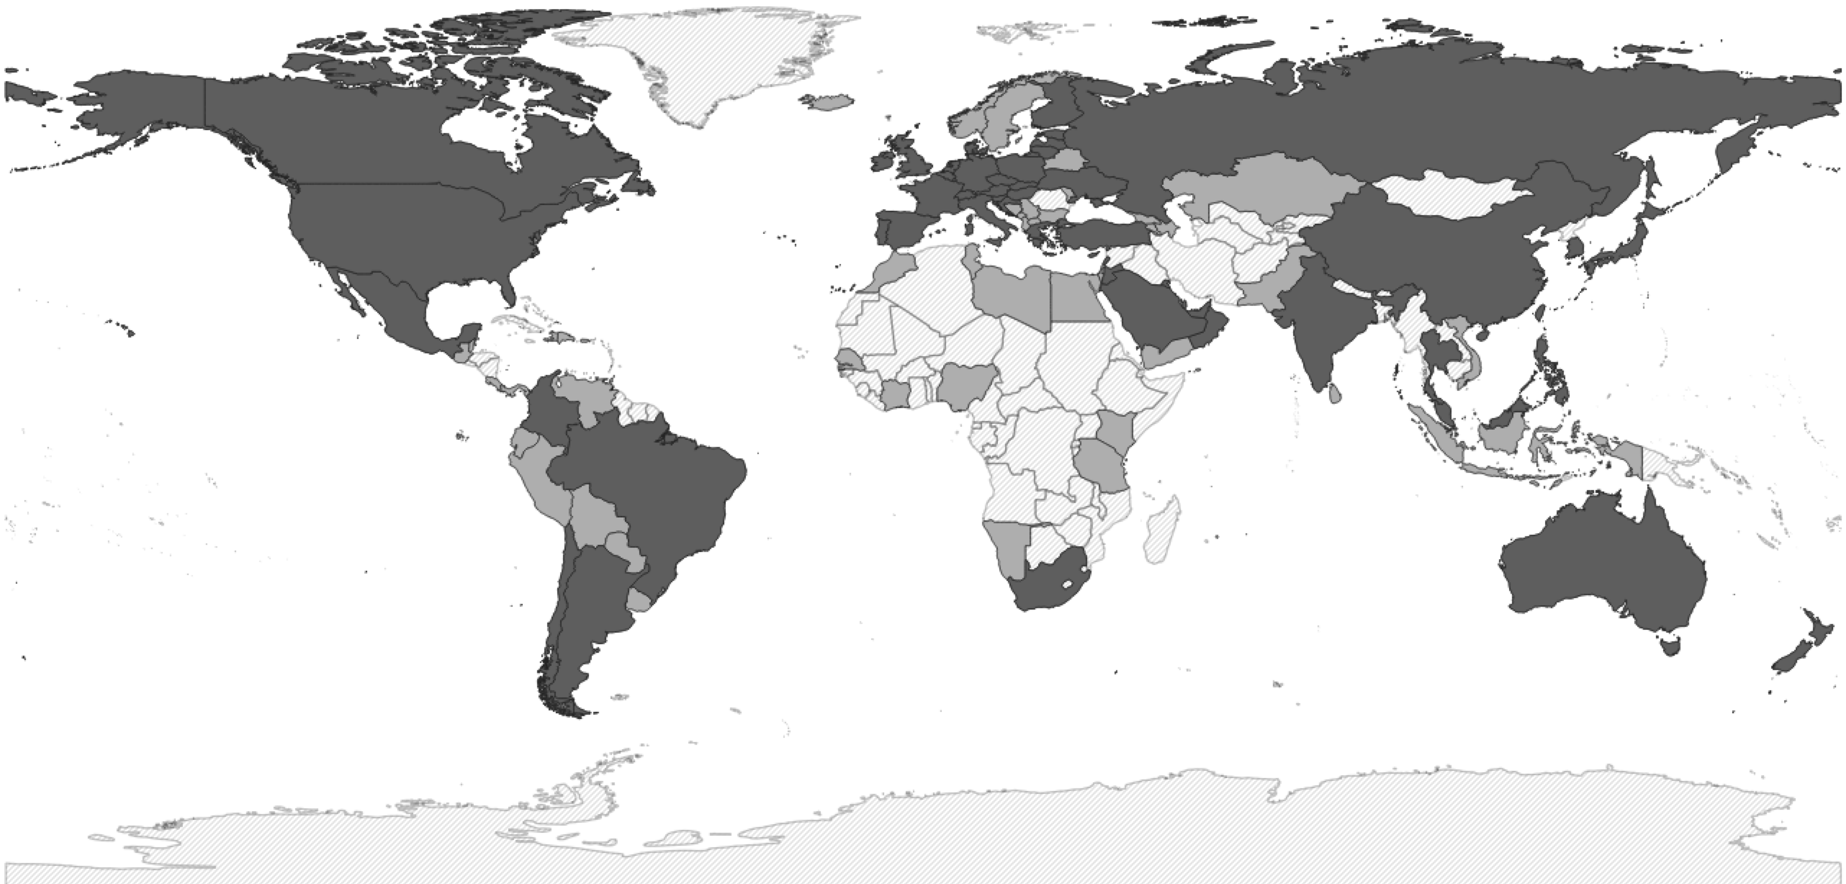

Notes: Figure indicates the origin countries included in our estimating sample. Dark-shaded are countries for which both newspaper and aggregated spending data is available, lightly shaded areas are countries for which only aggregated spending data is available.

Figure A2: Validation of aggregated spending data as a proxy for tourist arrivals: comparing subsets of data from the UN World Tourism Organisation

Panel A: Egypt

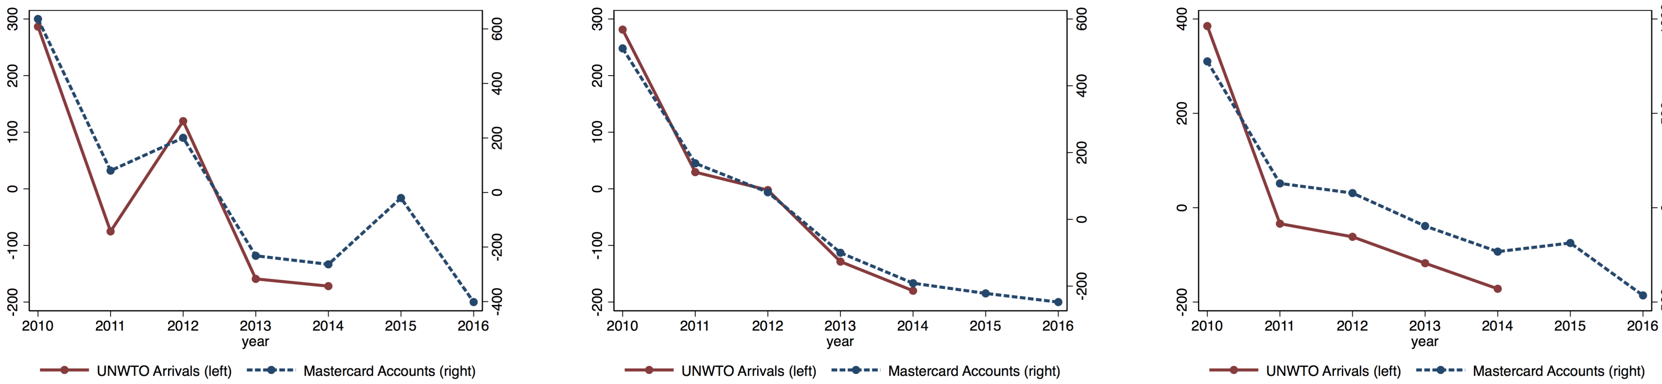

German French British

Panel B: Tunisia

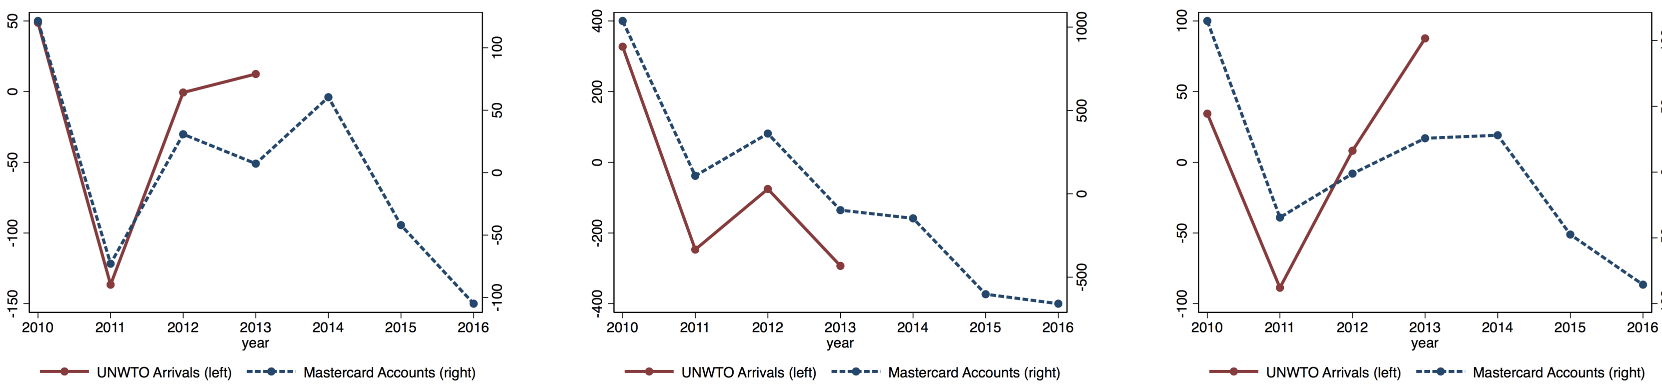

German French British

Notes: Figure plots dyadic data on tourist arrivals by destination country and by origin country, which is available annually for a small subset of countries from the UNWTO. The aggregate active accounts data has been further aggregated to the year level. The figures plotted are residuals obtained from removing dyad fixed effects as well as year fixed effects.

Figure A3: Distribution of Share of Articles on Fatal Violence and Violence Against Tourists

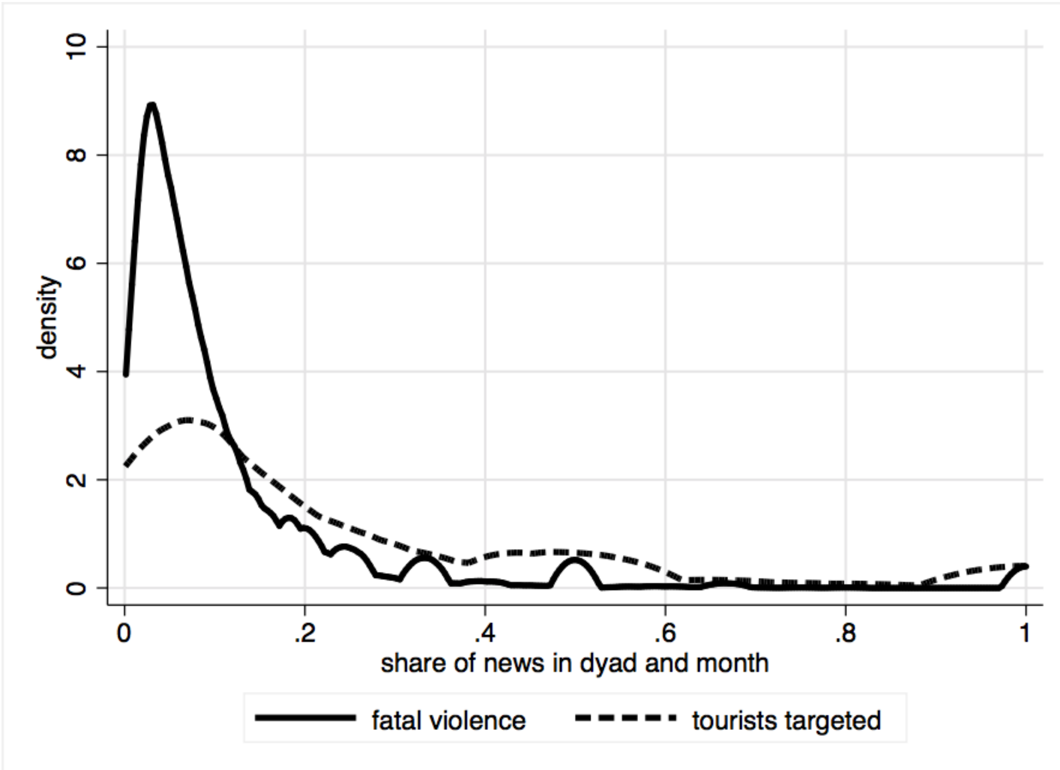

Notes: Figure plots kernel density plotting the distribution of share of newspaper reporting on (any) fatal violence or on violence directed towards tourists.

Figure A4: News reporting around known violent events in the GTD dataset: No evidence of diverging pre-trends prior to events

Panel A: Share of articles classified as indicating fatalities

Panel B: Share of articles indicating tourist targeted

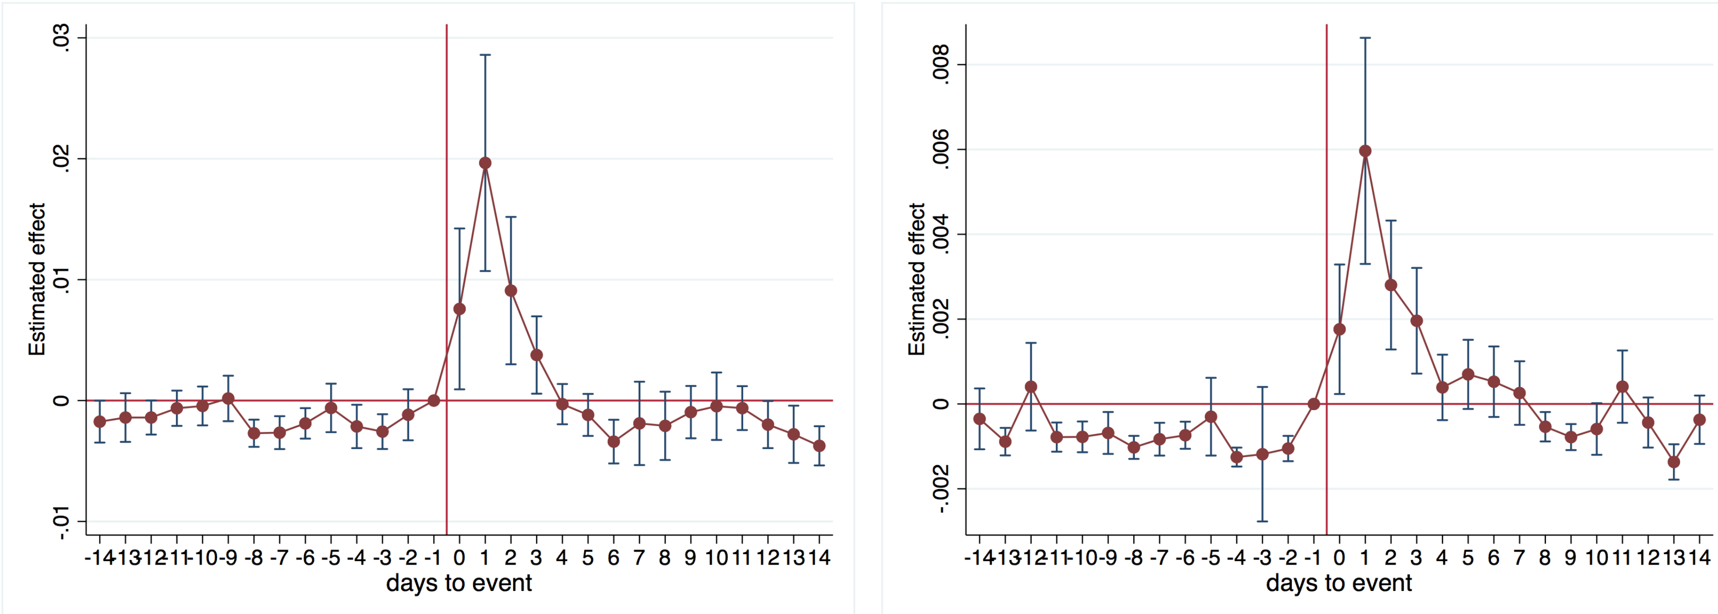

Notes: Figure plots point estimates from a regression that absorbs event, reporting dyad and day fixed effects. The dependent variable in Panel A measures the share of articles on a day and dyad that are classified as reporting on violence that involved fatalities. The dependent variable in Panel B measuring the share of articles on a day and dyad that are classified as tourism having been the target of violent events. The plotted point estimates capture the timing of reporting on a dyad relative to the timing of an individual event recorded in the GTD dataset. 90% confidence intervals obtained from clustering standard errors at the dyad level are indicated.

Figure A5: GTD Events and Reporting Activity: Noisy level effect on number of articles around violent events

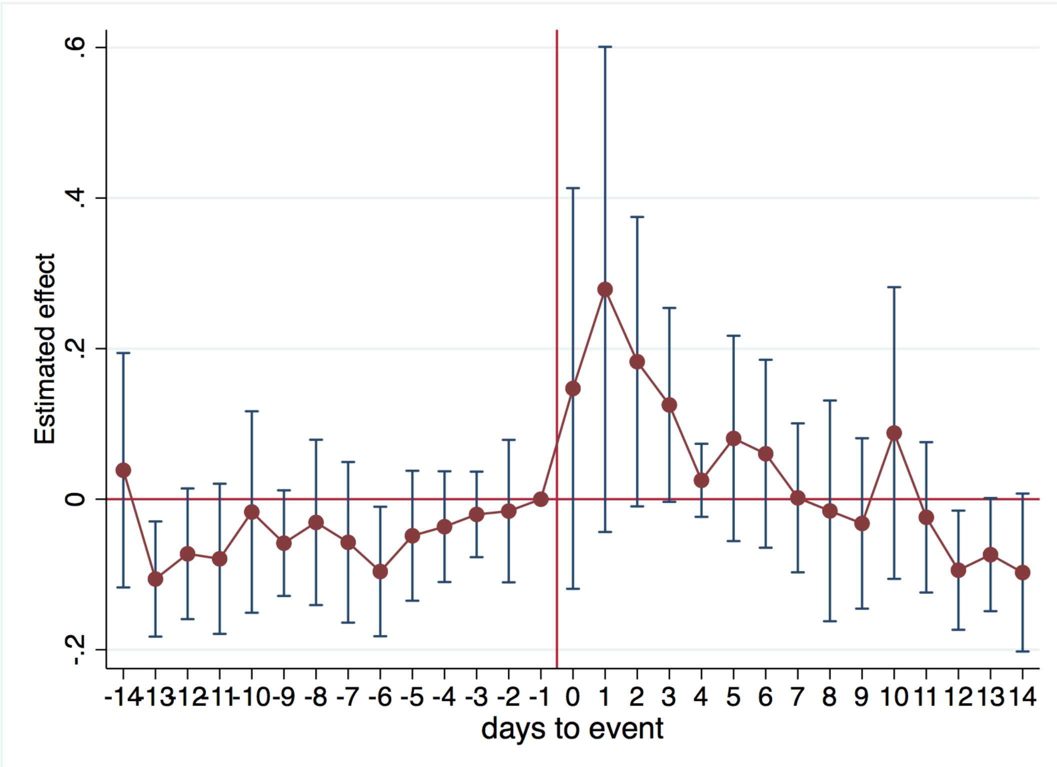

Notes: Figure plots point estimates from a regression that absorbs event, reporting dyad and day fixed effects. The plotted point estimates capture the timing of reporting on a dyad specific to the timing of an individual event recorded in the GTD dataset. 90% confidence intervals obtained from clustering standard errors at the dyad level are indicated.

Figure A6: Monthly share of articles classified as reporting on violent events across four destination countries averaged across the set of 57 tourist origin countries for which news reporting data is available

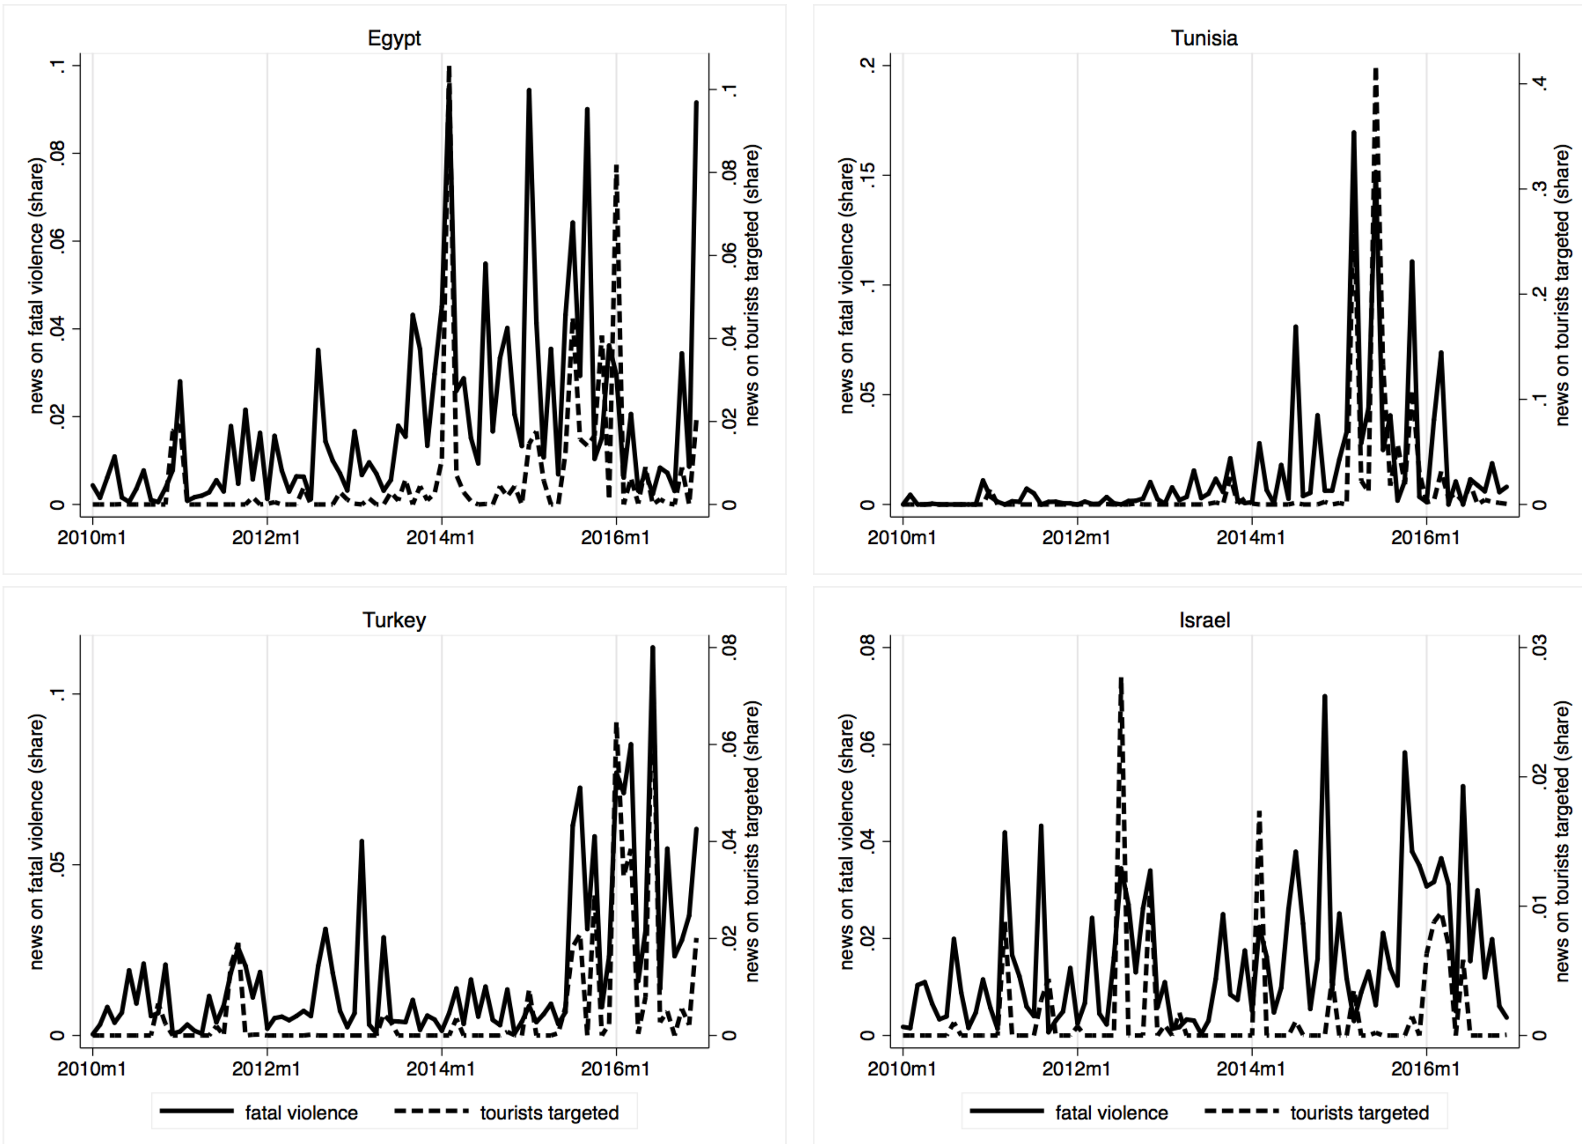

Notes: Figure plots the average reporting on violence across the 57 tourist origin countries for which data is available covering four main destination in our sample. The figure provide for each of destination the average share of news articles across the tourist-origin countries that are classified as reporting on violent events (left axis) or that are classified as covering violence against tourists (right axis). Note that the data was aggregated from daily- to monthly level as this coincides with the temporal resolution at which aggregated and anonymized credit card data has been made available.

Figure A7: Markov Chain Fitted Probability of Danger Across Sample Countries

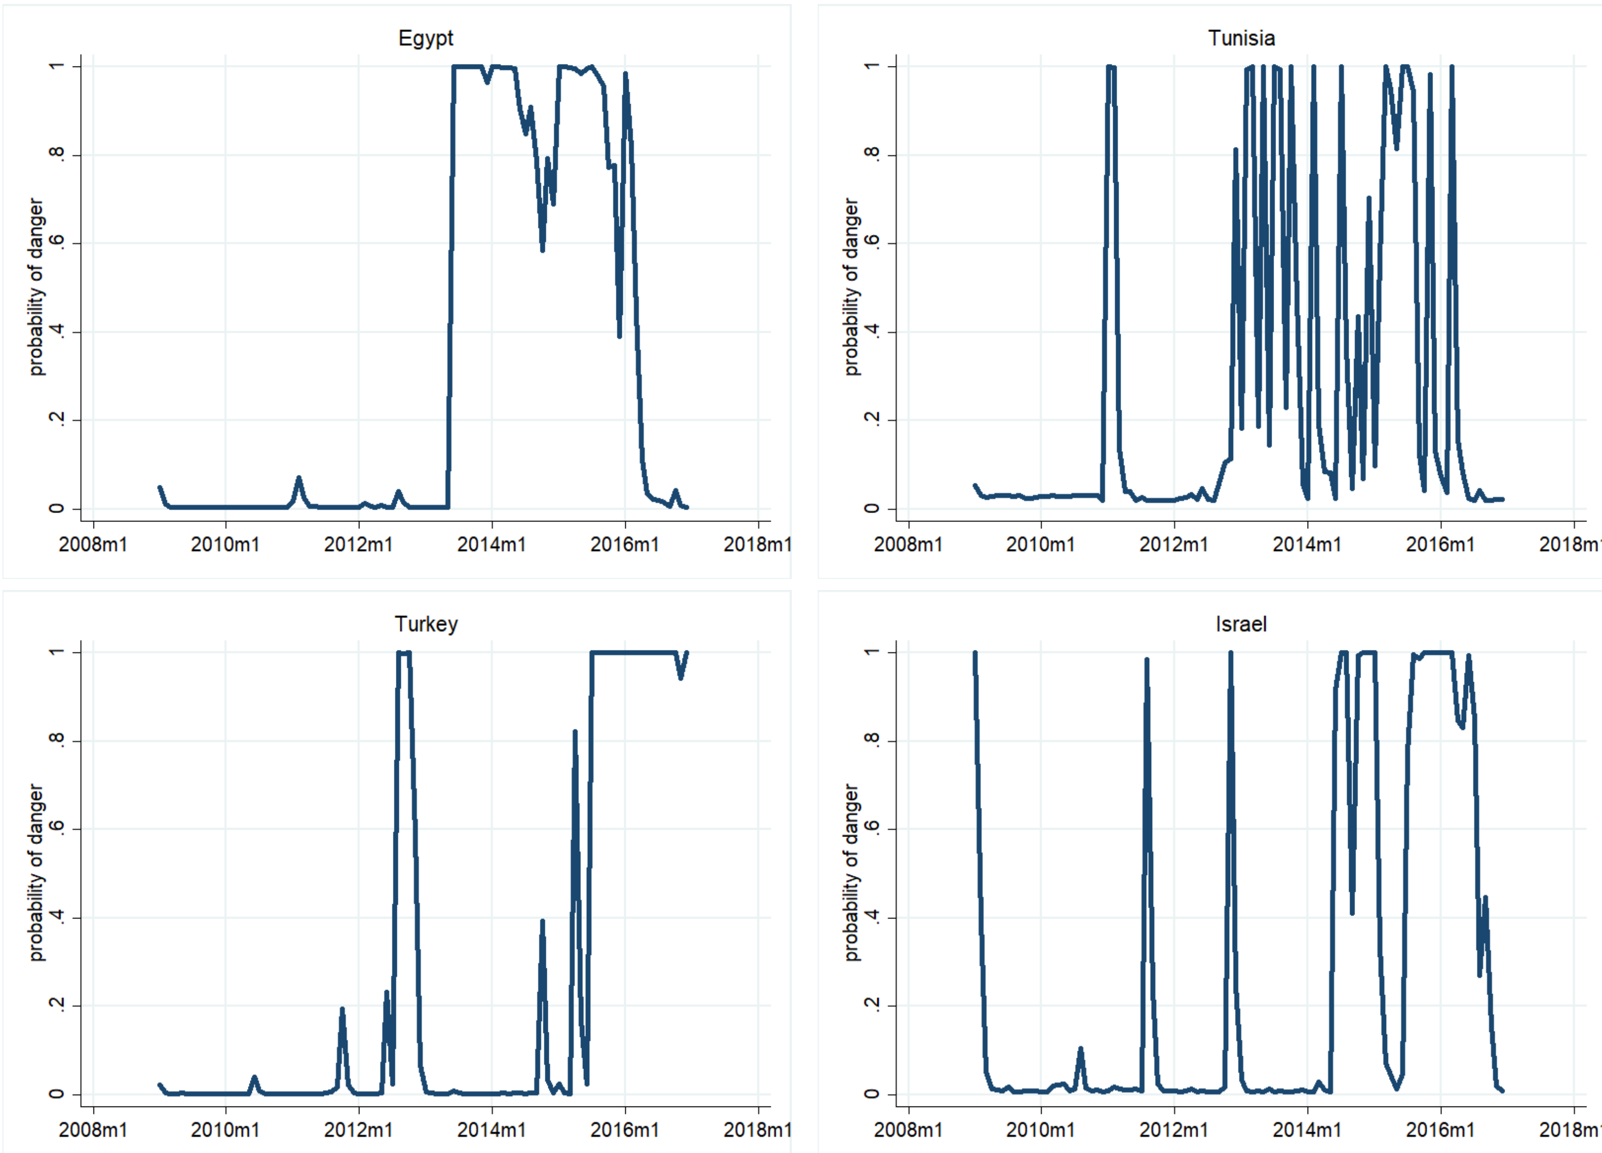

Notes: Figures plot out the probability of danger Π_{dt} as inferred by "Event-based" tourists from the time-series data on violent events. The figure plots the fitted Markov chain estimates across the four main tourist destinations in our data. The Markov switching model and model fitting exercise is described in more detail in the main text in Section 5. The estimated parameters of the Markov switching model are presented in Appendix Table A10.

Figure A8: Event study evidence of the average effect of violent events on tourist activity

Panel A: log(Spending) *Panel B*: log(Number of Cards)

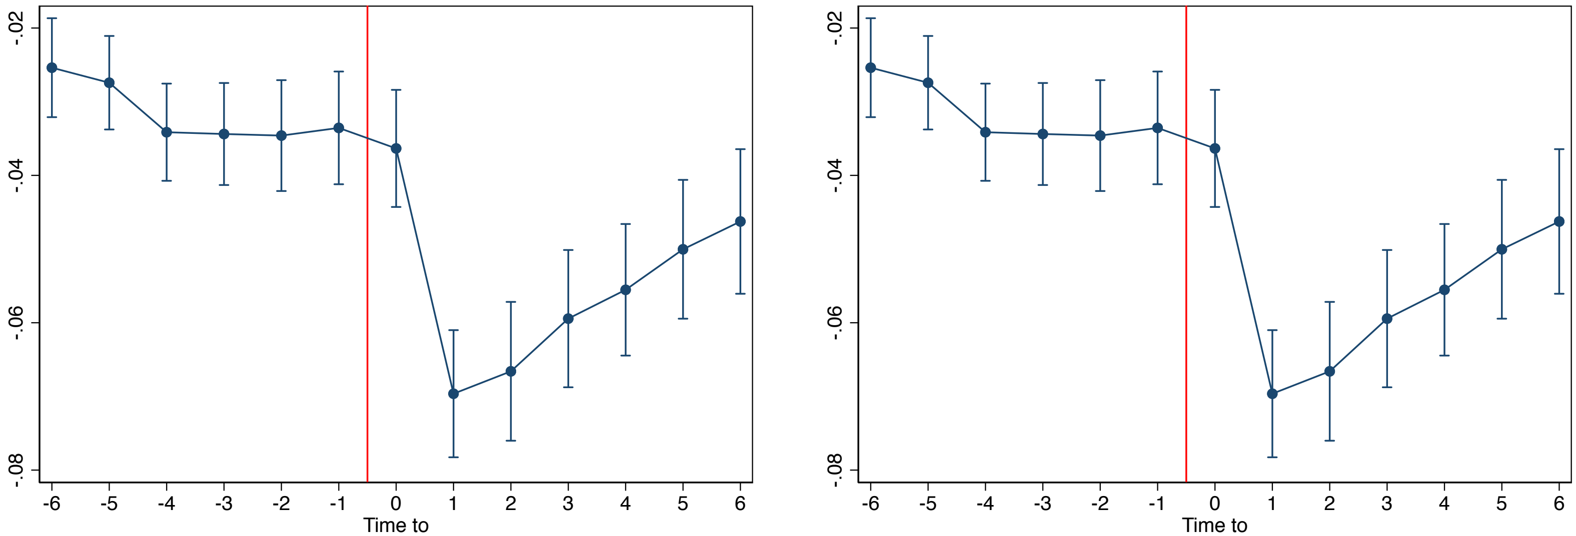

Notes: Figure plots results from an event study design exploring the effect of time series variation in the share of (any) fatal violence or on violence directed towards tourists across four main countries. The data set is an event-month level panel with each month with a violent event treated as an event-month. The regressions control for event fixed effects, dyad fixed effects and destination by month fixed effects. The figure plots the effect of an event occurring in month 0 on average card activity or the number of active cards across dyads. Standard errors are clustered at the dyad and time level with 90% confidence intervals indicated.

Figure A9: Heterogeneity of the Effect of News reporting on Aggregated Spending by Card Issuing Country

Panel A: News on Tourists Targeted (Share)

Panel B: News on Fatal Violence (Share)

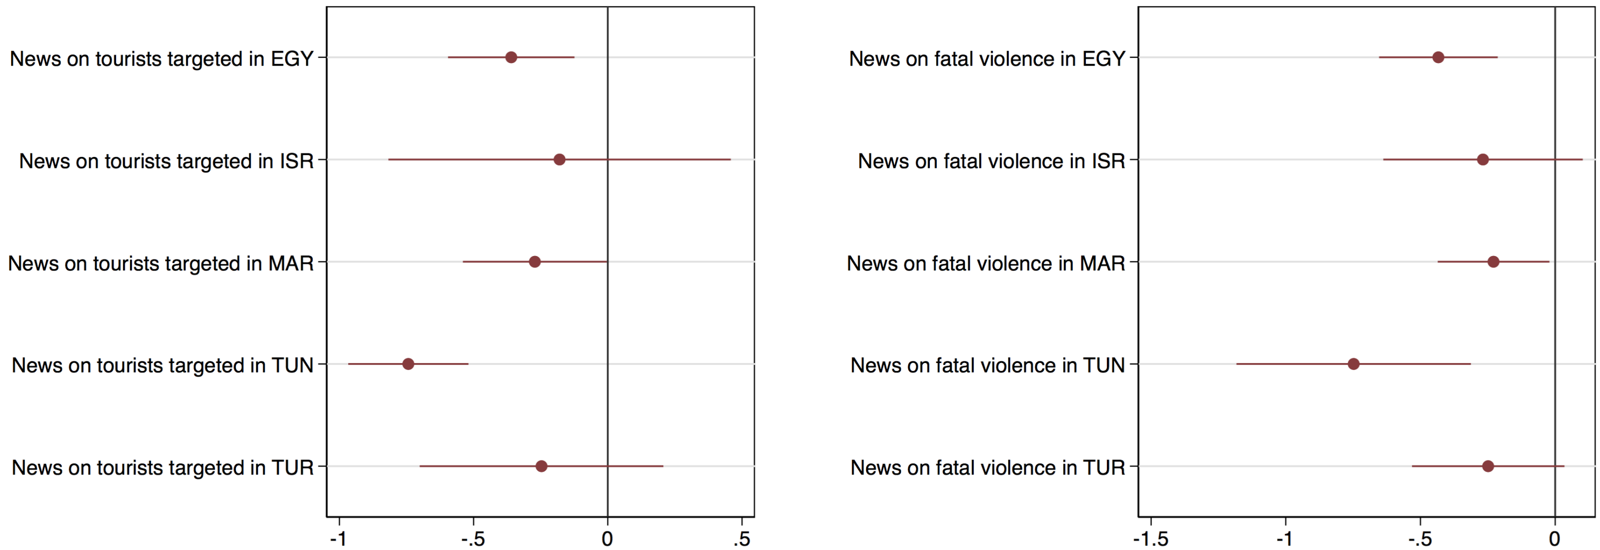

Notes: Figure plots point estimates from a regression that absorbs dyad, issuing-country by time fixed effects and destination by month fixed effects. 90% confidence intervals obtained from clustering standard errors at the dyad level are indicated.

Figure A10: Distribution of Beliefs about Safety or Danger

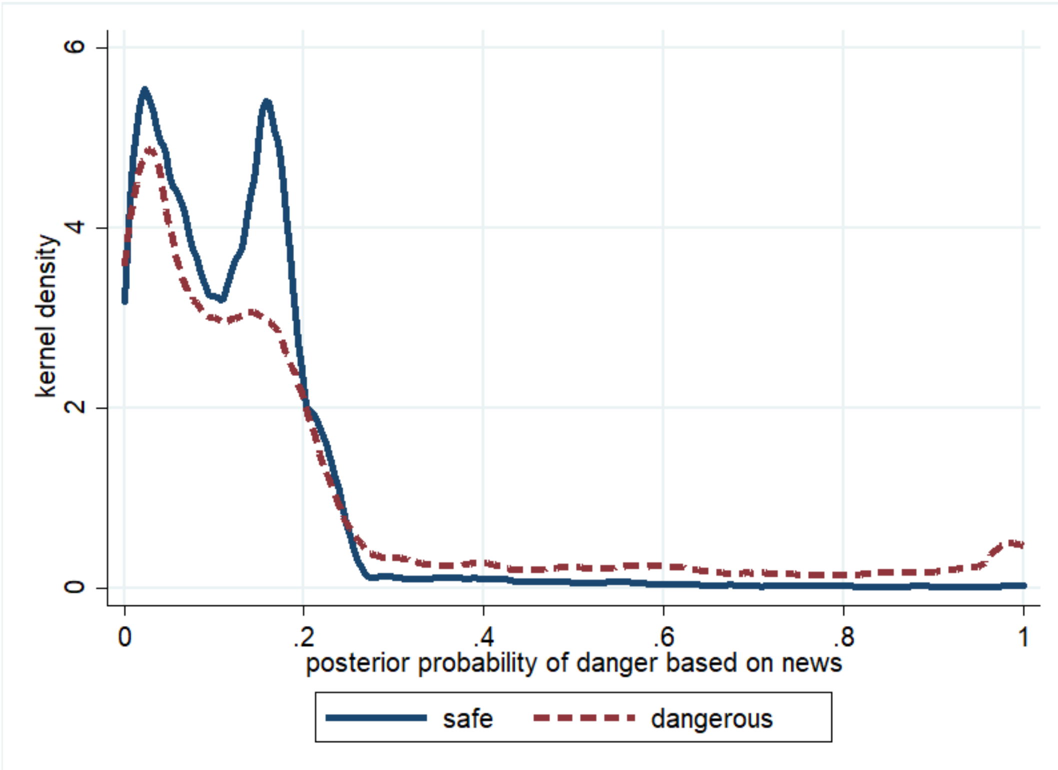

Notes: Figure plots kernel density plotting the distribution of share of newspaper reporting on (any) fatal violence or on violence directed towards tourists conditional on any reporting on fatal violence or tourists being targeted.

Figure A11: Evidence on the correlation between the level of news reporting and average beliefs about violence in months with violent events

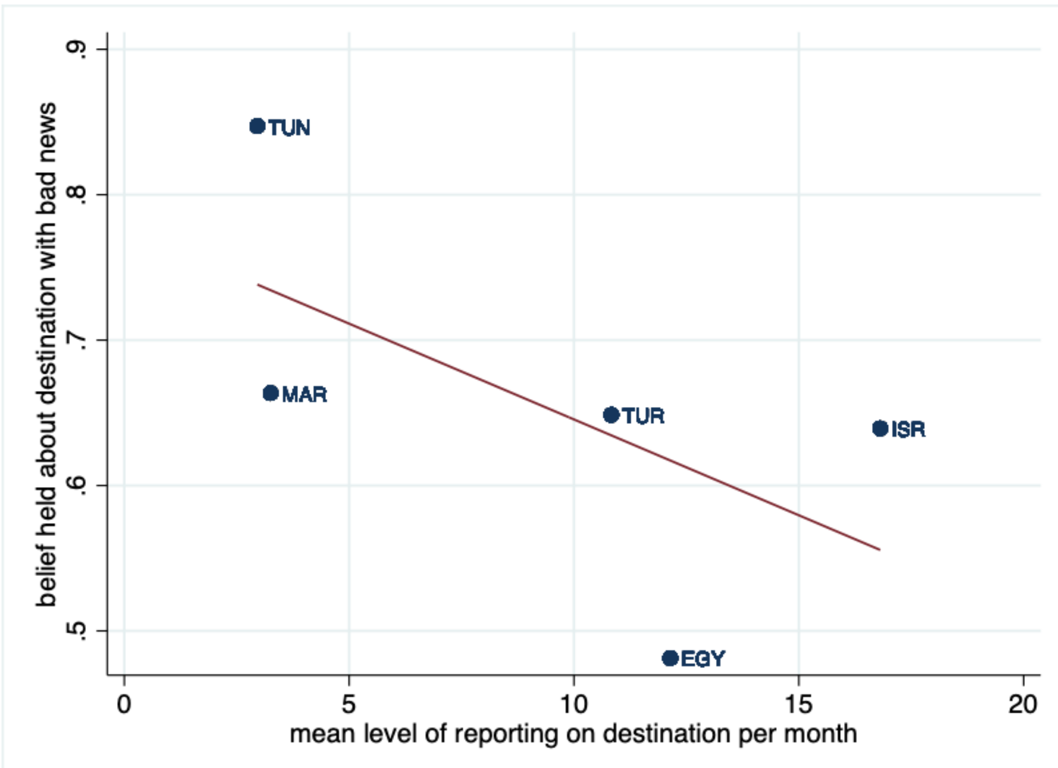

Notes: Figure highlights the interactive relationship between the mean-level reporting on the x-axis about a destination country and the average estimated belief about the state of the world being dangerous conditional on there being news on tourists being targeted on the y-axis. The figure shows that lower levels of mean reporting are associated with higher levels of beliefs about the underlying state of the world indicating danger.

Table A1: Summary Statistics

	Mean	SD	Observations
ACLED Events	0.612	0.934	31212
UCDP Events	0.278	0.968	55620
GTD Events	0.404	0.904	55620
ICEWS armed violence events	0.627	0.925	55105
GDELT armed violence events	0.718	0.983	49440
News on tourists targeted (count of articles)	0.035	0.856	30495
News on tourists targeted (share of all articles)	0.002	0.033	30495
News on violence with fatalities (share of all articles)	0.015	0.068	30495
Any tourist killed	0.001	0.031	61800
Same region x Any Casualties	0.006	0.075	61800
Common language x Any Casualties	0.007	0.082	61800

Table A2: Validation of aggregate spending data and official annual dyadic tourist arrival data available from UNWTO for a small subset of countries

	Cards (1)	Transactions (2)	Spend (3)
Panel A:			
arrivals	0.700***	1.189***	182.604***
	(0.197)	(0.262)	(60.961)
Dyads	294	294	294
Observations	1258	1258	1258
Panel B:			
arrivals	0.535***	1.048***	126.985***
	(0.097)	(0.190)	(39.719)
L.arrivals	-0.016	-0.021	0.271
	(0.067)	(0.108)	(20.866)
Dyads	290	290	290
Observations	974	974	974
Panel C:			
F.arrivals	0.344**	0.619**	83.950***
	(0.149)	(0.256)	(31.125)
arrivals	0.755***	1.183***	204.813***
	(0.206)	(0.299)	(62.734)
Dyads	286	286	286
Observations	961	961	961
Dyad FE	Yes	Yes	Yes
Time FE	Yes	Yes	Yes

Notes: The table reports regressions to validate that the anonymized and aggregated credit card data is a good proxy measure of tourism activity. The dependent variable measures the number of active cards in a destination by year in column (1), the number of transactions in column (2), and the total spending in column (3). The independent variables across panels A - C are the annual number of tourist arrivals in a destination obtained from annual data from the UN World Tourism Organisation. The data is not available for many dyads and is low frequency. Standard errors clustered at destination by time level in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A3: Effect of Country-level Violence measured by different event data sets on tourism spending and active cards

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	log(Spending)					log(Number of Cards)				
UCDP Events	-0.040***					-0.034***				
	(0.010)					(0.008)				
GTD Events		-0.076***					-0.076***			
		(0.017)					(0.016)			
ICEWS armed violence events			-0.068***					-0.054***		
			(0.020)					(0.019)		
GDELT armed violence events				-0.065***					-0.047***	
				(0.016)					(0.015)	
Armed violence component 1					-0.067***					-0.053***
					(0.017)					(0.017)
Armed violence component 2					-0.040***					-0.034**
					(0.014)					(0.015)
Armed violence component 3					0.017					0.010
					(0.015)					(0.015)
Armed violence component 4					-0.031**					-0.039***
					(0.013)					(0.014)
Observations	42254	42254	42254	42254	42254	42299	42299	42299	42299	42299
R2	.947	.947	.947	.947	.947	.969	.97	.969	.969	.97
Dest./Month FE	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Origin/Time FE	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Dyad FE	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES

Notes: Table presents regression capturing reduced form effect of destination-country specific violence on the dyadic (tourist-origin by destination) specific log values of card spending in columns (1)-(5) and the number of cards in a month in columns (6) - (10). The explanatory variables are lagged by one month to account for the lagged response of tourism activity to violent events. We also divide all explanatory variables by their standard deviation so that the coefficients can be interpreted as the response of a one standard deviation increase in violence regardless of the violence measure. The ”Armed violence” components are constructed by performing a principal component analysis of all violence data series that we have available. Robust standard errors clustered at at the destination by time level are provided in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A4: Event Characteristics and Reporting Intensity *across dyads*

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Share of articles indicating fatal violence					Share of articles indicating tourist targeted				
post	0.002***	0.002***	0.002***	0.011***	0.003*	-0.000	0.000***	0.000**	0.001**	-0.005***
	(0.000)	(0.000)	(0.000)	(0.002)	(0.002)	(0.000)	(0.000)	(0.000)	(0.000)	(0.001)
post × Casualties	0.000***				0.002***	0.000***				0.002***
	(0.000)				(0.000)	(0.000)				(0.000)
post × US Casualties		0.014***			0.016**		0.001**			0.001**
		(0.002)			(0.007)		(0.001)			(0.001)
post × Suicide attack			0.017***		0.020***			0.002***		-0.002
			(0.002)		(0.006)			(0.000)		(0.004)
post × Tourist targeted				0.014***	0.008***				0.022***	0.016***
				(0.003)	(0.003)				(0.003)	(0.002)
Observations	6033450	6122712	6122712	57855	57855	6033450	6122712	6122712	57855	57855
Number of Events										
Dyad FE	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Date FE	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Event FE	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES

Notes: The underlying data is an event-level balanced panel across dyads. Events are those recorded events with a specified date in the GTD dataset. For each event, we generate a balanced 14-day time window on either side of the event date and for each potential reporting country. The dependent variable in columns (1) - (5) measure the share of reporting on a day and dyad that is due to articles classified as indicating any fatal violence. The dependent variable in columns (6) - (10) measures the share of articles on day and dyad that are classified as indicating violence targeted at tourists. Robust standard errors clustered twoway at dyad and event level with stars indicating *** p<0.01, ** p<0.05, * p<0.1.

Table A5: Reduced form Results: Relative Nature of Variation in Reporting Measure

	(1)	(2)	(3)	(4)
	log(Spending)			
Panel A: News on tourist being targeted				
News on tourists targeted (count of articles)	-0.007**	-0.006**		
	(0.003)	(0.002)		
News on tourists targeted (share of all articles)	-0.512***	-0.494***		
	(0.095)	(0.101)		
News on tourists targeted (share of all articles) - first quartile			-0.005	0.051
			(0.051)	(0.048)
News on tourists targeted (share of all articles) - second quartile			-0.087*	-0.051
			(0.049)	(0.047)
News on tourists targeted (share of all articles) - third quartile			-0.160***	-0.123***
			(0.041)	(0.042)
News on tourists targeted (share of all articles) - fourth quartile			-0.304***	-0.297***
			(0.079)	(0.081)
Observations	23859	23859	23859	23859
R2	.966	.967	.966	.967
Panel B: News on any fatal violence				
News on violence with fatalities (count of articles)	-0.003	0.001		
	(0.002)	(0.003)		
News on violence with fatalities (share of all articles)	-0.422***	-0.351***		
	(0.097)	(0.102)		
News on violence with fatalities (share of all articles) - first quartile			-0.065	-0.019
			(0.047)	(0.044)
News on violence with fatalities (share of all articles) - second quartile			-0.063**	-0.018
			(0.025)	(0.025)
News on violence with fatalities (share of all articles) - third quartile			-0.085***	-0.035
			(0.024)	(0.025)
News on violence with fatalities (share of all articles) - fourth quartile			-0.113***	-0.069**
			(0.032)	(0.032)
Observations	23859	23859	23859	23859
R2	.966	.967	.966	.966
Dyad FE	YES	YES	YES	YES
Origin/Time FE	YES	YES	YES	YES
Dest./Month FE	YES	YES	YES	YES
Event controls	NO	YES	NO	YES

Notes: Table presents regression capturing reduced form effect of dyadic (tourist-origin by destination) specific news coverage on the dyadic log values of card spend. The table illustrates that a relative measure of news coverage is adequate in capturing the underlying relationship in the reduced form. Panel A uses as news measure the share of articles in a month on a dyad that is classified as capturing tourists being targeted by violent events. In Panel B the news measure captures the share of news in a month on a dyad that is classified as covering violent events. Explanatory variables are lagged by one month to account for the lagged response of tourism to events and news. Event controls are the first four principal components combining the main country-level violence data series studied in Table A8 and Appendix Table **??**. Robust standard errors clustered at the dyad level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table A6: Instrumental Variable Regression of Reduced Form Effect: Instrumenting news coverage intensity with known casualty-distribution for select events

	(1)	(2)	(3)
	News on tourists targeted		
Panel A: first stage with reporting spillovers			
Any tourist killed	0.181***	0.191***	
	(0.050)	(0.051)	
Contiguous country x Any Casualties		0.149***	0.146***
		(0.038)	(0.040)
Same region x Any Casualties		0.059**	0.061**
		(0.024)	(0.027)
Common language x Any Casualties		0.034**	0.033**
		(0.015)	(0.016)
R2	0.260	0.303	0.305
Panel B: Second stage: tourism spending			
News measure (share of articles)	-0.811	-0.790***	-0.628*
	(0.601)	(0.291)	(0.354)
R2	0.967	0.967	0.965
Weak IV	11.265	15.611	13.487
Panel C: Second stage: number of active cards			
News measure (share of articles)	-1.188*	-1.109***	-0.942***
	(0.663)	(0.297)	(0.335)
R2	.971	.971	.972
Weak IV	11.3	15.6	13.5
Observations	23859	23859	21527
Dyad FE	YES	YES	YES
Origin/Time FE	YES	YES	YES
Dest./Month FE	YES	YES	YES
Event controls	YES	YES	YES
Excluding direct treated dyads	NO	NO	YES

Notes: Tables presents first-stage (Panel A) and second stages (Panel B and C) of an instrumental variable regression that attributes variation in dyad-specific violent news coverage on tourism activity measured as the log value of tourism spending in Panel B and the log value of the number of active credit cards in Panel C. The instrument exploited in column (1) is an indicator taking the value 1 if a dyad has experienced a casualty due to a known terrorist event in a destination country in a specific month. In column (2) we augment this to code also tourist-origin months as treated if they are contiguous/have the same official language/ are located in the same geographic region. This implies that a German casualty in a tourist destination and month would also count as a casualty for Austria and Switzerland. In column (3) we drop all dyads that ever have a direct casualty themselves and identify the impact fully through spillovers, i.e. we identify the impact on tourism activity of a German casualty only on card spending of Austrian or Swiss cards through the casualties' impact on media reporting in these countries. Explanatory variables are lagged by one month to account for the lagged response of tourism to events and news. Event controls are the first four principal components combining the main country-level violence data series studied in Table A3. Robust standard errors clustered at destination/month level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table A7: Robustness of Effect of Violence and Aggregate Spending: Dropping each country in turn

	(1)	(2)	(3)	(4)	(5)	(6)
		Dropping a country in turn				
	All	EGY	TUN	TUR	MAR	ISR
Panel A:						
Armed violence component 1	-0.081***	-0.041***	-0.094***	-0.114***	-0.067***	-0.095***
	(0.008)	(0.010)	(0.008)	(0.010)	(0.008)	(0.013)
Observations	42268	33315	35610	32617	34241	32851
R2	.947	.954	.95	.945	.952	.953
Panel B:						
GTD Events	-0.076***	-0.030***	-0.089***	-0.110***	-0.063***	-0.085***
	(0.008)	(0.010)	(0.008)	(0.010)	(0.008)	(0.012)
Observations	42254	33315	35596	32600	34227	32834
R2	.947	.954	.95	.945	.952	.953
Panel C:						
ICEWS armed violence events	-0.068***	-0.025**	-0.081***	-0.118***	-0.056***	-0.070***
	(0.008)	(0.010)	(0.008)	(0.010)	(0.008)	(0.010)
Observations	42254	33315	35596	32600	34227	32834
R2	.947	.954	.95	.945	.952	.953
Panel D:						
UCDP Events	-0.040***	-0.034***	-0.043***	-0.042***	-0.039***	-0.045***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.016)
Observations	42254	33315	35596	32600	34227	32834
R2	.947	.954	.95	.944	.952	.953
Dest./Month FE	YES	YES	YES	YES	YES	YES
Origin/Time FE	YES	YES	YES	YES	YES	YES
Dyad FE	YES	YES	YES	YES	YES	YES

Notes: Robust standard errors clustered at destination by time level in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Explanatory variables are lagged by one month to account for the lagged response of tourism to events. We also divide all explanatory variables by their standard deviation so that the coefficients can be interpreted as the response of spending to one standard deviation increase in violence regardless of the violence measure. Components are coming from principal component analysis of all different violence data sub-categories.

Table A8: Relationship between Reporting and Aggregate Spending

	(1)	(2)	(3)	(4)
	log(Spending)			
News on tourists targeted (count of articles)	-0.006**	-0.002		
	(0.002)	(0.002)		
News on tourists targeted (share of all articles)	-0.494***	-0.199**		
	(0.101)	(0.091)		
Armed violence component 1	-0.084***		-0.084***	
	(0.009)		(0.009)	
Armed violence component 2	-0.026***		-0.028***	
	(0.006)		(0.006)	
Armed violence component 3	0.018***		0.013***	
	(0.004)		(0.004)	
Armed violence component 4	-0.015*		-0.016*	
	(0.009)		(0.009)	
Violent events with fatalities (count of articles)			0.001	0.003
			(0.003)	(0.003)
News on violence with fatalities (share of all articles)			-0.351***	-0.222**
			(0.102)	(0.096)
Observations	23859	23859	23859	23859
R2	.967	.972	.967	.972
Dyad FE	YES	YES	YES	YES
Origin/Time FE	YES	YES	YES	YES
Dest./Month FE	YES	NO	YES	NO
Event controls	YES	NO	YES	NO
Dest./Time FE	NO	YES	NO	YES

Notes: Explanatory variables are lagged by one month to account for the lagged response of tourism to events and news. Event controls are the first four principal components of the eight violence measures from Table 2, column (5). Robust standard errors clustered at destination/month level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table A9: Robustness to additional control variables and different violent news coding cutoffs

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Additional controls				Alternative classification cutoffs & not relying on hand coding					
					$c = 0.5$		$precision = 0.90$		$precision = 0.95$	
Panel A: News on tourist being targeted										
News on tourists targeted (share of all articles)	-0.529***	-0.559***	-0.205**	-0.153*	-0.286***	-0.102**	-0.852***	-0.327**	-0.963***	-0.303
	(0.098)	(0.097)	(0.090)	(0.080)	(0.045)	(0.044)	(0.168)	(0.143)	(0.257)	(0.203)
Observations	23859	23859	23859	23859	23859	23859	23859	23859	23859	23859
R2	.967	.967	.972	.976	.967	.972	.967	.972	.967	.972
Panel B: News on any fatal violence										
News on violence with fatalities (count of articles)	-0.337***	-0.380***	-0.196**	-0.199**	-0.167***	-0.085**	-0.337***	-0.196**	-0.244***	-0.073
	(0.098)	(0.098)	(0.093)	(0.080)	(0.040)	(0.035)	(0.098)	(0.093)	(0.082)	(0.078)
Observations	23859	23859	23859	23859	23859	23859	23859	23859	23859	23859
R2	.967	.967	.972	.976	.967	.972	.967	.972	.966	.972
Dyad FE	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Origin by Time FE	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Destination by Month FE	YES	YES	NO	NO	YES	NO	YES	NO	YES	NO
Destination by Time FE	NO	NO	YES	YES	NO	YES	NO	YES	NO	YES
Dyad Linear Trend	NO	NO	NO	YES	NO	NO	NO	NO	NO	NO
Exchange rate	NO	YES	NO	NO	NO	NO	NO	NO	NO	NO

Notes: Explanatory variables are lagged by one month to account for the lagged response of tourism to events. We employ two alternative classification cutoffs c as discussed in Appendix section D. In columns (5) - (10) we rely on the news measures constructed not using the secondary hand coding procedure we described in . The columns explore alternative classification cutoffs to highlight results are robust. All explanatory variables by their standard deviation so that the coefficients can be interpreted as the response of aggregate spending to one standard deviation increase in violence regardless of the violence measure. Robust standard errors clustered at destination by time level in parentheses. ** p<0.01, * p<0.05, * p<0.1.

Table A10: Markov Chain Estimates of Parameters

	hrid EGY	TUN	TUR	ISR	MAR
mean violence in danger	0.46	0.300	0.50	0.431	0.27
mean violence in safety	0.31	0.270	0.30	0.309	0.27
difference(danger-safety)	0.14	0.030	0.20	0.122	0.00
persistence of danger	0.96	0.601	0.93	0.814	0.27
persistence of safety	0.99	0.856	0.97	0.929	0.84

Notes: Table reports estimates of the parameters for the Markov chain switching model. For definitions see the main text in section 5.

Table A11: Effect of Markov Chain Fitted Probability of Dangerous State on Spending across destination countries

probability of danger (news-based) in Egypt	-0.951***
	(0.092)
probability of danger (news-based) in Israel	-0.508***
	(0.166)
probability of danger (news-based) in Morocco	-0.214
	(0.207)
probability of danger (news-based) in Tunisia	-1.065***
	(0.185)
probability of danger (news-based) in Turkey	-0.524***
	(0.095)
Observations	23859
R2	.967

Notes: Table reports estimates of the parameters for the Markov chain switching model. For definitions see the main text in section 5.

Table A12: Heterogenous effects: Role of geographic distance or other similarity

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	log(Spending)				log(Number of cards)			
probability of danger (tourist news-based)	-0.286**	-0.415***	-0.354***	-0.311***	-0.327***	-0.398***	-0.398***	-0.361***
	(0.115)	(0.048)	(0.070)	(0.079)	(0.113)	(0.051)	(0.063)	(0.075)
probability of danger (tourist news-based) * christian share	-0.001				-0.001			
	(0.002)				(0.002)			
probability of danger (tourist news-based) * muslim share		0.005				0.003		
		(0.003)				(0.003)		
probability of danger (tourist news-based) * far away			-0.023				0.068	
			(0.102)				(0.102)	
probability of danger (tourist news-based) * distance				-0.199				-0.024
				(0.193)				(0.197)
Observations	23859	23859	23859	23859	23869	23869	23869	23869
R2	.966	.966	.966	.966	.971	.971	.971	.971
Dyad FE	YES	YES	YES	YES	YES	YES	YES	YES
Origin/Time FE	YES	YES	YES	YES	YES	YES	YES	YES
Dest./Month FE	YES	YES	YES	YES	YES	YES	YES	YES

Notes: Table presents results from a regression explaining variation in tourist activity measured either as the log value of card spending in columns (1)-(3) and the log number of active cards in columns (4) - (6) on a dyad over time with the probability of a country being in the latent state of being "dangerous" from "event-based tourists" Π_{dt} and those of "news-based tourists" π_{hdt}. Please refer to section 5.1 for how we leverage the violence and news reporting data to estimate these. The table explores whether the impact of violence shocks is heterogenous across destinations in the extent to which the tourist origin countries are a) more distant or if b) the share of muslims in an origin country are higher or lower. Robust standard errors clustered at destination/month level are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

D Description of Machine Learning Method and Validation

We developed the algorithm for spotting fatal violence and attacks on tourists to identify articles out of sample. Training was always conducted on a balanced set, i.e. 1:1 negatives and positives, but we knew that the final dataset would be very imbalanced. This was a particularly important concern for attacks against tourists which is why we checked the coding results by hand for these news.

In all applications, we exclude tokens that appear in less than 100 articles. We sometimes improved the fit by choosing a higher cutoff of around 150 articles but we also wanted to use an additional method of dimension reduction, the singular value decomposition (SVD), so that we chose 100 as a default. We looked at unigrams, up to bigrams, and up to trigrams and experimented extensively with them. Generally, up to trigrams performed clearly the best and we therefore stick to them.

We use three ways to classify the two news items. First, we use a random forest of depth 12 for fatal violence and depth 9 for attacks against tourists. Second, we use a random forest of the same depth but only after running the SVD. In addition, we use a naïve bayes classifier. All steps and hyperparameters were checked using cross validation. A worry we had was that headlines would repeat similar key words so that cross validation used a relatively small training sample (three folds). Figure B1 illustrates the grid-search for the optimal tree depth for attacks against tourists without the SVD, for example. We kept increasing the maximum tree depth and recalculated the AUC on the testing sample of three folds. The figure illustrates quite nicely how the AUC first rises significantly but then stagnates and falls with rising depth. In order to maintain out-of-sample performance, we picked a relatively general tree depth of 9. The higher tree depth of 12 for fatal violence reflects the fact that we have many more circumstances of fatal violence and an algorithm that is able to pick it up a lot better.

Figure B2 shows the performance of the three models, two random forest models and one NB model, on three different folds and on two different samples with sampling rate of 1:1 and 1:10. The y-axis shows the AUC and the x-axis the average precision the respective classifier reached in the sample. The green dots show the result of the ensemble classifier, the simple average of the other three scores. Three things are clear from this. First, the NB performs a lot worse than the random forest - both in terms of AUC and precision. Second, the ensemble is performing better than the random forest, despite the fact that it uses the NB. Thirdly, the precision of the ensemble is less affected by the imbalance and therefore a lot more stable across the different samples. This is an important reason for us to adopt the ensemble method.

Figure B3 shows the precision recall curve for fatal violence on three random folds and figure B4 shows the same curve for violence against tourists. These figures are particularly important in a context with imbalanced data in which we are worried about precision. "Recall" on the x-axis is the true positive rate, i.e. the share of all actual articles with violence which the algorithm picks up. "Precision" on the y-axis is the rate at which articles which were identified as articles with violence were actually articles with violence. Clearly, precision is a lot lower when trying to identify violence against tourists. While the average precision is close to 0.85 for fatal violence, it is between 0.59 and 0.71 for violence against tourists. This implies that our precision cut-offs of 90 percent will exclude a lot more articles when we try to identify violence against tourists.

This is why we added an additional layer of hand-coding for violence against tourists. The analysis of mistakes made by the algorithm reveals something interesting about the task of spotting violence against tourists in newspaper articles. Some of these mistakes were difficult judgement calls such as news on shark attacks (on tourists) or an attack on a military bus in Egypt in which no tourists died. Some other manual recodings were driven by the text we downloaded, for example, declarations by our destination countries about events in other countries or when citizens of our destination countries conducted attacks elsewhere. These were downloaded as events in the destination countries and identified by the algorithm as being about tourists being attacked. Most remaining mistakes were driven by reactions to attacks, such as governments investigating the attackers, tourists fleeing the attack or reports about the court cases. We kept many of these codings if they were in the direct aftermath of an attack but excluded them if they were news reports on actions that were not taken in the direct context of an attack.

We also used the hand-coding to get an impression of the error rate that we imposed through our cutoffs. Between the cutoffs of $c = 0.8$ and $c = 0.75$ we hand-code 626 negatives and 192 positives, i.e. the proportions change considerably compared to the sample above $c = 0.8$. Also, we find a rapidly declining rate of false negatives in the remaining recodings. In 4,000 additional observations below $c = 0.75$, we found only an additional 416 positives. This is a rate decline of positives per article of 0.904 to 0.235 and 0.104 so that we suspect that the remaining articles will not contain a lot of actual positives. The resulting distribution of coded attacks is displayed in Figure B5, which shows two kernel densities. The first kernel density displays the overall distribution of predictions for violence against tourists coming out of the ensemble. Clearly, the predictions indicate that attacks against tourists is a rare event with most mass at low predictions. The red curve then shows the distribution of probabilities for the articles we identified through the hand coding. This provides a good confirmation of the decreased rate at which positives could be in the sample.[2]

The final proof that the algorithm works comes from a feature of the downloading process. We hand-coded only on articles downloaded from Lexis Nexis but, as we realized the method would work, later downloaded twice as much articles from different countries from Faktiva. In the sample that we confirmed by hand, the algorithm had spotted 535 attacks against tourists in the Lexis Nexis sample and 1,052 attacks against tourists in the Faktiva sample. This is a true out-of sample test, which implies that we have managed to develop a automated detection of attacks on tourists from the news.

[2]The few positives for low scores were identified from sources with very few articles as our RA sampled the 100 top articles from all sources.

Figure B1: Bias-variance trade-off in search for random forest tree depth

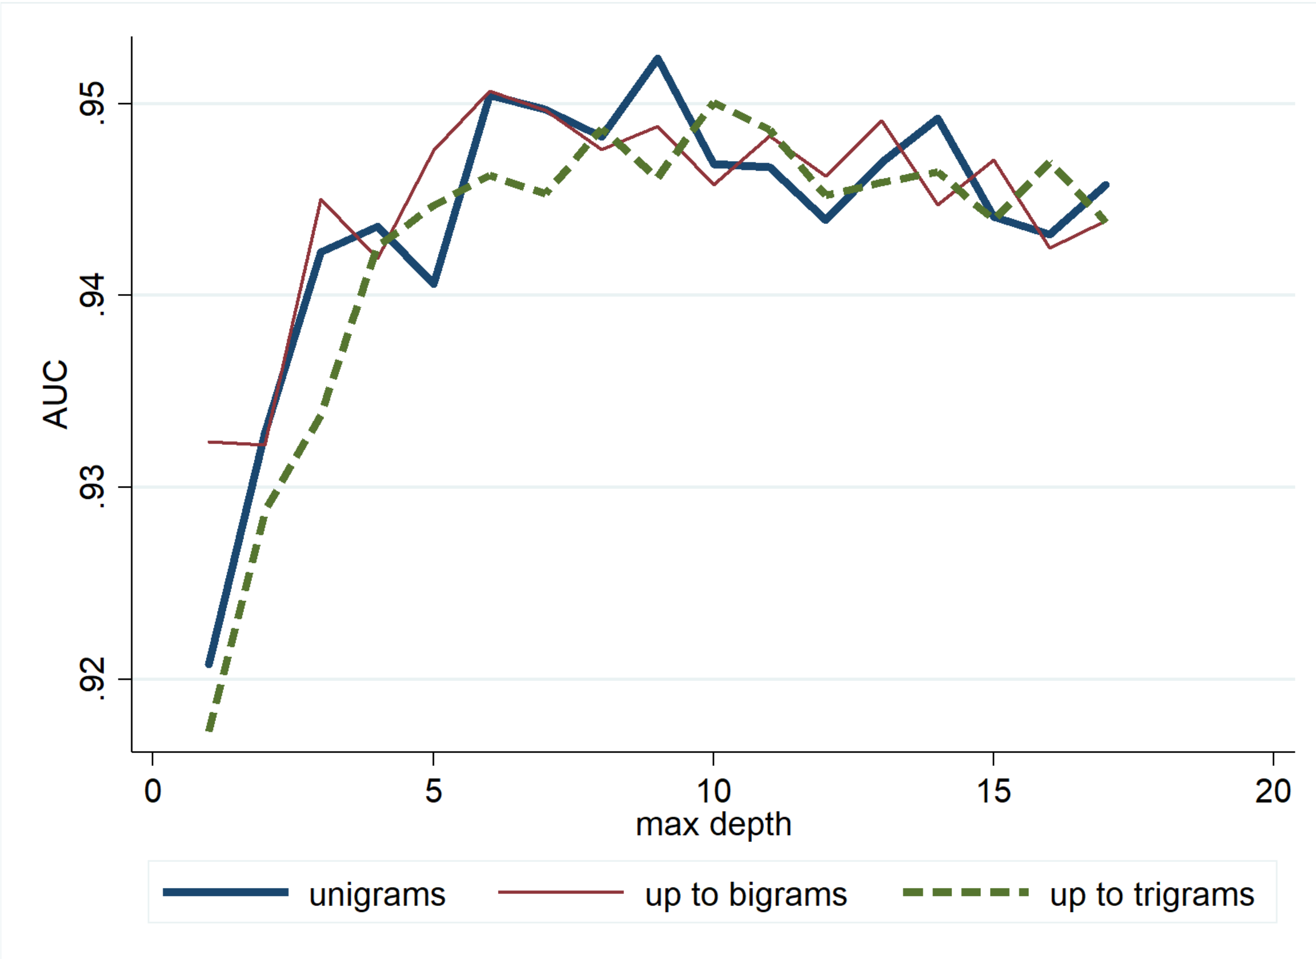

Notes: Figure plots out the AUC curve obtained on the testing sample. For each different maximum tree depth, the AUC is recalculated the three folds. The figure highlights the bias-variance trade-off as the AUC first rises as models are allowed to become more flexible but then stagnates and starts falling as models become too flexible and start overfitting the training data, resulting in worse accuracy in the testing samples.

Figure B2: AUC and Precision across a set of different models

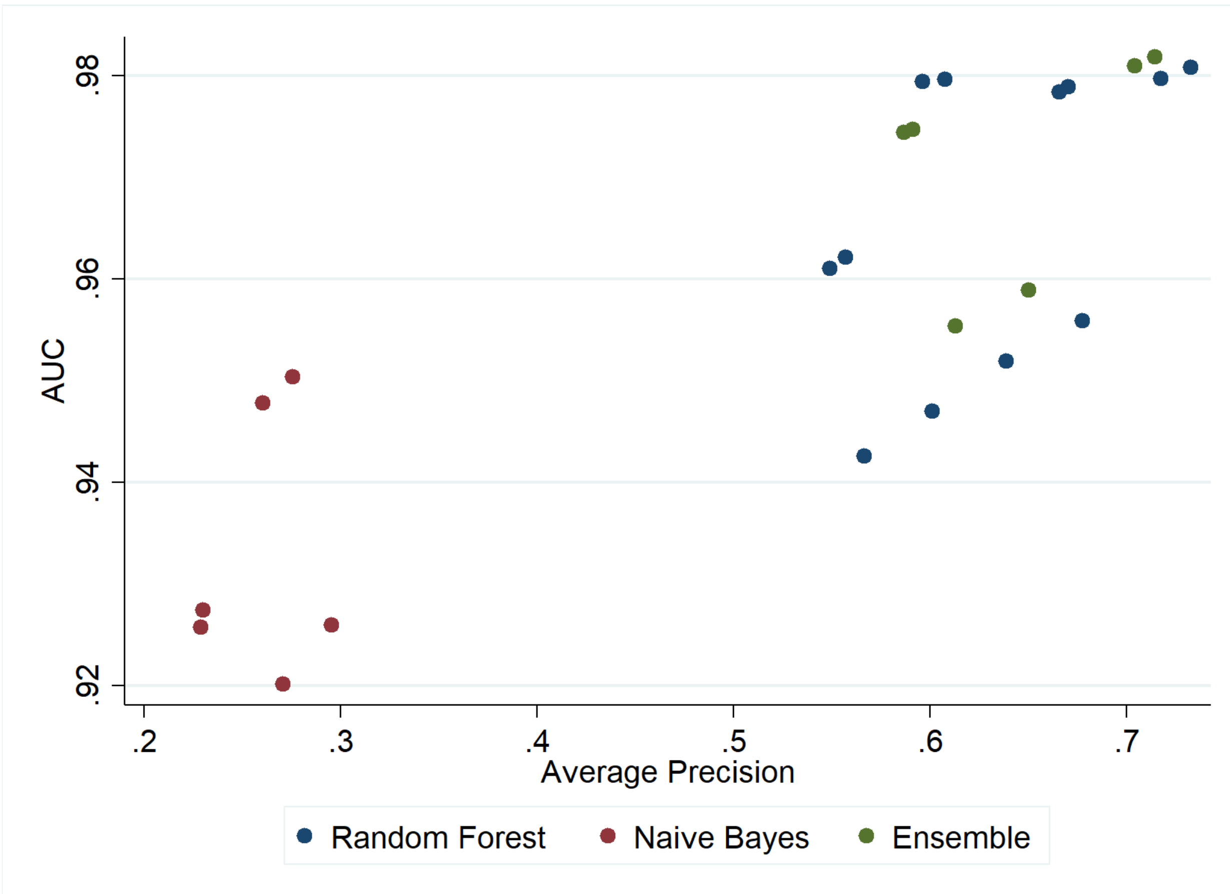

Notes: Figure plots the AUC scores obtained on the testing sample. The AUC is the area under the curve of the ROC curve. For each different maximum tree depth, the AUC is recalculated the three folds. The figure highlights the bias-variance trade-off as the AUC first rises as models are allowed to become more flexible but then stagnates and starts falling as models become too flexible and start overfitting the training data, resulting in worse accuracy in the testing samples.

Figure B3: Classification of articles capturing fatal violence: precision and recall across folds

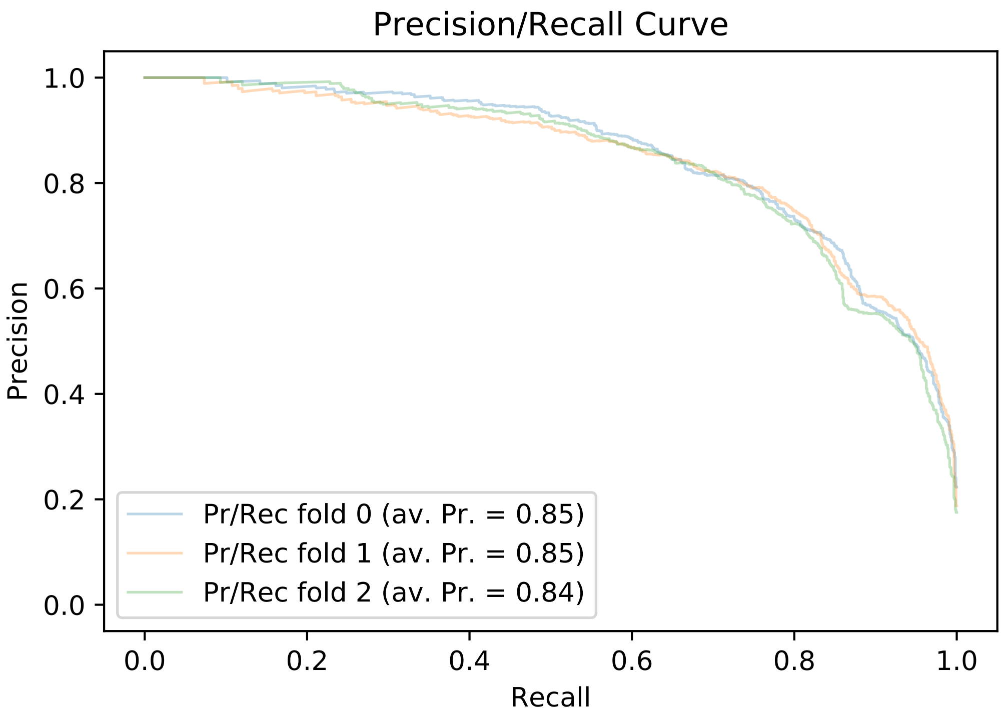

Notes: Figure plots the AUC and precision measured obtained on different folds of data highlighting that naïve Bayes performs worse compared to Random Forests, but that the ensemble of models outperforms the two. AUC is the area under the curve of the ROC curve. Precision of the algorithm captures the share of all correctly classified articles indicating violent events with fatalities among all articles that the algorithm classifies as indicating violence with fatalities.

Figure B4: Classification of articles capturing violence against tourists: precision and recall across folds

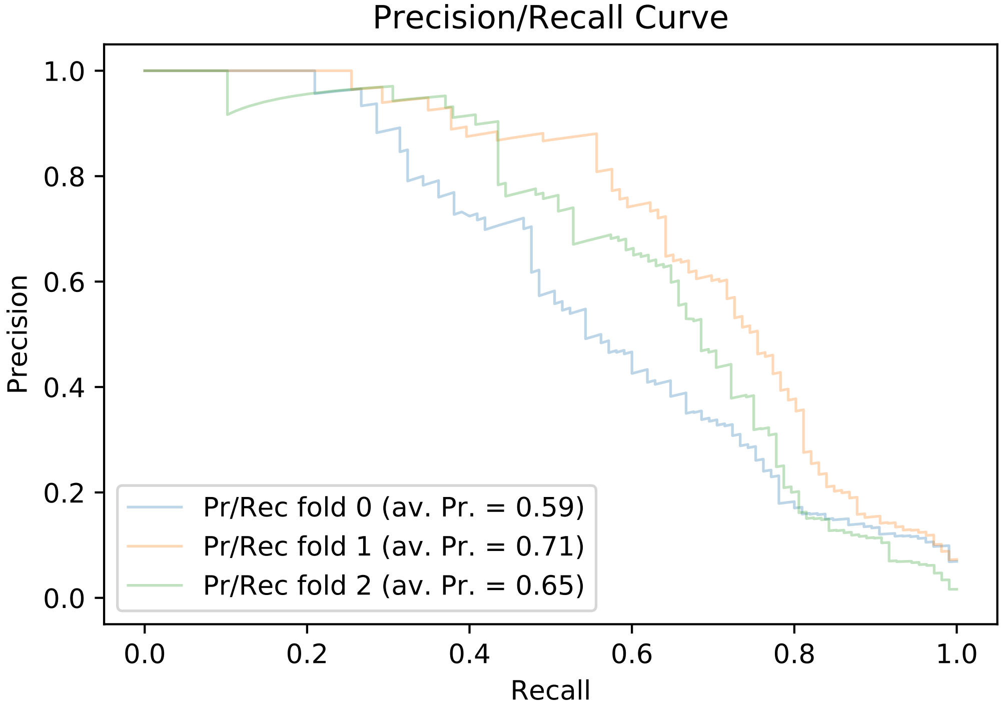

Notes: Figure plots recall on the x-axis. Recall is the true positive rate capturing the share of all actual articles with violence that are correctly classified by the algorithm. The vertical axis plots out precision, which is the share of all correctly classified articles indicating violence against tourists among all articles that the algorithm classifies as indicating violence against tourists.

Figure B5: Predicted conditional probabilities of class labels after computing ensemble: clear separation of data is achieved

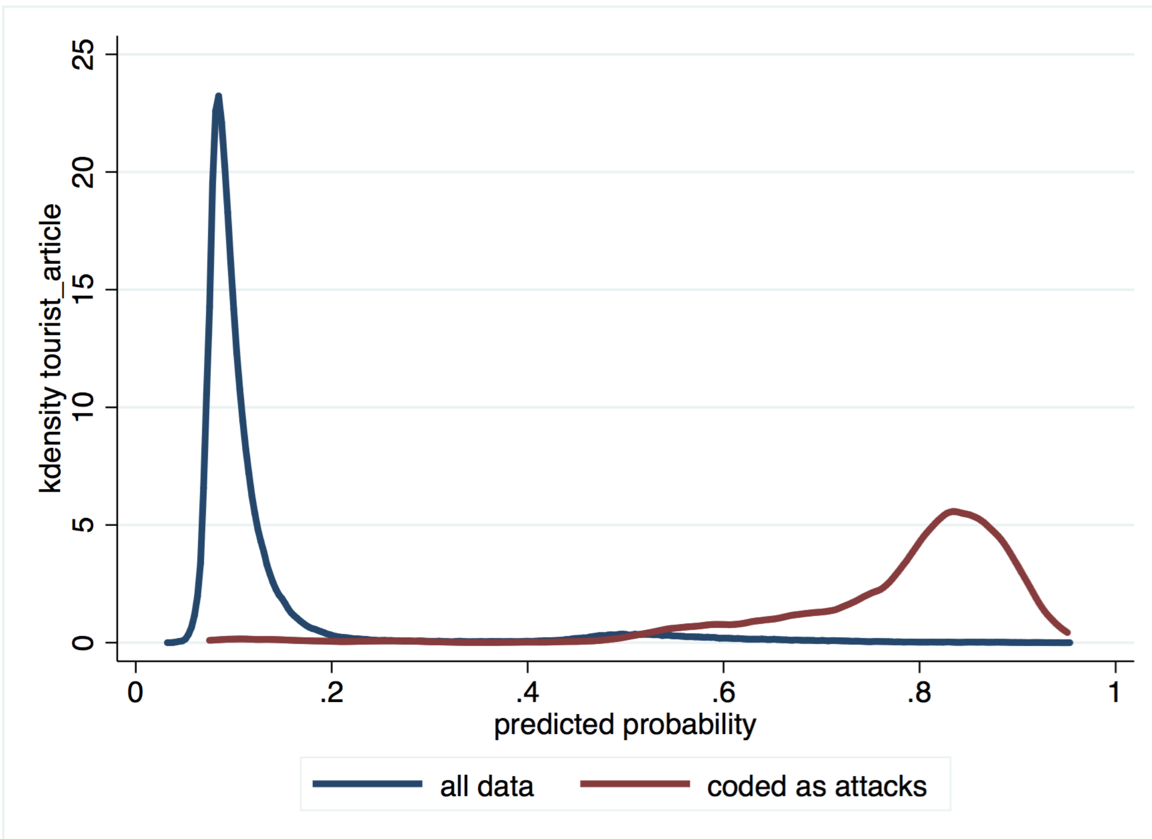

Notes: Figure displays the kernel densities associated with the empirical distributions of the predicted class labels after the ensemble. The data are clearly separated between the two classes.

Table B1: Newspaper coverage sample and sources included

Country	Main Source Name	Article count	Language	Main Source	Flag
ARE	Gulf News	4712	arabic	LexisNexis	
ARG	Source: La Nación (Argentina, Spanish Language)	822	spanish	Factiva	
AUS	Syndey Morning Herald	4492	english	LexisNexis	
AUT	Der Standard	5901	german	Factiva	
BEL	Agentschap Belga (Belgium, Dutch Language)	2140	dutch	Factiva	
BHR	Akhbar Al Khaleej.com (Bahrain, Arabic Language)	20243	arabic	Factiva	agency
BRA	O Globo	3534	portuguese	LexisNexis	
CAN	The Toronto Start	4563	english	LexisNexis	
CHE	Neue Zürcher Zeitung	4427	german	Factiva	
CHL	La Nación (Chile, Spanish Language)	2707	spanish	Factiva	
CHN	Xinhua News Agency	70554	english	LexisNexis	agency
COL	El Tiempo (Colombia, Spanish Language)	593	spanish	Factiva	
CYP	Cyprus Mail	4143	english	Factiva	
CZE	CIA Daily News	920	english	LexisNexis	
DEU	Die Welt	5380	german	Factiva	
DNK	Politiken / Politiken Weekly	6186	danish	LexisNexis	
ESP	El País	29187	spanish	Factiva	
EST	Baltic Business Daily	272	english	Factiva	
FIN	Helsinki Times	736	english	LexisNexis	
FRA	Le Figaro	16072	french	Factiva	
GBR	Daily Telegraph	5755	english	Factiva	
GRC	Athens News Agency	2334	english	Factiva	
HKG	South China Morning Post	855	english	Factiva	
HRV	HINA (Croatia)	723	english	Factiva	
HUN	MTI - EcoNews (Hungary)	736	english	Factiva	
IND	Hindustan Times	3861	english	Factiva	
IRL	The Irish Times	5352	english	Factiva	
ITA	Correrie della Sera	8682	italian	LexisNexis	
JPN	The Tokyo Shimbun	137	japanese	Factiva	
JOR	Addustour (Jordan, Arabic Language)	38402	arabic	Factiva	
KOR	Chosun Ilbo	1129	korean/english	Factiva	
KWT	Kuwait News Agency (Arabic Language)	37195	arabic	LexisNexis	agency
LBN	Tayyar.org (Arabic Language)	10977	arabic	Factiva	
LTU	Lithuanian News Agency - ELTA	969	english	Factiva	
LUX	Tageblatt (Luxembourg, German Language)	2777	german/french	Factiva	
LVA	Vesti Segodnya (Latvia, Russian Language)	793	russian	Factiva	
MEX	Reforma (Mexico, Spanish Language)	1511	spanish	Factiva	
MYS	Berita Dalam Negeri	745	malay	Factiva	
NLD	De Telegraaf	4691	dutch	Factiva	
NZL	The New Zealand Herald	1935	english	LexisNexis	
OMN	Al Shabiba (Oman, Arabic Language)	9201	arabic	Factiva	
PHL	Manila Bulletin (Philippines)	1354	english	Factiva	
POL	Gazeta Wyborcza	738	polish	Factiva	
PRT	Jornal de Notícias	1870	portuguese	Factiva	
QAT	Qatar Tribune	446	english	Factiva	
ROM	AGERPRES (Romania)	1212	english	Factiva	agency
RUS	RIA Novosti (Russia, Russian Language)	55008	russian	Factiva	agency
SAU	Arab News	2655	english	Factiva	
SGP	The Straits Times	1395	english	Factiva	
SVK	TASR - Tlacova Agentura Slovenskej Republiky	407	slovak	LexisNexis	
SVN	STA (Slovenia)	757	english	LexisNexis	
THA	The Nation (Thailand)	1352	english	Factiva	
TUR	Dunya (Turkey, Turkish Language)	27786	turkish	Factiva	agency
TWN	Liberty Times (Taiwan, Chinese Language - Traditional)	1180	chinese	Factiva	
UKR	Delo.ua (Ukraine, Russian Language)	2387	russian	Factiva	
USA	New York Times	18783	english	Factiva	
ZAF	Cape Times	2989	english	Factiva	

Notes: Table presents the names of the main newspaper sources used by country in the paper, along with the original source language and the number of articles covered.

Table B2: Example News Headlines coded as covering violence with fatalities

Country	**Year**	**Month**	**Headline**	$\hat{P}_k(Y_i = 1 \mid D_i)]$
TUR	2011	8	One soldier killed in clash with PKK rebels in southern Turkey	0.987
TUR	2012	10	3 police officers killed in clashes with PKK in Turkey	0.987
EGY	2015	2	17 killed in security raids in Egypt's Sinai	0.987
EGY	2014	4	7 extremists killed, 20 injured in Egypt's Sinai raids	0.987
TUR	2016	9	Two soldiers killed in clashes with PKK in SE Turkey	0.986
TUR	2012	8	2 PKK members killed in southeast Turkey	0.986
TUR	2015	8	2 soldiers killed in PKK attack in SE Turkey	0.986
TUR	2016	9	2 soldiers killed in clash with PKK in SE Turkey	0.986
TUR	2012	10	6 PKK members killed in operation in SE Turkey	0.986
TUR	2016	6	6 soldiers killed in PKK attacks in SE Turkey	0.986
TUR	2016	3	4 soldiers killed in PKK bomb attack in SE Turkey	0.986
TUR	2012	10	3 PKK rebels killed in clash in eastern Turkey	0.986
TUR	2012	8	21 killed, 7 wounded in clashes after mine blasts in SE Turkey	0.986
TUR	2012	11	5 PKK rebels killed in military operation in SE Turkey	0.986
TUR	2013	1	One soldier killed in clashes in SE Turkey	0.986
EGY	2013	9	1 soldier killed, 9 injured by militants in Egypt's Sinai	0.986
TUR	2015	10	3 soldiers killed in clashes with PKK in SE Turkey	0.986
EGY	2013	7	3 terrorists killed in car bomb explosion in Egypt's Sinai	0.986
EGY	2014	9	18 extremists killed in security raid in Egypt's Sinai	0.985
TUR	2012	11	5 Turkish soldiers killed in clash with PKK militants	0.985
EGY	2013	7	2 policemen killed by extremists in Egypt's Sinai	0.985
TUR	2012	12	42 PKK militants killed in eastern Turkey	0.985
TUR	2016	4	1 soldier killed in PKK bomb attack in SE Turkey	0.985
EGY	2015	9	2 killed in suicide car bombing in Egypt's Sinai	0.985
EGY	2013	9	Several militants killed in military raid in Egypt's Sinai: security source	0.985
TUR	2016	4	2 soldiers killed in PKK bomb attack in SE Turkey	0.985
EGY	2013	9	Urgent: Several militants killed in military raid in Egypt's Sinai: security sou	0.985
EGY	2015	2	15 extremists killed in security raid in Egypt's Sinai	0.985
TUR	2012	7	1 Turkish soldier killed, 3 wounded in clashes with PKK	0.985
EGY	2013	8	Urgent: 5 soldiers killed, 8 injured by gunmen in Egypt's Sinai	0.985
EGY	2013	7	Urgent: 2 policemen killed by extremists in Egypt's Sinai	0.985
TUR	2016	3	Update: 4 soldiers, 1 policeman killed in PKK attacks in SE Turkey	0.985
TUR	2010	8	Five PKK rebels killed in clash in southeast Turkey	0.985
EGY	2013	9	9 militants killed in Egypt's Sinai raid: army	0.985
TUR	2012	10	3 soldiers killed in PKK attacks on outposts	0.985
TUR	2011	10	Village guard killed in clash with PKK in southeast Turkey	0.984
TUR	2012	8	4 soldiers killed, 2 wounded in mine blast in SE Turkey	0.984
EGY	2015	10	Police killed in blast in Egypt's Sinai	0.984
EGY	2014	6	8 extremists killed in security raids in Egypt's Sinai	0.984
TUR	2012	7	15 PKK members killed in clashes with troops in southeastern Turkey	0.984
EGY	2013	8	25 policemen killed in attack in Egypt's Sinai: official	0.984
TUR	2016	9	5 soldiers killed, 6 wounded in PKK attack in SE Turkey	0.984
EGY	2013	9	Urgent: 1 soldier killed, 9 injured by militants in Egypt's Sinai	0.984
EGY	2013	7	2 policemen killed by gunmen in Egypt's Sinai	0.984
EGY	2015	7	5 soldiers killed in Egypt's north Sinai in clash with IS branch	0.984
TUR	2012	12	3 PKK members killed in eastern Turkey	0.984
TUR	2011	9	One policeman and wife killed by PKK in eastern Turkey	0.984
TUR	2012	6	Two killed in clashes in southeastern Turkey	0.984
TUR	2016	7	3 police killed in PKK bomb attack in SE Turkey	0.984
TUR	2016	3	26 PKK militants killed in SE Turkey	0.984

Notes: Table presents some example headlines of articles that are classified as being covering violence along with the estimated $\hat{P}_k(Y_i = 1 \mid D_i)]$.

Table B3: Example News Headlines coded as covering violence

Country	Year	Month	Headline	$\hat{P}_k(Y_i = 1 \mid D_i)]$
TUN	2015	3	Spanish couple escapes Tunisia attack by hiding in cupboard for 23 hours	0.948
TUN	2015	6	Kuwait Embassy in Tunisia: no Kuwaiti nat'ls in Tunisia terrorist attack	0.934
TUN	2015	6	Urgent: Armerd men attack Sousse hotel in Tunisia	0.889
TUN	2015	6	Austrian Chancellor's expresses sorrow over Kuwait, Tunisia and France attacks	0.881
TUN	2015	6	Tunisia apprehends culprits behind Sousse resort attack	0.870
TUN	2015	3	1st LD: 19 killed, including 17 tourists, in Tunisia's museum attack: PM	0.865
EGY	2012	2	Three South Korean tourists held by locals in Egypt's Sinai, kidnapper identifie	0.862
TUN	2013	11	Suicide bomber targets top Tunisian tourist destination	0.861
TUN	2016	3	Roundup: Jihadist attacks shiver Tunisia's calm, eliciting casualties	0.857
EGY	2014	2	S. Korea censures terrorist attack on tourist bus in Egypt	0.846
MAR	2011	4	Sarkozy condemns Marrakech attack	0.844
TUN	2015	6	Thousands of European tourists are evacuated from Tunisia	0.835
TUN	2015	3	8 tourists killed in Tunisia museum attack	0.835
TUN	2015	11	A new attack is enraged with the Tunisian transition	0.832
EGY	2014	2	Urgent: Tourist bus explodes in Egypt's Taba, casualties feared	0.828
TUN	2015	6	Thousands of European tourists are evacuated from Tunisia ¡¿	0.826
TUR	2016	6	A suicide attack causes at least 36 deaths at the Istanbul airport	0.825
TUN	2015	3	Slovak gov't sends plane to evacuate Children's Folk Group from Tunisia	0.818
TUR	2016	1	The jihadist attack on the hotel in Burkina causes 23 dead	0.817
EGY	2012	2	Three South Korean tourists held by locals in Egypt's Sinai	0.814
TUN	2015	3	Third French tourist probably killed in Tunis attack: Hollande	0.814
TUR	2016	6	A suicide attack causes at least 28 deaths at the Istanbul airport	0.808
TUN	2015	6	Thousands of visitors are evacuated from Tunisia after the attack	0.806
TUN	2015	3	Two Spanish pensioners die in the attack against the Bardo Museum	0.805
TUN	2015	6	Bloody Friday Jihadism shows its cruelty in the attacks in Tunisia Lyon and Kuwa	0.803
TUN	2015	3	Feature: Italy mourns four victims in Tunisia's museum attack	0.798
TUN	2015	6	Tunisia's transitional priority target of terror	0.794
TUN	2015	3	We thought we were going to die, we've had a terrible time	0.793
TUN	2015	3	Roundup: Tunisia tries to restore national image after deadly museum attack	0.792
TUN	2015	3	Tunisia ... Hostages taken after attack at Bardo museum	0.791
EGY	2012	2	Urgent: Egypt's Bedouins release three South Korean tourists	0.791
TUN	2015	6	Gunman Focused on Tourists in Slaughter at a Tunisian Beach Hotel	0.786
TUN	2015	6	I could hear the bullets whining Gary Pine English tourist on the beach in Souss	0.786
TUN	2015	3	2nd LD: 21 killed, including 17 foreigners, in Tunisia's museum attack: PM	0.783
TUN	2015	3	Militants hold tourists hostages inside Tunisia museum	0.780
TUN	2015	6	Irish woman among fatalities in Tunisia attack	0.762
TUN	2015	6	Scores Die in Attack at Tunisian Beach Hotel	0.758
TUN	2015	6	5th LD: Death toll rises to 37 in catastrophic hotel attack in Tunisia	0.756
TUN	2015	6	Norway condemns attacks in Tunisia, France, Kuwait	0.753
TUN	2015	3	Belgium to open own investigation into Tunisia attacks	0.752
TUN	2015	3	Bardo museum reopens a week after killings; Tunisia sends out message country s	0.744
TUN	2015	6	4th LD: 28 killed, 36 injured in terror attack on Tunisia hotel	0.739
TUN	2015	6	Thousands of visitors are evacuated from Tunisia after the attack ¡¿	0.735
TUN	2015	3	The attack of the Bardo museum in Tunisia .. What do we know about the nationali	0.734
TUN	2015	6	Germany condemns deadly hotel attack in Tunisia	0.732
TUN	2015	6	After Tunisia attack, UK ups Wimbledon security	0.729
TUN	2015	3	Hollande expresses solidarity with Tunisia after deadly attack	0.726
TUN	2015	6	Deaths of British nationals in Friday's attack in Tunisia rise to 15: FCO	0.725
EGY	2012	2	1st LD Egypt's Bedouins release three South Korean tourists	0.717
TUN	2015	6	3rd LD: Terrorist suspect in Tunisia's hotel attack arrested: official	0.711

Notes: Table presents some example headlines of articles that are classified as being covering violence along with the estimated $\hat{P}_k(Y_i = 1 \mid D_i)]$.