

Manuscript version: Author's Accepted Manuscript

The version presented in WRAP is the author's accepted manuscript and may differ from the published version or Version of Record.

Persistent WRAP URL:

<http://wrap.warwick.ac.uk/181695>

How to cite:

Please refer to published version for the most recent bibliographic citation information. If a published version is known of, the repository item page linked to above, will contain details on accessing it.

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk.

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

Timothy Besley

Thiemo Fetzer

Hannes Mueller*

November 11, 2023

Abstract

This paper estimates the size of the *media multiplier*, an easily generalizable model-based measure of how far media coverage magnifies the economic response to shocks. We combine monthly aggregated and anonymized credit card activity data from 114 card issuing countries in 5 destination countries with a large corpus of news coverage in issuing countries reporting on violent events in the destinations. To define and quantify the media multiplier we estimate a model in which latent beliefs, shaped by either events or news coverage, drive card activity. According to the model, media coverage can more than triple the economic impact of an event. We document, through our model, that this effect is highly heterogenous and depends on the broader media representation of countries in each others news. We speculate about the role of the media in driving international travel patterns an.

JEL Classification: O1, F5, D8, F1, L8

Keywords: media, economic behavior, news shocks

*Besley is based at the London School of Economics, Fetzer at Warwick University and University of Bonn and Mueller is based at IAE (CSIC) and the Barcelona GSE. This work was conceived in 2014. Insane. 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 financial support from grant PID2021-124256OB-I00 funded by MCIN/AEI/10.13039/501100011033 and by ERDF *A way of making Europe*, and the Severo Ochoa Programme for Centres of Excellence in R&D (CEX2019-000915-S). Fetzer acknowledges funding by the European Union (ERC, MEGEO, 101042703). Views and opinions expressed are however those of the author only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them. acknowledges support from grant 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 and political behavior.¹ These effects could be particularly pronounced for dramatic, news-worthy events such as protests, major crimes, violence, accidents or public health scares where individual's 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 amplifying effect of negative news shocks.

This paper introduces a model of the *media multiplier*, which measures how much media coverage amplifies the behavioral response to shocks. The core idea in the model is that it is the process of belief formation that determines how the news environment translates into behavior and hence affects economic activity. In particular, we model how events affect economic behavior because they change what agents believe about an underlying, but largely unobserved state. We estimate this model in the context of sudden outbreaks of violence, which we hypothesise change behavior by altering risk perceptions. Our estimates suggest that, in the context of terrorism and violence, assessments of risk get amplified by a factor of over 3 with consistent, long-lasting negative news coverage. However, the actual effect depends on how intensive negative reporting is and whether or not it is drowned out by other "unrelated" news coverage.

We analyze the impact of violent events on credit card activity in five travel destinations: Turkey, Egypt, Tunisia, Israel and Morocco. The data is dis-aggregated so that we

¹See Strömberg (2015) and Prat and Strömberg (2011) for reviews. See Besley and Burgess (2002) for evidence on government accountability, Stromberg (2004) on redistributive spending, Eisensee and Strömberg (2007) on US disaster relief, 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 and Durante et al. (2019) on the proclivity towards populist rhetoric.

²In 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*.

measure spending by tourist *origin* in each of these five *destinations*.³ A unique feature of our approach is that we analyze spending responses of tourists from different origin countries in different destination countries in response to differential media reporting in tourist origin countries on the same set of violent events occurring in destinations.⁴ The direct economic effect that an event has on economic outcomes in a destination can therefore be separated from the impact of the intensity of news coverage across the card origin countries. This constitutes 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 news coverage in different news environments reporting on what is happening in the same destination.

One contribution of the paper is to construct dyadic – or network data – that has an origin and destination dimension both on economic activity – tourism as a form of service sector trade – and news coverage.⁵ To achieve this, our data set combines monthly aggregated and anonymized credit card activity data – measured through the number of active cards and the amount spent – from 114 origin countries at each destination with a corpus of news coverage of the destinations across the different origins. We analyze more than 450,000 news articles from 57 different origin countries to track how media outlets cover the same violent events in the five destination countries by using a machine learning approach that enables us to identify news coverage of fatal violence and tourist harm. Specifically, 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

³This terminology labelling countries as either origin or destination countries is adopted for convenience.

⁴A part of the spending response we observe will be driven by tourists in the past not booking travel. We therefore use the terms *potential tourist* and *tourists* interchangeably.

⁵A dyad for this purpose is an origin and destination pair.

news. Our 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. \(2018\)](#).⁶ 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, can be extremely valuable even in challenging settings.

(Figure 1)

The empirical regularity that we study in this paper is best illustrated visually in [Figure 1](#) which shows data from Tunisia as a destination country. To construct the figure, we identify two types of origin countries according to their news coverage of violent events in Tunisia: those that report more intensively compared to those that report relatively less. For these two types of origin countries the figure then plots – after removing a set of dyad, origin-by-time and destination-by-time fixed effects – the average residuals measuring card spending in panel a), and the number of active cards – a proxy measure for the number of tourists – in panel b). Tunisia saw a persistent increase in frequency of violence from 2013 onward which is indicated by the vertical line. We observe that after 2013, there is a clear shift in the relative spending patterns and in the number of active cards, i.e. those being used to make transactions in the country. For the group of countries that, on average, report more on violence in Tunisia, we see that the dashed line falls more sharply when the violence starts and only recovers slowly towards the end of the sample – relative to countries whose media reports less on violence in Tunisia. This highlights the variation that we can exploit using our dyadic dataset which will allow us to document a robust reduced-form differential impact of violent events, driven by different intensities of news coverage. We then develop a theoretical model conceptualizing the data generating process which we fit to the data.

⁶The main advantage is that we can use cross-validation to show that an ensemble of naïve Bayes, XGboosted trees and random forests provides relatively large gains over dictionaries or simpler methods (see [Manela and Moreira, 2017](#); [Becker and Pascali, 2019](#)).

This permits a formal definition of the *media multiplier* which characterizes the news-coverage induced amplification effect that this figure illustrates.

Our theoretical approach follows [Besley and Mueller \(2012\)](#) who develop a model of latent beliefs to examine the reaction of house prices to violence in Northern Ireland. The key insight that we build on is that changes in economic activity are mediated by how agents form beliefs about the risk of *future* violence. This implies that some changes in violence will have little effect on activity whereas others have a dramatic effect. Such models of beliefs as mediators have already served to better explain variation in transport costs ([Besley et al., 2015](#)), child health in conflict zones ([Tapsoba, 2023](#)) and educational choice [Alfano and Görlach \(2022\)](#). Here we look at the importance of news coverage for the formation of these beliefs explicitly in a unified model. We also show that news-based beliefs in countries with free media seem to provide a much better explanation for changes in card activity, suggesting that tourists rely more on the news coverage in their country of origin when they form beliefs about the prospect of violence before making travel plans. The theoretical model not only provides a way of interpreting heterogeneity and the time path of media reporting compared to a reduced-form approach but also allows us to project the implications of our findings to other contexts as well as considering counterfactual paths. In our model, news reporting matters because it is used by some agents to update their beliefs on what is happening in a destination. It is the prevalence of agents who rely exclusively on media reports when forming their beliefs that generates a media multiplier.

The approach that we take contributes to a wider 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 \(2017\)](#) show 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. Fetzer et al. (2020) show how the extent and intensity of media reporting on own-country – compared to other NATO casualties – in Afghanistan affects public support for continued military engagement across the different troop sending countries. Alfano and Görlach (2022) show that access to media content amplifies and spreads the impact of terror attacks on schooling choices. In all of these cases, changing beliefs is the mechanism behind changes in behavior.

However, none of these papers has an underlying model of how beliefs are formed.⁷ Our model-based approach suggests that when a country is perceived as dangerous by all potential visitors, card activity falls by close to 60 percent, with more than two thirds of this due to the media multiplier. The closest available data on such an effect comes from the COVID-19 pandemic where lockdowns imposed legal barriers to economic activity. The hypothetical numbers from our model are close to the effects of these lockdowns found by Carvalho et al. (2020). In other words, our model suggests that persistent mortality risk along with persistent news coverage could impose economic damage similar to a lockdown during a pandemic. In a somewhat speculative out-of-sample prediction, we estimate the consequences of such effects being extended to entire regions in Asia and Africa.

The idea that changes in expectations are the mechanism through which news affects economic activity is a central pillar of work on “news shocks” in macroeconomics (Ramey, 2011). Arezki et al. (2017) document that news reports on resource discoveries

⁷Alfano and Görlach (2022) are closest in spirit to our model. They estimate a scaling factor on the reaction to violence by households which they call the *fear factor*, π_i , and distinguish households with and without access to the media. The implicit media multiplier they derive in this way is around 12 and the authors argue that this is due to an overestimate of violence risk on the part of the households with media access.

have a 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 Fourniaies \(2014\)](#) 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 ([Glick and Leduc, 2012](#)). Our model of heterogeneous beliefs is able to explain large swings in spending which are macro-economically significant. Back of the envelope calculations based on our estimates indicate losses between 2011 and 2016 of over 38 billion USD due to violence with close to 12 billion USD being due to the additional effect of negative news reporting.

We also find a *drowning out* effect in which other, neutral news can lower the impact of negative news. Conversely, this implies that destinations that are not typically covered in the news experience larger economic responses to violent-event induced news shocks. This is consistent with behavioral mechanisms suggesting that individuals may use heuristic approaches or simple mental models to evaluate information.⁸ We demonstrate the possibilities offered by a calibrated model by simulating the effect of negative news shocks with various levels of other, “neutral” news. The resulting strong, nonlinear response of beliefs to events in Tunisia, a country with a low baseline level of media coverage, is relevant for other contexts as well. International news coverage of most developing countries is rare and tends to follow adverse events. Our model suggests that persistent negative news without any other news to counterbalance it can explain stickiness in perceptions. Given the economic and social benefits of openness and economic integration in general (see [??, Fra; Melitz and Trefler, 2012](#)) and the prominence of tourism as a development strategy for some countries in particular (see [Faber and Gaubert, 2016](#)), the way media reports on countries could therefore have detrimental consequences for economic development. This is particularly relevant in

⁸See [Azeredo da Silveira and Woodford \(2019\)](#); [Bordalo et al. \(2016, 2018\)](#); [Fetzer et al. \(2020\)](#); [Handel and Schwartzstein \(2018\)](#); [Zhu et al. \(2020\)](#) for a suite of behavioral models and related empirical evidence.

the Middle East and North African region, which is one of the least economically integrated regions (Rouis and Tabor, 2012) and where growing economic ties to Europe are potentially important.

The drowning out effect could also 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 (terrorists) the impact of their actions.⁹ This generates incentives to produce good or at least neutral news when faced with negative news - a strategy that has been found to work for individuals as well (Lewandowsky et al., 2020).¹⁰

The remainder of the paper is organized as follows. The next section – section 2 – discusses 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 3, we present reduced form evidence. In section 4, we propose a statistical model and fit this on both the news and the tourism activity data. Concluding comments are provided in section 5.

2 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.

⁹We observe in our sample that news coverage sharply responds to violent events, but does not increase in anticipation.

¹⁰Countries aware of these mechanisms may also adopt strategies to change their outside perceptions and build “soft power”. In a service sector trade escalation – similar to trade escalations around goods trade (see Fetzer and Schwarz, 2021; Fetzer et al., 2023) – narratives and spin can be used to shape public opinion and may be weaponized to foster both geoeconomic, geopolitical or domestic policy interests more broadly. Poorly governed social media may provide an open door for such malign interference, stoking divisions in more heterogeneous and in particular liberal societies (Müller and Schwarz, 2021).

2.1 Aggregated Spending Data

Mastercard provided us access to an anonymized and aggregated monthly data set, which includes an index measure capturing the number of transactions, the number of active cards and spending levels in five different countries (Egypt, Israel, Morocco, Tunisia and Turkey) broken down by the country of origin of the card, where the latter is identified based on the financial institution that issued the card. The data is for the period 2010 to 2016, a total of 84 months. Where Mastercard’s confidentiality or disclosure controls may have resulted in, for example, a blank monthly-dyad observation, e.g. due to a low count of observations, 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 5 destination dyads. The origin countries in our sample span all continents but tend to be higher income countries and those that are geographically closer to the destination countries.¹¹ Figure A1 maps all of the origin (card-issuing) countries that we have in the sample.

There are notable differences between origin countries with a low volume of cards active per month in tourism spending in countries such as Haiti and Namibia compared to higher volume tourism spending from countries such as Germany and the United States. The type of research that we do here would not be feasible using data that statistical agencies make available – which is typically patchy, aggregated and annual at best. Yet, for a small sample of countries, we can validate the Mastercard data by aggregating it to the level that official tourism statistics are published. Specifically, a small set of countries that produce annual data on travel flows is provided to the United Nations World Tourism Organisation (UNWTO) broken down by the most important tourist sending countries. Appendix Figure A3 highlights that the annually-aggregated

¹¹We also include cards originating in the destination countries themselves although these data can be dropped without affecting the results.

card data correlates well with the travel flows data. Appendix Table A2 presents further regression evidence highlighting the closeness of the fit.¹²

2.2 Data on Violent Events

Data on violent events serves as an important anchor dataset capturing whether an event is occurring in a location at a point in time. To achieve broad coverage, we leverage five different conflict event data sources, three of which are hand-coded event data while the other two are constructed using information retrieval methods.

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 (2020), those datasets identify conflict events from news articles by processing raw text leveraging

¹²We cannot validate this for all countries since arrival statistics are typically only provided for larger countries. That said, our results are robust to the composition of tourist origin countries being considered.

¹³Results 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 (Egypt, Tunisia, Morocco) we do not include them in the analysis.

natural language processing techniques that, in essence, extract three pieces of information: 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, observing every moment of the day stretching back to January 1, 1979 to produce data on events.¹⁴ GDELT is more inclusive, but 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 of violent events based on objective criteria. Neither of them provide 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.

2.3 Data on News

The news data variable that we construct is intended to proxy the news coverage that potential travelers 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

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

¹⁵To the best of our knowledge, such a dataset does not exist.

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 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

2.4 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

¹⁶Appendix 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.

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.

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 E. 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 $\hat{P}_k(Y_i = 1|D_i)$, 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 and are 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 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 $\hat{P}_k(Y_i =$

$\mathbb{1}(D_i)$ according to the function:

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

To choose 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) across 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, and with a rapidly declining rate, so that we suspect the remaining arti-

¹⁷Our 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.

¹⁸We 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.

cles 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 E, we describe the classification approach in greater detail, while Appendix Tables B2 and B3 provide some sample headlines of articles coded as covering violence with fatalities and flagged up as capturing that tourists are targeted by a violent event, respectively. 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 respectively, 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.

2.5 Patterns in the Reporting Data

We approach news reporting on violent events as standing in 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

$$n_{hdt-1} = \frac{B_{hdt-1}}{N_{hdt-1}} \quad (2)$$

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. Therefore, 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 four 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 5.9 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:

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

where k 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 news 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 aggregate our news measures to the monthly level, as the dyadic card spend data is only available at the monthly level. Figure A6 reports the mean shares at the 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 (2) for violence against tourists – the dashed line with the axis on the right hand side – and (2) for fatal violence on a monthly basis – the solid line with the axis on the left hand side.

Figure A6 indicates that there is 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 overall reporting on Tunisia. At the time of the Sousse attack, for example, violence against tourists occupied around 40 percent of all news about Tunisia. Reporting in Egypt, Turkey and Israel is

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

more intense for fatal violence than it is for violence against tourists. However, this coverage rarely occupies more than 10 percent of the news. The most extreme example is Israel where news on violence never exceeds 13 percent of reporting and violence against tourists never more than 3 percent.

3 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 tourists 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

3.1 Core Findings

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

$$y_{hdt} = \alpha_{hd} + \alpha_{ht} + \alpha_{dm(t)} + \tilde{\zeta}_1 n_{hdt-1} + \tilde{\zeta}_2 v_{dt-1} + \varepsilon_{hct} \quad (3)$$

where the origin countries are denoted h , destination countries are denoted d and time is denoted t . The dependent variable, y_{hdt} , is a measure 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/calendar month fixed effects, $\alpha_{dm(t)}$, to account for destination-specific seasonality in demand. In some specifications we will use

²⁰Appendix 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.

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 reported 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.

Prior to presenting the results, it is worth stressing that our analysis is only exploiting within-dyad variation which absorbs all factors such as 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 . This implies that all magnitudes are based on comparing the attractiveness of each destination relative 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. This is a somewhat conservative approach since there could be reputational externalities which lead potential travelers to shy away from the entire region due to the turmoil in one of the five destination countries that we study.

Finally, it is important to note 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

²¹Appendix 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 a broad measure of objective measured violence in a destination.

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 3.3 that idiosyncratic factors, such as the cultural proximity to the victims of an attack, drive a lot of this dyad-specific variation.

3.2 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 dependent 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 0.5 log points, or 60 percent. 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.

(Table 1)

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 origin 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 0.18 log points or 20 percent. This continues to suggest that a significant part of the overall effect in columns (1) and (2) is driven by pure news report-

ing. 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 spending is less, but also at the extensive margin as fewer tourists travel to a destination 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 in magnitude compared to those presented in Panel A. We also find statistically and economically significant relationships throughout, although we do lose statistical significance in the most saturated specification shown in column (6). The results suggest that differential intensities in media coverage of violence may have an important independent 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

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 4 and to test robustness of the reduced form results, we implemented a set of additional tests.

²²A 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.

3.3 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 most suitable way of measuring news in the reduced-form exercise. In section 4, 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 and immediate negative effect which becomes stronger when we add longer lags. This is not surprising given that tourists book travel in advance and will not book travel to a place with negative reporting at the time of booking. Importantly, we also find no pre-trends in card activity before the news events that we study.

(Figure 2)

However, Figure 2 also highlights that a simple reduced form regression analysis as in specification 3 is unlikely to capture the overall dynamic impact of news coverage completely 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 4 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 There are two main challenges to the view that the intensity of reporting on violent incidents is a source of exogenous variation for studying the impact of media reporting on tourism. The first is the possibility of strategic reporting whereby the level of media coverage is scaled up or down in a given country to reflect the likely level of public interest in a specific violent incident.

Moreover, this could be correlated with the extent to which a country is a popular tourist destination for that country. The second is that those who commit acts of violence are motivated by getting media coverage in specific countries to whom they wish to convey a political message. Both of these considerations lead to the possibility that there is a correlation between incident reporting and the error term in (3). However, in either case, the direction of bias is not clear *a priori*.

To address such concerns, we present results from an IV approach in Appendix Table A6. This exploits 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, it is not clear that risks for the specific nationalities have suddenly increased. So the idea is that such an instrument is relevant in driving media-reporting and plausibly excludable as the distribution of nationalities of casualties is random conditional on fixed effects and the violent event having occurred.

We can also relax this identifying assumption by focusing exclusively on *media reporting spillovers* where reporting is higher for other origin countries, which are culturally similar.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 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 there are strong reporting spillovers through shared

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

language and contiguity.²⁴ The instrument provides a Weak IV identification F-statistic above 10 for all of the exercises. In the second stage, we find that point estimates on the reporting share increase relative to our reduced form estimates in Table 1, which 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 or not a news story is picked up in a home country that leads to a response in card activity.

Further Robustness We have also conducted 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 checks. 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). Furthermore, the 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.

4 Modelling the Impact of News Coverage on Card Activity

In this section, we posit a model in which agents form beliefs about a latent binary state (dangerous or safe). News and violence only matter to the extent that agents dislike booking travel to dangerous locations because these locations tend to stay dangerous. Our model supposes that there are two types of tourist that together generate

²⁴The results are presented in Appendix Table A6.

aggregate behavior. One type uses *event-based* beliefs about violent events based on information databases, while *news-based* beliefs come from observing domestic media reports. The reaction of card activity to these two different ways of forming beliefs is the basis on which we define and quantify the media multiplier. We construct a fitted model and show that this model of beliefs provides a good fit to the pattern of card activity and spending. It is also useful for out-of-sample prediction exercises that we illustrate.

4.1 The Model

Model Overview Figure 3 presents a graphical overview of the model inspired by plate notation common in machine learning. 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 is based on data about violent events, v_{dt} , and the latter is based on news coverage, i.e. bad news, B_{hdt} , and all news, N_{hdt} .

(Figure 3)

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. We also 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 , whereas q_d captures the persistence of safety.

The state of danger is central to our model 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 the parameter η_s . We assume that all agents, including those that base their beliefs on domestic 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 the likelihood that there will be violence expectation in 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:

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

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

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, news (B_{hdt} , N_{hdt}) and violent events (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

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

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 by 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 sources such as UCDP, GTD and the GDELT event database that track all available reporting. Since this information is not country-specific, beliefs based on this 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 and depend only on violent events. They provide a benchmark against which to calibrate the media multiplier.

We construct a measure of such beliefs based on data on the history of violent events up to t , Ω_{dt} . To do so, assume that violent events are distributed normally, i.e. $v_{dt} \sim N(\mu_{sd}, \sigma_{sd}^2)$ with mean, μ_{sd} , and variance, σ_{sd}^2 , varying with the state, s_{dt} . At each date, there is a destination-specific transition probability between states, where p_d denotes the probability of transitioning from dangerous to safe and q_d denotes the probability of transitioning from safe to dangerous. This gives a parameter vector for the model with six elements for each destination country d , summarized as

$$\theta_d = \{\mu_{0d}, \sigma_{0d}^2, \mu_{1d}, \sigma_{1d}^2, p_d, q_d\}.$$

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

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

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

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

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

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

is the prior from the previous period. For v_{dt} , we use the principal components across the data from different sources on violent events.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. To do so, we assume that news-based beliefs can also rely 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 reported in Appendix Table A10. They show

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

strong persistence in the state (safe or dangerous) 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 out as the two underlying 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 at any date based on the level of violence at each destination. Thus, unlike equation (13), the effect of a given change in v_{dt} is heterogeneous across different destinations depending on the history and persistence of violence. This makes sense; what would be deemed to be a level of violence indicating that a place is dangerous in a country like Israel (bottom right) is different from what would be considered to be indicative of a dangerous situation in Tunisia (top right). These differences in turn reflect heterogeneous estimates of μ_{sd} . The probabilities based on equation (5) also show that Tunisia was the first of our destinations to become dangerous based on its level of violence and was followed by Egypt and Turkey. A dangerous state in Israel, based on this method, is less persistent.

Spending by individuals with event-based beliefs depends on the beliefs displayed in Figure A7. We test in table A14 whether the kind of non-linearities that the Markov model suggests can be seen in the spending and cards data. Indeed, we find strong evidence for shifts in card activity with the shifts we show in Figure A7 up-and-above the fluctuations that can be explained by the raw violence data. For the number of cards, the raw violence data even becomes insignificant when we add the Markov estimate of the probability of danger, $\hat{\theta}_d$.

Deviations in card activity from event-based beliefs will allow us to understand how large the media multiplier is. However, in order 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 whether a country is dangerous or safe as event-based beliefs and use the same underlying persistence parameters, \hat{p}_d and \hat{q}_d . The key 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 are taking place in a destination country.

As shown in Figure 3, we allow news reporting to react to the underlying state so that 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 choose this distribution to capture the fact that news reporting of B_{hdt} relative to N_{hdt} has fat tails.²⁸ If $\hat{\eta}_1 > \hat{\eta}_0$, there is a higher frequency of “violent news” relative to the total number of news articles when a country is dangerous so that the extent of such news coverage can serve as a signal about whether a destination is dangerous.

²⁷In terms of popular psychological models, this could be interpreted as individuals forming their views using an “availability heuristic” where news at home is more available and salient.

²⁸It could be justified by supposing that there is a “stopping rule” for media consumption but we follow it purely for pragmatic reasons.

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

$$f(B_{hdt}, N_{hdt} | s) = \binom{N_{hdt}}{B_{hdt}} (\eta_s)^{B_{hdt}} (1 - \eta_s)^{(N_{hdt} - B_{hdt})}. \quad (6)$$

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 $\hat{\eta}_s$ are based on the following weighted average:

$$\hat{\eta}_{sd} = \frac{\sum_{hdt} \Pi_{dt} \left[\frac{B_{hdt}}{N_{hdt}} \right]}{\sum_{hdt} \Pi_{dt}} \quad (7)$$

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

We estimate that $\hat{\eta}_1 = 0.02$ and $\hat{\eta}_0 = 0.002$, i.e. the share of violent news in all news is around 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:

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

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

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

which evolves according to the Bayesian recursion:

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

where, as before, $E_{t-1} [\pi_{hdt}]$ is derived from the Markov chain governing the underlying categorical state.²⁹

Discussion of Beliefs When assessing whether this model is empirically successful, it should be remembered that the model is quite specific in the way that it allows news coverage to matter and hence there is nothing built into the model that will *guarantee* a good fit to patterns in the data. Thus, what we present below constitutes 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 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 - we show this below in detail.

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 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

²⁹Appendix 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.

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 it means to be dangerous in different places but by how information from domestic news coverage is processed to update beliefs. This implies that risk perceptions from news can depart quite substantially from risk perceptions that one would get from the event data.

Importantly, we can show that event-based beliefs, Π_{dt} , indeed anticipate violent events whereas news-based beliefs, π_{hdt} do not once we control for Π_{dt} . Appendix Table A13 shows this when forecasting violence one to three months ahead. We find a very strong, robust association of violence with lags of Π_{dt} always but not with π_{hdt} . Event-based beliefs are much more persistent and capture risk's true nature much better.

4.2 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):

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

where χ is the weight on beliefs formed by observing violence. This gives rise to a natural formal representation of the media multiplier as:

$$multiplier_{hdt} = 1 + \frac{1 - \chi}{\chi} \left[\sum_{\tau=0}^{-9} \omega_{\tau} (\pi_{hdt-\tau}) \right] \quad (10)$$

which will be equal to 1 if all news-based beliefs are that a country is safe, i.e. $\pi_{hdt-\tau} = 0$. The maximum value of the media multiplier is $1/\chi$ if all news-based agents think

that a destination is dangerous, $\sum_{\tau=0}^{-9} \omega_{\tau} (\pi_{hdt-\tau}) = 1$. The total effect of a violent event is then $1/\chi$ times larger when negative event-based beliefs get complemented by news-based beliefs. Note that the formulation in equation (10) sets $\Pi_{dt-\tau} = 1$. This prevents situations 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 of violence.

We interpret (9) as modelling the beliefs of a “representative” tourist. But χ and ω_{τ} could also reflect heterogeneity in the population in terms of booking behavior, or the prevalence of tourists that form beliefs based on observing the news or based on observing data on violent events. In that case, the parameter χ would correspond to the share of tourists that form beliefs based on observing data on violent events from wider sources than domestic media.

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 violent events and news reporting. Second, they depend on the entire history of violence rather than on one-period lagged values. Third, they allow for a lag structure to reflect the timing of booking decisions. At the same time, the weights (ω, χ) reflect the importance of different kinds of latent beliefs for overall tourist activity.

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.³⁰ 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 equation (9). 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 0.36 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 fatal violence, rather than events 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.

To estimate the media multiplier, we need to combine news-based and violence-based beliefs in the same regression. We do so in column (4) of Table 2. Now we find a strong amplifying effect of news coverage on behavior 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 interpret this as empirical support for the specific form of news-based beliefs captured by our model and it also provides a theoretical interpretation of the media multiplier reflected in π_{hdt} .

(Table 2)

³⁰This 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.

Columns (5) and (10) use 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 and number of cards data using a grid search over the weights in equation (9) to maximize goodness of fit.³¹ We find that $\hat{\omega}_0 = 0.125, \hat{\omega}_1 = 0.175$ so that 12.5 percent of spending is driven by contemporaneous beliefs and 17.5 percent are coming from the first lag. After that, the weight based on the best fit falls (the weight sequence is 0.075, 0.15, 0.1, 0.125, 0.1, 0.025, 0.05, 0.075). Our estimate of the maximum media multiplier is $1/\hat{\chi} = 3.33$, i.e. the effect of an event is more than tripled if it is accompanied by intense negative media coverage.

Our estimates in columns (5) and (10) imply that if all tourists switched their categorical beliefs that a destination is dangerous from zero to one, spending would fall by over 0.8 log points - close to 60 percent. Of course, this is a 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 would be the case in ongoing terror campaigns in Afghanistan or Iraq or, as 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 columns (5) and (10) to quantify the separate impact of news on spending below.

4.3 Interpreting the Results

We find that card activity and beliefs are strongly correlated with each other, even after controlling for dyad and origin/time fixed effects. We also observed that there is no clear evidence of pre-trends either in reporting or in card activity. The heterogeneity in economic outcomes that we find is almost exclusively related to heterogeneity in beliefs; we do not find, for example, heterogeneity in the spending response to π_{hdt}

³¹See Appendix B and the replication code for more details.

with respect to distance to a destination or the share of Muslims in an origin country (see Appendix Table A12).³²

Our approach identifies the media multiplier from the differential response of card activity to changes in news-driven beliefs. One issue with interpreting this response as a media multiplier is the possibility that risk perceptions might shift on specific dyads with the attack on a certain type of tourist. For example, tourists from Europe or English-speaking countries could react more strongly to an attack because they feel they are a group that is more likely to be targeted by a terror attack. This does present a challenge to our interpretation of the card activity because it implies a shift in dyad-specific, event-based beliefs not news-based beliefs.

That said, this type of shift in card spending is likely to be much more persistent since it would lead to updating whether a destination is dangerous, regardless of the intensity of news reporting. We have shown that beliefs based on news reporting react sharply to events with strong peaks as in Figure A4. We estimate a media multiplier from a response of card activity that reacts sharply to the news shock but recovers quickly within a few months - a response that appears to be inconsistent with an event-based, level-shift in risk evaluations. As discussed above, we provide evidence that news-based beliefs are not good predictors of event risk in Table A13.³³

4.4 The Media Multiplier and Press Freedom

Our results suggest that up to 70 percent of the impact of violent events are accounted for by the media multiplier. But it is well known that that many countries around the world suppress media outlets. This provides an interesting natural experiment and a further test of the role of news coverage on spending since we would expect the media multiplier to be less important in media-repressed countries.

³²The only source of heterogeneity that we find is with respect to press freedom which we turn to in the next sub-section.

³³The reader should also bear in mind the reduced form response we showed in Figure 2.

In the framework of our model, suppression of information through news censorship would affect the use of news-based beliefs, π_{hdt} , by agents, implying that π_{hdt} should be less able to predict responses to violent events in countries with censored media. Even if censorship does not actually repress reporting of violence, citizens may give less credence to media reports in countries where media is censored. To test this idea, we collected data from Reporters without Borders (RSF), which scores countries according to the extent of 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. This allows us to test whether the variable π_{hdt} enters heterogeneously between free and censored countries.

(Table 3)

The results are 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 media – when free to report – 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 number of cards as the dependent variable.

These findings further corroborate the idea that news-based beliefs, π_{hdt} , capture the impact of signals transmitted through domestic news reporting, giving rise to a media multiplier. Even if this is deemed to be an “overreaction” to news about violent events, we would not use this finding to support media censorship. Public trust in signals coming from the news media is a hugely useful asset in most circumstances. Also, there are many beneficial aspects of societies with free media that transcend the issues

that we are dealing with here, with free media being a major hallmark of open societies.

34

4.5 The Media Multiplier and Background News

A key feature of news-based beliefs is that non-violent or background news coverage, $N_{hdt} - B_{hdt}$, affects beliefs. Figure 4 illustrates the role of background news by considering the number of cards 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 column (10) of Table 2 then yields the visitor response represented by the dashed line in Figure 4. According to this, the immediate effect would be that the 12.5 percent of tourists who react immediately to danger ($\omega_0 = 0.125$) do not travel to or leave the destination and tourism spending would fall by around three percentage points.³⁵ The dashed line shows the persistent effect on tourism spending, i.e. the effect of $\Pi_{d,0} = 1$ first increases in the first lag and then falls only slowly because most tourists book their travel in advance.

(Figure 4)

We now show the *additional* impact of the media multiplier via a change in news-based beliefs in a scenario in which the violent event is reported in the news, i.e. $B_{hd0} = 1$. 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 close to 10 percentage points in the first month. In other words, the economic impact of the violent event on spending more than triples due to the media multiplier. Moreover, since there is lead time on tourist bookings, this effect

³⁴However, there is something to be said about sensationalism in the sense that the news response is not consistent with providing a good risk update on fatality risks.

³⁵The exact calculation is $0.125 \times 0.3 \times 0.834 = 3.1\%$.

persists over time.

Contrast this with $N_{hd0} = 100$, i.e. a case where there is a lot of news reporting about a destination that is unrelated to violence. In this case, the news about tourist violence is “drowned out”. Concretely, we find that the contemporaneous effect is now less than 6 percentage points. The reason for this in the model is that news-based beliefs update according to the density in equation (6) and reporting of $B_{hd0} = 1$, $N_{hd0} = 100$ is relatively more likely to have come from a safe country. The size of the media multiplier, therefore, depends not only on the estimated parameters of the model but also on the overall news landscape and the amount of media attention given to a specific destination.

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.³⁶

4.6 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,

$$\hat{\xi} \sum_{\tau=0}^{-9} \hat{\omega}_{\tau} \hat{\chi} \hat{\Pi}_{dt-\tau}, \quad (11)$$

with the overall effect including news-based beliefs,

$$\hat{\xi} \sum_{\tau=0}^{-9} \hat{\omega}_{\tau} (\hat{\chi} \hat{\Pi}_{dt-\tau} + (1 - \hat{\chi}) \hat{\pi}_{hdt-\tau}). \quad (12)$$

We illustrate in the case of Tunisia. Figure 5 Panel A shows the effect that we would

³⁶Their 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 which destination countries are stereotyped as dangerous if they are covered by bad news without any other background news coverage.

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, a year that saw several terrorist attacks, it alone accounts for a spending decline of close to 0.2 log points. This illustrates the media multiplier at work in Tunisia during that year.

(Figure 5)

Figure 5 Panel B explores how well the model-based approach fits the data. We present the predicted effect, averaged over all potential tourists' origin countries, for Tunisia based on (9) and compare it to the average of the residuals in the spending data after having conditioned on 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.

Our estimates indicate 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 38 billion USD due to violence with close to 12 billion USD being due to the independent effect of negative news reporting.³⁷ Tunisia is 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.

³⁷For calculations see the Appendix C.

4.7 The Media Multiplier: Out-of-Sample Findings

In our main results section 4.2, we speculated about the size of the media multiplier in situations with continuous reporting of mortality risks and compared the economic effect to lockdowns implemented during the COVID-19 pandemic. But a much more direct out-of-sample exercise is to consider the impact that international news coverage could have on international travel. This is especially true for developing countries where the drowning out of bad news with background news is less likely to be important given the relative sparsity of reporting on events that happen there.

We investigate this empirically using a news corpus of close to 5 million international news articles that cover over 180 countries but are not dyadic, i.e. origins and destinations are not matched. For this, we use our access to the topic model (LDA) output at the article level from the webpage <https://conflictforecast.org/> to classify all articles in the underlying database into bad news, B_{dt} , and background news, $N_{dt} - B_{dt}$ (see ?).³⁸ We focus on countries that report more than 100 fatalities and have sufficient news coverage, defined as having had at least 1000 articles over the sample period covered in (?). As a result, we get a monthly time series of news reporting for 67 destination countries, d . We then use the UCDP armed conflict fatality data as a measure of violence. This allows us to derive our two series of beliefs $\{\Pi_{dt}, \pi_{dt}\}$ for each country.

As in our core analysis, we use the estimated beliefs to calibrate the average media-multiplier, $multiplier_{dt}$ in equation (10), during the period 1989 to 2022. The result is shown in Figure 6; we see notable variation across countries. In particular, news-based beliefs for Afghanistan, Pakistan and Somalia are extremely negative throughout the period so that the multiplier reaches close to its maximum of 3.33. But we also get very high values for other countries. For many countries like Cambodia, India, Algeria,

³⁸The coding method that we use is described in the Appendix.

Turkey, Iraq, Sri Lanka, Nepal, the Philipines, Kenia, Angola and Nigeria we get an average multiplier close to 2. The United States and Israel are towards the middle of the distribution - here our model suggests that the media amplifies the effect of violent events by around 30 percent on average. The contrast between India with a lot of bad media coverage and a multiplier over 2.5 and China with almost entirely neutral coverage and a multiplier of 1 is particularly striking.

Entire regions are consistently covered negatively in the media data. The average multiplier in countries with violence in South-Eastern Asia is 1.9, in Southern Asia it is 2.3 and the average multiplier for Eastern and Middle Africa is 1.6. If our main estimates on card activity were to be applied to violent countries in these regions this would suggest large economic losses driven by a significant reduction in foreign visits. With our estimates of the overall beliefs in section 4.2, a media multiplier of 2 is associated with a reduction in the number of tourists – proxied by the active card measure – of over 50 percent!³⁹

Although interesting, the findings in Figure 6 should be treated with some caution. The automated labelling we conduct with the help of the topic model is more likely to capture military violence and not just violence that may put tourists in harms way. As tourist harm is directed at visitors, we expect a larger impact on negative beliefs. Moreover, violence against tourists is harder to predict whereas military conflict might be covered internationally and may well be restricted to specific locations within a country. Tapsoba (2023) suggests that higher uncertainty of violence magnifies the economic impact so that countries with highly localized conflicts are less affected in their visits than Figure 6 suggests. In particular, this will affect countries like Turkey and India that have high levels of militarized violence in some localities.

These caveats notwithstanding, these out-of-sample calibrations support the idea

³⁹The coefficient in Table 2, Column 8 is -0.888 . The effect on the number of cards is therefore $(e^{-0.888} - 1)(0.3 + 2 \times 0.3) = -0.53$.

that international news reporting on violence could be shaping international travel patterns in economically meaningful ways. Taken at face value, the results suggest that some countries are effectively cut-off from tourism flows if they experience violence as the risks make international headlines without much offsetting background news. And it illustrates one of the key risks associated with building up tourism as a central plank of a development strategy.

5 Concluding Comments

This paper contributes to our understanding of the power of the media in influencing economic decisions by introducing the concept of a media multiplier - which captures how agents' responses to news reporting can amplify the effect of an event. We have offered a theory-based measurement approach based on a specific model of belief formation. The 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 seventy 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.

The work in this paper focused on quantifying the media multiplier for a specific type of event – episodes of violence. Alternative event-types may give rise to very different forms of media multipliers and it would be important to explore this case-by-case. Given the fact that information markets have a high degree of national media dominance and most of the challenges humanity face are transnational – a natural extension of this work would consist of developing dyadic data more generally. As most existing research linking media and political economy focus on a small sample of individual countries, we lack an answer to many important research questions. This is particularly

problematic when studying issues that ultimately are transcending national borders – such as events related to climate change.

It would be useful in the future to use lab experiments to explore in greater detail how news media coverage plays out in influencing behavior more broadly. It would also be interesting to study diversion of tourism spending documented here to other locations. 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 suggests that this could be important in affecting tourists' experiences.

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 – potentially seeding anxiety and panic (Fetzer et al., 2020). Besley (2020) argues 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 censored media when forming beliefs.

The media multiplier is also relevant to some ongoing debates about deglobalization. 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, creates 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 a wide range of contexts.⁴⁰

⁴⁰This is envisioned in the “Media, Economics, and Geopolitics (MEGEO)” project funded by the European Research Council (see <https://www.trfetzner.com/megeo>).

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 (2022). Terrorism, Media Coverage, and Education: Evidence from al-Shabaab Attacks in Kenya. *Journal of the European Economic Association* 21(2), 727–763.
- Amodio, F. and M. Di Maio (2017). Making Do With What You Have: Conflict, Input Misallocation and Firm Performance. *The Economic Journal* 128(615), 2559–2612.
- 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.
- Azaredo 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.
- 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.
- 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. Fourinaies (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. (2020). Can Workfare Programs Moderate Conflict? Evidence from India. *Journal of the European Economic Association* 18(6), 3337–3375.
- Fetzer, T., B. Feld, P. J. Lambert, and F. Didier (2023). On the unintended consequences of trade blockades. *mimeo*.
- Fetzer, T., L. Hensel, J. Hermle, and C. Roth (2020). Coronavirus Perceptions and Economic Anxiety. *The Review of Economics and Statistics*, 1–36.
- Fetzer, T. and C. Schwarz (2021). Tariffs and Politics: Evidence from Trump’s Trade Wars. *The Economic Journal* 131(636), 1717–1741.
- Fetzer, T., O. Vanden Eynde, P. C. Souza, and A. Wright (2020). Losing on the home front: Evidence from NATO in Afghanistan. *Mimeo*.

- 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.
- Lewandowsky, S., M. Jetter, and U. K. Ecker (2020). Using the president’s tweets to understand political diversion in the age of social media. *Nature Communications* 11(1), 1–12.
- 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.
- Müller, K. and C. Schwarz (2021). Fanning the Flames of Hate: Social Media and Hate Crime. *Journal of the European Economic Association* 19(4), 2131–2167.
- Neumayer, E. (2004). The Impact of Political Violence on Tourism. *Journal of Conflict Resolution* 48(2), 259–281.

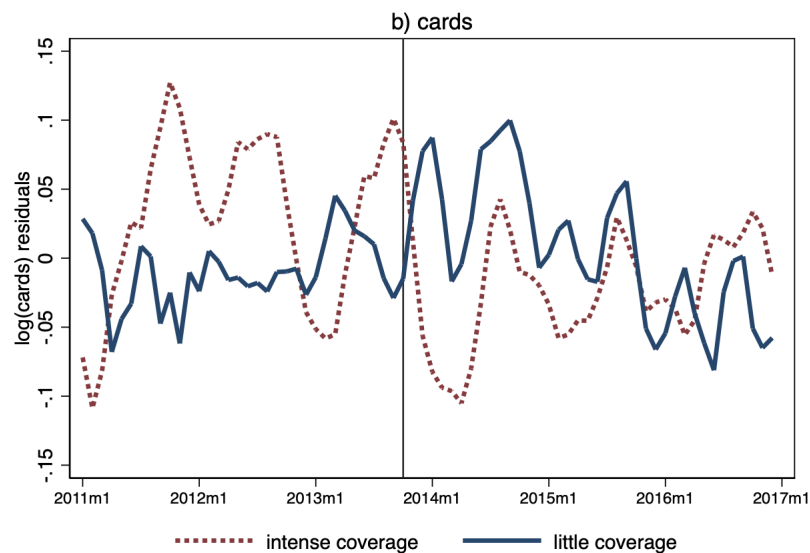
- Pedregosa, F., V. Michel, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, J. Vanderplas, D. Cournapeau, G. Varoquaux, A. Gramfort, B. Thirion, O. Grisel, V. Dubourg, A. Passos, M. Brucher, M. Perrot, and E. Duchesnay (2011). Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 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.
- Tapsoba, A. (2023). The cost of fear: Impact of violence risk on child health during conflict. *Journal of Development Economics* 160(August 2022), 102975.
- 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: Card activity in Tunisia on dyads with more/less news coverage of violence

Panel A:



Panel B:




52

Notes: Figure shows average log(spending) and log(cards) residuals cleaned for dyad, destination/time and origin/time fixed effects. Two lines in each Panel each show average residuals for half the sample (28 and 27 dyads respectively). The dashed line shows the average for dyads with intense news coverage of violence in Tunisia. The solid line shows the average for dyads with less intense news coverage of violence in Tunisia. The vertical line indicates the beginning of intensified violence in Tunisia with an attack against police and a terror attack on a resort in October 2013.

Figure 2: Lead and lagged effects of violence targeted against tourists and fatal violence on tourism spending

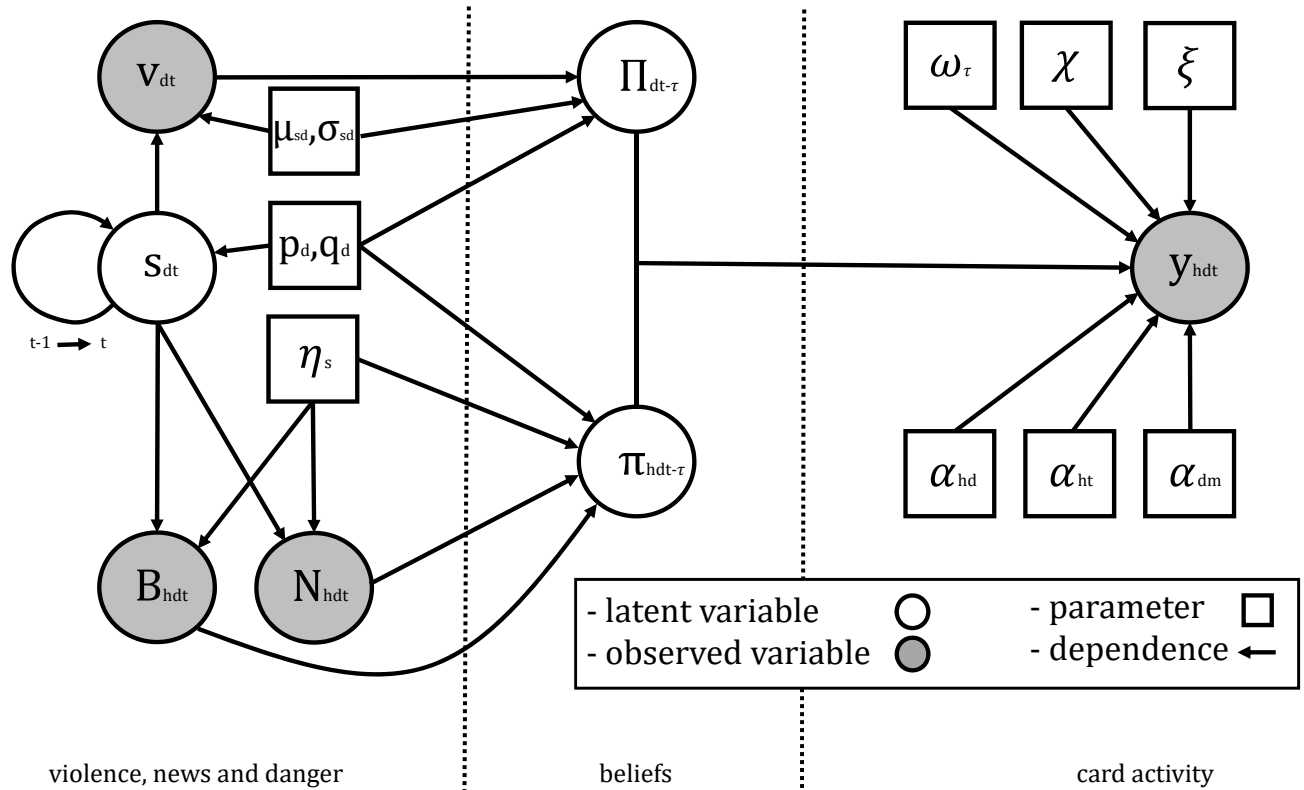
Panel A: News on tourist being targeted

Panel B: News on any fatal violence



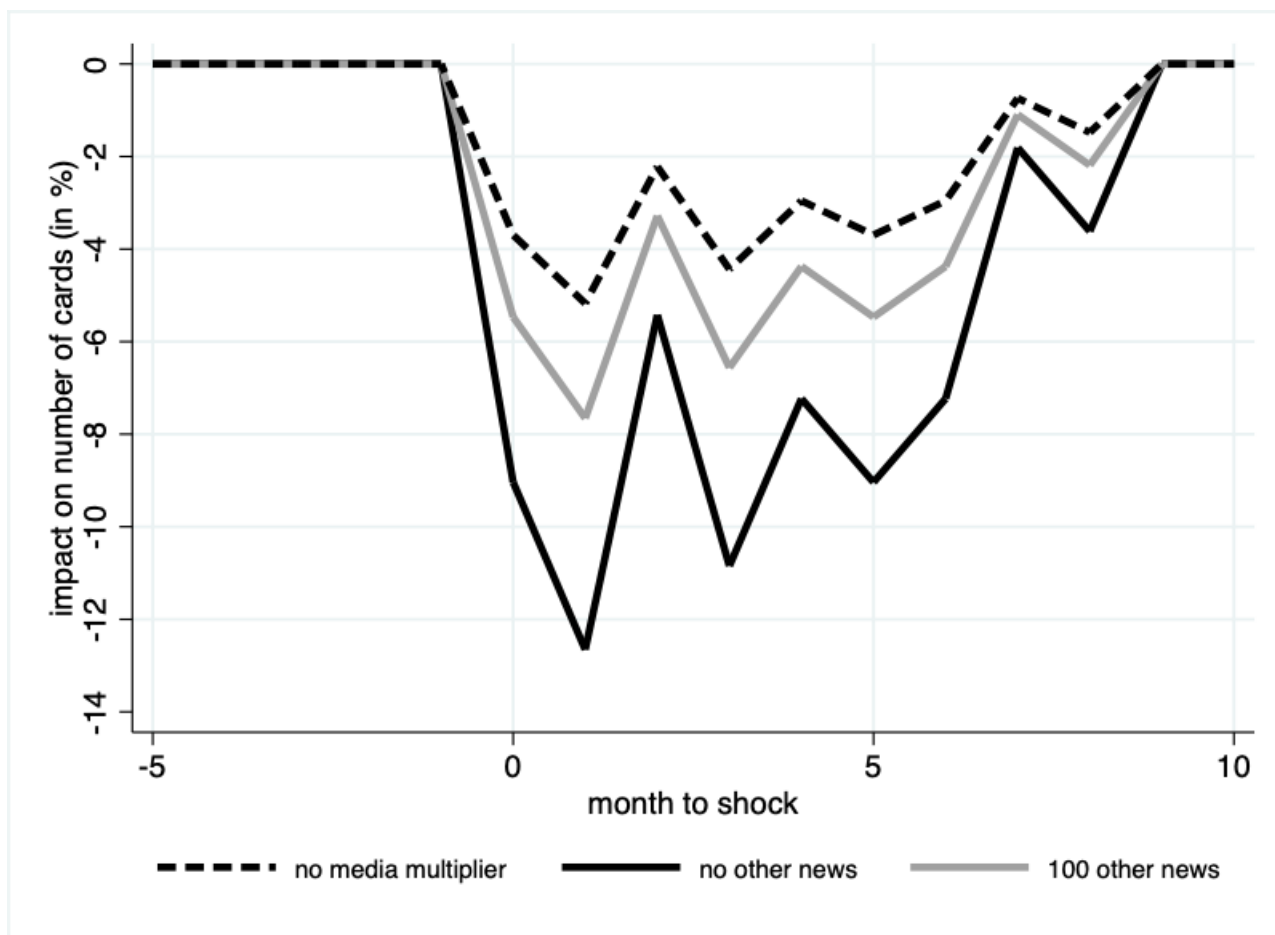
Notes: Figure plots the results from estimating a linear regression controlling for dyad fixed effect, issuing-country by time fixed effect and destination by month seasonality. The coefficients that are reported are the estimates on different leads and lags of the violence reporting measure on the log-value 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 the 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.

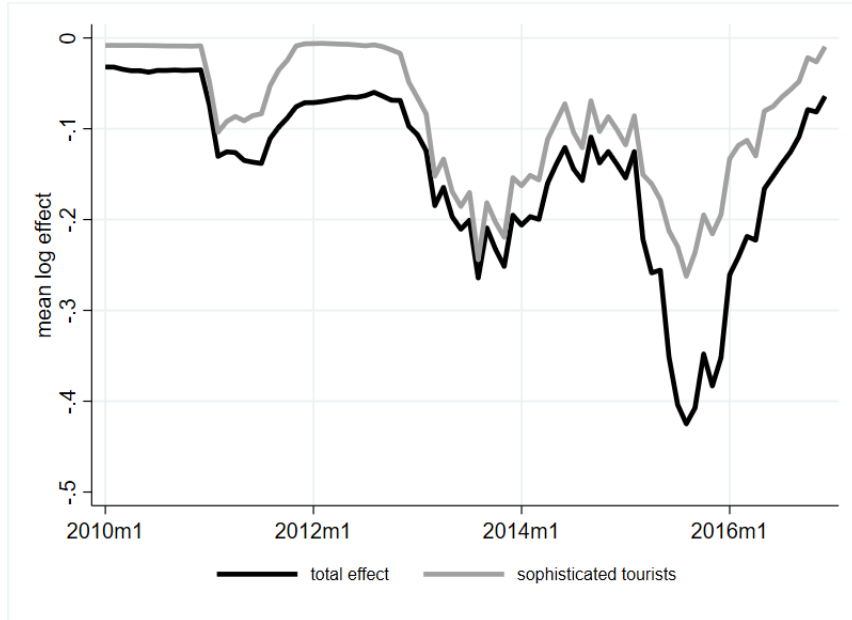
Figure 4: Impact of the media multiplier at different levels of background news




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 is 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: Visual evaluation of model fit in the case of Tunisia

Panel A: Relevance of news-based beliefs




Panel B: Model fit versus data



Notes: Figure plots the overall model-based impact of violent news shocks on aggregate card spending for Tunisia over time. Panel A documents the relevance of the news-based beliefs in driving the total effect by contrasting the full model with the optimal mixture of both types of beliefs against a model fit in which there are only agents that hold beliefs based on the observed violence data, i.e. “event-based beliefs”. This difference drives the media multiplier. Panel B illustrates the model fit vis-a-vis the residualized data.

Figure 6: The media multiplier internationally



Notes: The figure plots shows the average $multiplier_{dt} = 1 + \frac{1-\lambda}{\lambda} \left[\sum_{\tau=0}^{-9} \omega_{\tau} (\pi_{dt-\tau}) \right]$ for destination countries d in the period 1989 to 2022. The events data here is the best estimate of fatalities from UCDP. The news-based beliefs are constructed from a corpus of 6 million international news articles that we categorize into good and bad news with a topic model. We code an article as bad news, B_{dt} , if topics covering armed conflict, police abuses or the military dominate the article.

Table 1: Reduced form results: impact of news reporting on tourism activity

	(1)	(2)	(3)	(4)	(5)	(6)
	log(Spending)			log(Number of cards)		
<i>Panel A: News on tourist being targeted</i>						
News on tourists targeted (share of all articles)	-0.542*** (0.102)	-0.506*** (0.109)	-0.179* (0.099)	-0.615*** (0.076)	-0.593*** (0.082)	-0.166** (0.074)
Observations	23859	23859	23859	23869	23869	23869
R2	.966	.967	.972	.97	.971	.979
<i>Panel B: News on any fatal violence</i>						
News on violence with fatalities (share of all articles)	-0.295*** (0.041)	-0.197*** (0.043)	-0.091** (0.037)	-0.281*** (0.039)	-0.208*** (0.040)	-0.048 (0.032)
Observations	23859	23859	23859	23869	23869	23869
R2	.966	.966	.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 regressions capturing the 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 with fatalities. 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)	-0.208*** (0.019)			-0.190*** (0.019)		-0.203*** (0.020)			-0.186*** (0.019)	
probability of danger (tourist news-based)		-0.360*** (0.050)		-0.268*** (0.049)			-0.344*** (0.047)		-0.255*** (0.046)	
probability of danger (fatal news-based)			-0.230*** (0.041)					-0.210*** (0.046)		
weighted probability of danger					-0.834*** (0.068)					-0.888*** (0.073)
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 regressions explaining variation in tourist activity measured either as the log value of card spending in columns (1)-(5) or 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 and 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)). Columns (5) and (10) explore the weighted average of the two. Please refer to section 4.1 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)	-0.235*** (0.072)	-0.251*** (0.096)	-0.162** (0.068)	-0.204*** (0.068)	-0.171** (0.078)	-0.138** (0.065)
probability of danger (tourist news-based) * free press	-0.237** (0.095)	-0.221* (0.115)	-0.204** (0.094)	-0.267*** (0.089)	-0.300*** (0.097)	-0.220** (0.089)
probability of danger (based on violence data)			-0.187*** (0.027)			-0.169*** (0.028)
probability of danger (based on event data) * free press			-0.003 (0.035)			-0.033 (0.036)
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) or 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_{ldt} . Please refer to section 4.1 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 country has a press freedom score above average. The results highlight that the impact of violent news shocks (as opposed to violent events) on 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 to E are Supplementary Material

November 11, 2023

This Appendix is subdivided into five 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. We then describe the speculative out-of-sample exercise we conduct to estimate the media multiplier across continents. Lastly, section **E** 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:

$$y_{hdt} = \alpha_{hd} + \alpha_{ht} + \alpha_{dm(t)} + \tilde{\zeta}v_{hd,t-1} + \varepsilon_{hdt} \quad (13)$$

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 $v_{hd,t-1}$, is lagged by one month to capture the possibility that international travel reacts to past violence. We expect to find that $\xi < 0$ in (13), i.e. violence deters travel activity. We use four sources of data on violent events at the country level in different versions of (13). 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 (13) 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 four measures of violence and the level of tourism activity in a country. The size of the coefficients are in the range of 4% to 7.6% decreases 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 (13) 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 the 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 fixed 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 (13) 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

A.2 Event Regression Evidence for News

Complementary to Figure A4 we run the following specification:

$$n_{hdt} = \alpha_k + \alpha_{hd} + \alpha_t + \beta \times \text{Post}_{k,t} + \gamma_k \times (\text{Post}_{k,t} \times z_k) + \epsilon_{hdt} \quad (14)$$

where we have defined a dummy variable $\text{Post}_{k,t} = 1$ for $t = 0, 1, 2$ for up to two days following a violent event. Estimating equation (14) also allows us to explore whether this average effect is heterogenous across a range of event characteristics z_k : the level of casualties, whether Americans are among the casualties and attacks involving tourists.

We present results from specification (14) in Table A4. In columns (1) through (5), the dependent variable is the share of news articles on a day classified as indicating

¹See Appendix Figure A8.

violence with fatalities. In columns (6) through (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 differences approach relies on there being a common underlying trend in spending between places that experienced violent events and those that 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:

$$y_{khd t} = \alpha_k + \alpha_{hd} + \alpha_{ht} + \sum_{\tau=-6}^6 (\beta_{\tau} \times \text{Time to event month}_{k,t-\tau}) + \epsilon_{khd t} \quad (15)$$

where, as above, $y_{khd t}$ 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 (15) permits us to trace out the patterns of aggregate spending around an event month. The results are depicted in Figure A8 for both the log of spending and the log of active card accounts. In both cases, there is no evidence of any anticipation of the event. Moreover, the observed pattern suggests a sharp contraction in card spending and the number of active cards, with a one month lag as in the 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

$$y_{hdt} = \alpha_{hd} + \alpha_{ht} + \alpha_{dm} + \zeta \sum_{\tau=0}^9 \omega_{\tau} (\chi \Pi_{dt-\tau} + (1 - \chi) \pi_{hdt-\tau}) + \epsilon_{hct} \quad (16)$$

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

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

which we then use as a regressor in equation (16). 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. We pick the parameter values that yield the highest within R-squared and then average the parameters we get from the regression of spending and number of cards to generate a general model.

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 take a similar shape as in the reduced form. In particular, we get the weight sequence:

$$0.125, 0.175, ,0.075, 0.15, 0.1, 0.125, 0.1, 0.025, 0.05, 0.075.$$

For χ , we get a value of 0.3 which implies that only 30 percent of the agents in our model are estimated to have event-based beliefs. However, event-based tourists will nonetheless drive most spending movements as the shifts in their beliefs are a lot more persistent.

C Calculations of Total Loss

Assume that we have 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:

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

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: 0.043

(0.027) billion USD in Tunisia, 0.065 (0.041) billion USD in Israel, 0.159 (0.130) billion USD in Egypt, 0.264 (0.168) billion USD in Turkey. Numbers in brackets indicate the losses from event-based tourists alone. This means a total loss of 38.23 billion USD; 26.35 billion USD of this is driven by events and 11.88 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 $3.48/12 - e^{((\ln(3.48/12))+(-0.165))}$, where -0.165 is the average treatment on dyads into Tunisia in the period 2011-2016.

D Out-of-Sample Method


In order to construct a panel of news-driven beliefs for all countries with sufficient news coverage (more than 1000 articles written in the period 1989 to 2021), we use the topic model outputs at the article level to distinguish bad from background news. We proceed as follows. We first define an article as containing bad news if the largest estimated topic share of the article is either topic 5, a topic we label as *military conflict*, or topic 10, which is labeled as *military technology*.² The labeling is somewhat arbitrary but it leads to reasonable timelines as shown for the case of Tunisia in Figure A12. We use UCDP fatalities (the best estimate) per capita as our main outcome measure to define the event-based beliefs. In order to do this, we run the Markov chain model for each of the 168 countries with more than 100 fatalities and more than 1000 news articles.

Our derivations of the two sets of beliefs otherwise completely follow the main results section. The Markov chain results then form the foundation for our analysis of the news-based beliefs as well. The only difference is that we do not have several dyads but only one time series per country. Note that our model of event-based beliefs fits to each country time series by identifying escalations in the violence time series.

²We vary the topic model over time and this means the topics on the website and codebook can deviate from these topics.

This means something very different across countries. News-based beliefs are then calibrated using these models of violence. One way to check the plausibility of the model is therefore to check the binomial distribution parameters imposed by the event-based model on the the news-based model. We find that even in a sample of over 63 destination countries with very different violence time-series, $0.163 = \hat{\eta}_1 > \hat{\eta}_0 = 0.061$, so that bad news shares close to 16% will start to signal violence risk.


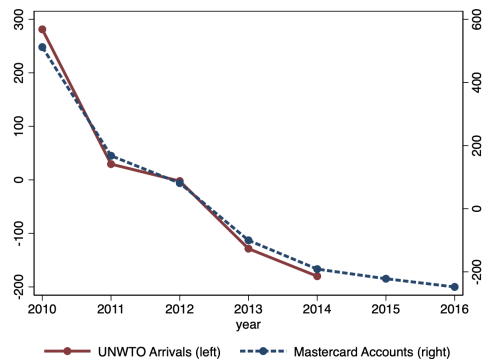
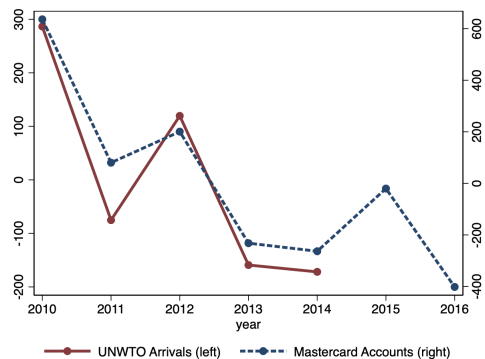
Figure A1: Map of countries included in our estimation samples across the exercises



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



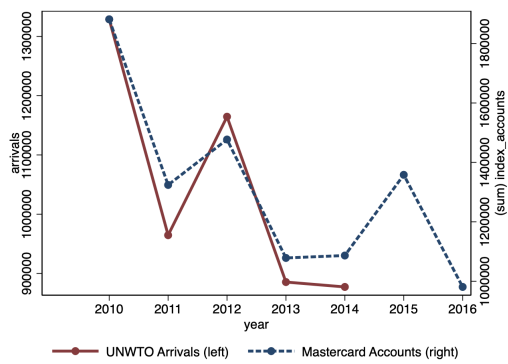
French

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

Panel A: Egypt

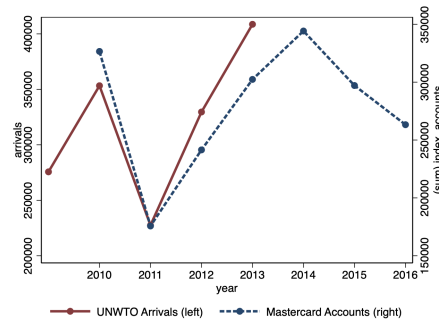

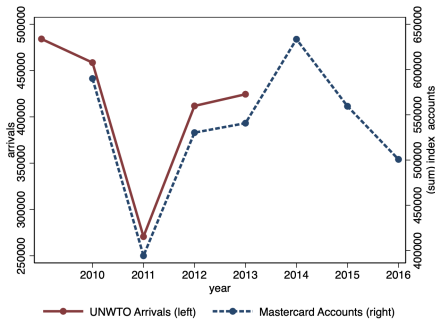


German

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.

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



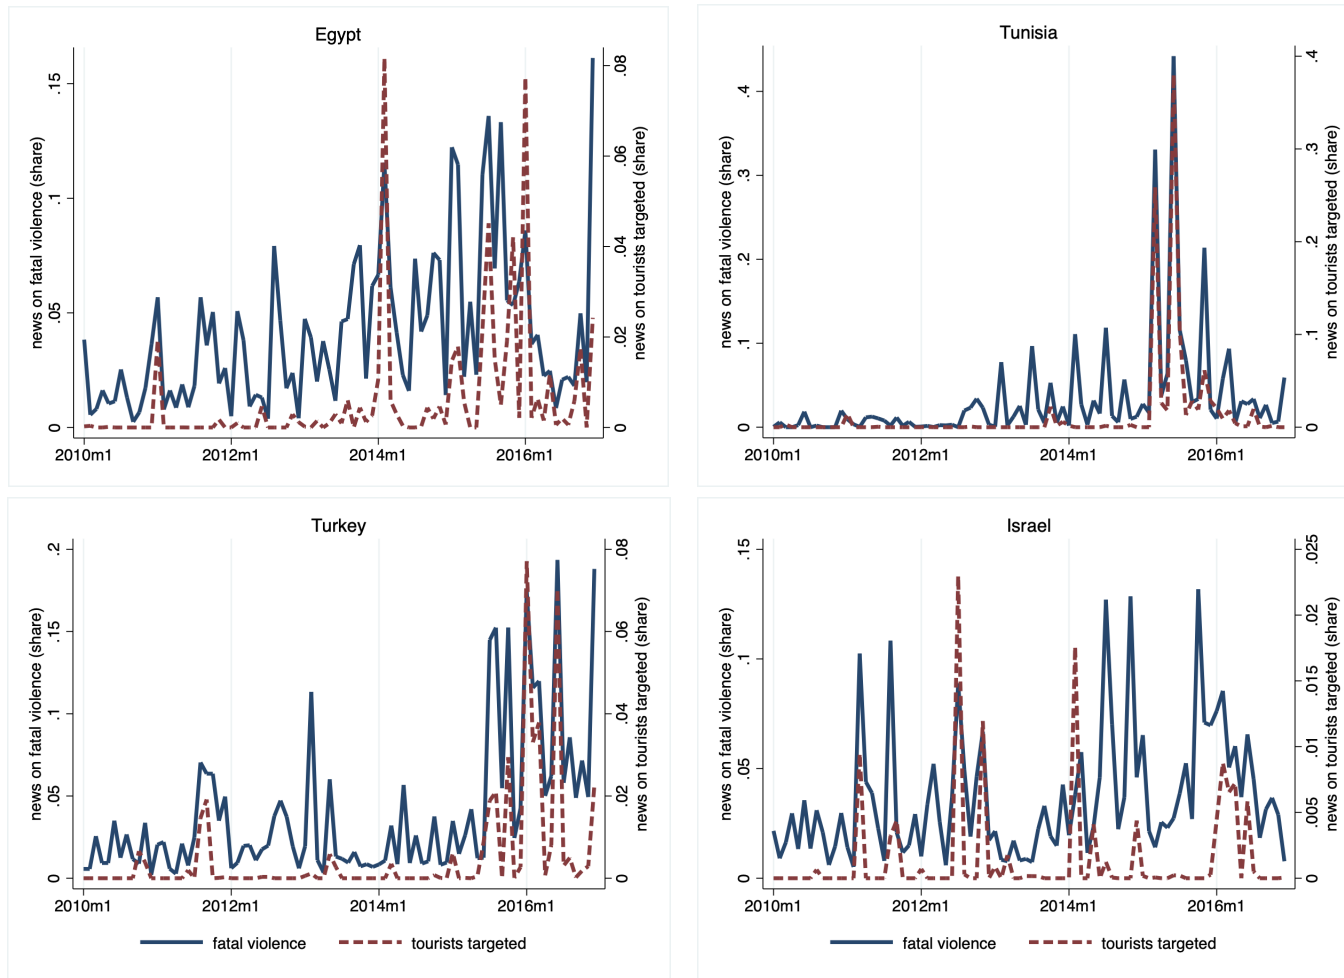
Notes: Figure plots point estimates from a regression that absorbs event, reporting dyad and day fixed effects. The dependent variable in Panel A measures the share of articles on a day and dyad that are classified as reporting on violence that involved fatalities. The dependent variable in Panel B measuring 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 the four main destination countries in our sample. The figure provides for each of destinations the average share of news articles across the tourist-origin countries that are classified as reporting on violent events with fatalities (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 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 4. 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)



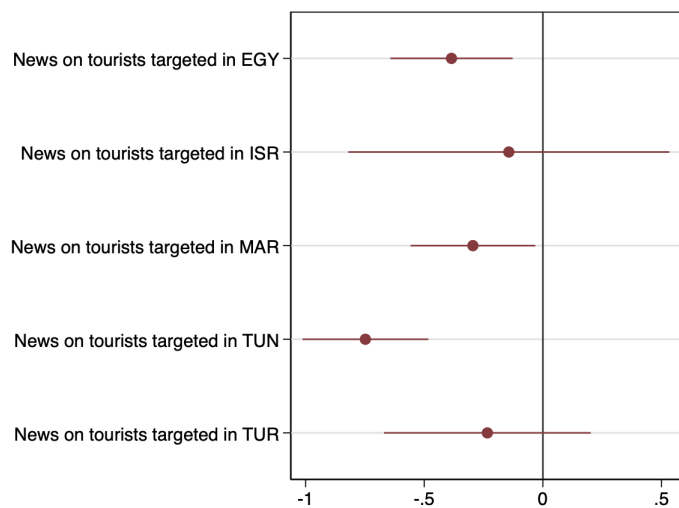
Panel B: log(Number of Cards)




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 kernel density of the posterior news belief for safe and dangerous dyad/months conditional on event beliefs indicating safety or danger.

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.

Figure A12: News-based beliefs in Tunisia based on topic model



Notes: Figure shows the news-based beliefs for Tunisia based on an international news corpus from <https://conflictforecast.org/> that covers over 180 destination countries but does not distinguish news origins. Articles in this news corpus are classified using a topic model. We count an article as bad news, B_{dt} if the topic share on *armed conflict* or *military technology* is larger than any other of the 15 topics used on the webpage. Event-based beliefs are calculated from UCDP fatalities. The Figure then shows the outcome of running our model with its parameter estimates from the main results section but using these new ways of capturing news and event-based beliefs.

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)
<i>Panel A:</i>			
arrivals	0.700*** (0.197)	1.189*** (0.262)	182.604*** (60.961)
Dyads	294	294	294
Observations	1258	1258	1258
<i>Panel B:</i>			
arrivals	0.535*** (0.097)	1.048*** (0.190)	126.985*** (39.719)
L.arrivals	-0.016 (0.067)	-0.021 (0.108)	0.271 (20.866)
Dyads	290	290	290
Observations	974	974	974
<i>Panel C:</i>			
F.arrivals	0.344** (0.149)	0.619** (0.256)	83.950*** (31.125)
arrivals	0.755*** (0.206)	1.183*** (0.299)	204.813*** (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*** (0.010)					-0.034*** (0.008)				
GTD Events		-0.076*** (0.017)					-0.076*** (0.016)			
ICEWS armed violence events			-0.068*** (0.020)					-0.054*** (0.019)		
GDELT armed violence events				-0.065*** (0.016)					-0.047*** (0.015)	
Armed violence component 1					-0.067*** (0.017)					-0.053*** (0.017)
Armed violence component 2					-0.040*** (0.014)					-0.034** (0.015)
Armed violence component 3					0.017 (0.015)					0.010 (0.015)
Armed violence component 4					-0.031** (0.013)					-0.039*** (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 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*** (0.000)	0.002*** (0.000)	0.002*** (0.000)	0.011*** (0.002)	0.003* (0.002)	-0.000 (0.000)	0.000*** (0.000)	0.000** (0.000)	0.001** (0.000)	-0.005*** (0.001)
post × Casualties	0.000*** (0.000)				0.002*** (0.000)	0.000*** (0.000)				0.002*** (0.000)
post × US Casualties		0.014*** (0.002)			0.016** (0.007)		0.001** (0.001)			0.001** (0.001)
post × Suicide attack			0.017*** (0.002)		0.020*** (0.006)			0.002*** (0.000)		-0.002 (0.004)
post × Tourist targeted				0.014*** (0.003)	0.008*** (0.003)				0.022*** (0.003)	0.016*** (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) measures 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 two-way at the 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)			
<i>Panel A: News on tourist being targeted</i>				
News on tourists targeted (count of articles)	-0.009*** (0.003)	-0.008*** (0.003)		
News on tourists targeted (share of all articles)	-0.492*** (0.106)	-0.465*** (0.112)		
News on tourists targeted (share of all articles) - first quartile			0.050 (0.060)	0.113* (0.058)
News on tourists targeted (share of all articles) - second quartile			-0.115*** (0.034)	-0.077** (0.035)
News on tourists targeted (share of all articles) - third quartile			-0.133*** (0.039)	-0.096** (0.038)
News on tourists targeted (share of all articles) - fourth quartile			-0.259*** (0.070)	-0.245*** (0.072)
Observations	23859	23859	23859	23859
R2	.966	.967	.966	.967
<i>Panel B: News on any fatal violence</i>				
News on violence with fatalities (count of articles)	-0.003 (0.002)	0.001 (0.002)		
News on violence with fatalities (share of all articles)	-0.270*** (0.043)	-0.205*** (0.044)		
News on violence with fatalities (share of all articles) - first quartile			-0.069*** (0.026)	-0.033 (0.025)
News on violence with fatalities (share of all articles) - second quartile			-0.058*** (0.021)	-0.014 (0.021)
News on violence with fatalities (share of all articles) - third quartile			-0.072*** (0.025)	-0.018 (0.024)
News on violence with fatalities (share of all articles) - fourth quartile			-0.096*** (0.020)	-0.058*** (0.020)
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 regressions capturing the reduced form effect of dyadic (tourist-origin by destination) specific news coverage on the dyadic log values of card spending. 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 with fatalities. 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 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		
<i>Panel A: first stage with reporting spillovers</i>			
Any tourist killed	0.171*** (0.043)	0.180*** (0.044)	
Contiguous country x Any Casualties		0.126*** (0.032)	0.122*** (0.035)
Same region x Any Casualties		0.038** (0.019)	0.040* (0.022)
Common language x Any Casualties		0.042*** (0.014)	0.040*** (0.015)
Constant	0.005*** (0.000)	0.004*** (0.000)	0.004*** (0.000)
R2	0.258	0.294	0.294
<i>Panel B: Second stage: tourism spending</i>			
News on tourists targeted (share of all articles)	-0.850 (0.609)	-0.897*** (0.334)	-0.701 (0.435)
R2	0.026	0.025	0.027
Weak IV	14.015	15.138	12.510
<i>Panel C: Second stage: number of active cards</i>			
News on tourists targeted (share of all articles)	-1.248* (0.692)	-1.298*** (0.347)	-1.132*** (0.422)
R2	.0244	.0235	.0268
Weak IV	14	15.1	12.5
Observations	23869	23869	21537
Dyad FE	YES	YES	YES
Origin/Time FE	YES	YES	YES
Dest./Month FE	YES	YES	YES
Event controls	YES	YES	YES
Excluding directly treated dyads	NO	NO	YES

Notes: Table 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 on aggregate spending: dropping each country in turn

	(1)	(2)	(3)	(4)	(5)	(6)
	Dropping a country in turn					
	All	EGY	TUN	TUR	MAR	ISR
<i>Panel A:</i>						
Armed violence component 1	-0.077*** (0.008)	-0.023** (0.011)	-0.090*** (0.008)	-0.112*** (0.009)	-0.063*** (0.008)	-0.040** (0.018)
Armed violence component 2	-0.036*** (0.007)	0.057*** (0.011)	-0.036*** (0.007)	-0.052*** (0.007)	-0.032*** (0.007)	-0.080*** (0.011)
Observations	42254	33315	35596	32600	34227	32834
R2	.947	.955	.951	.945	.952	.953
<i>Panel B:</i>						
GTD Events	-0.076*** (0.008)	-0.030*** (0.010)	-0.089*** (0.008)	-0.110*** (0.010)	-0.063*** (0.008)	-0.085*** (0.012)
Observations	42254	33315	35596	32600	34227	32834
R2	.947	.954	.95	.945	.952	.953
<i>Panel C:</i>						
ICEWS armed violence events	-0.068*** (0.008)	-0.025** (0.010)	-0.081*** (0.008)	-0.118*** (0.010)	-0.056*** (0.008)	-0.070*** (0.010)
Observations	42254	33315	35596	32600	34227	32834
R2	.947	.954	.95	.945	.952	.953
<i>Panel D:</i>						
UCDP Events	-0.040*** (0.005)	-0.034*** (0.005)	-0.043*** (0.005)	-0.042*** (0.005)	-0.039*** (0.005)	-0.045*** (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: 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 a one standard deviation increase in violence regardless of the violence measure. Components are coming from a principal component analysis of all different violence data sub-categories. Robust standard errors clustered at destination by time level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table A8: Relationship between reporting and aggregate spending

	(1)	(2)	(3)	(4)
			log(Spending)	
News on tourists targeted (count of articles)	-0.008*** (0.003)	-0.002 (0.002)		
News on tourists targeted (share of all articles)	-0.465*** (0.112)	-0.172* (0.101)		
Armed violence component 1	-0.084*** (0.009)		-0.083*** (0.009)	
Armed violence component 2	-0.026*** (0.006)		-0.028*** (0.006)	
Armed violence component 3	0.017*** (0.004)		0.013*** (0.004)	
Armed violence component 4	-0.015* (0.009)		-0.016* (0.009)	
News on violence with fatalities (count of articles)			0.001 (0.002)	0.002 (0.002)
News on violence with fatalities (share of all articles)			-0.205*** (0.044)	-0.105*** (0.038)
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)
	Additional controls				Alternative cutoffs & not relying on hand coding			
					$c = 0.5$		$c = 0.90$	
<i>Panel A: News on tourist being targeted</i>								
News on tourists targeted (share of all articles)	-0.506*** (0.109)	-0.547*** (0.107)	-0.179* (0.099)	-0.126 (0.087)	-0.920*** (0.215)	-0.365** (0.184)	-0.920*** (0.215)	-0.365** (0.184)
Observations	23859	23859	23859	23859	23859	23859	23859	23859
R2	.967	.967	.972	.976	.967	.972	.967	.972
<i>Panel B: News on any fatal violence</i>								
News on violence with fatalities (share of all articles)	-0.197*** (0.043)	-0.227*** (0.044)	-0.091** (0.037)	-0.083** (0.034)	-0.207*** (0.051)	-0.073 (0.045)	-0.209*** (0.067)	-0.065 (0.059)
Observations	23859	23859	23859	23859	23859	23859	23859	23859
R2	.966	.967	.972	.976	.966	.972	.966	.972
Dyad FE	YES	YES	YES	YES	YES	YES	YES	YES
Origin by Time FE	YES	YES	YES	YES	YES	YES	YES	YES
Destination by Month FE	YES	YES	NO	NO	YES	NO	YES	NO
Destination by Time FE	NO	NO	YES	YES	NO	YES	NO	YES
Dyad Linear Trend	NO	NO	NO	YES	NO	NO	NO	NO
Exchange rate	NO	YES	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) - (8), we rely on the news measures constructed not using the secondary hand coding procedure we described in section 2. The columns explore alternative classification cutoffs to highlight that results are robust. All explanatory variables are divided by their standard deviation so that the coefficients can be interpreted as the response of aggregate spending to a 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	0.45	0.296	0.49	0.423	0.26
mean violence in safety	0.31	0.265	0.30	0.303	0.26
difference(danger-safety)	0.15	0.031	0.19	0.120	0.00
persistence of danger	0.96	0.586	0.93	0.812	0.85
persistence of safety	0.99	0.855	0.97	0.929	0.83

Notes: Table reports estimates of the parameters for the Markov chain switching model. For definitions see the main text in section 4.

Table A11: Effect of Markov chain fitted probability of dangerous state on spending across destination countries

probability of danger (news-based) in Egypt	-1.018*** (0.102)
probability of danger (news-based) in Israel	-0.596*** (0.195)
probability of danger (news-based) in Morocco	-0.349 (0.251)
probability of danger (news-based) in Tunisia	-1.161*** (0.212)
probability of danger (news-based) in Turkey	-0.613*** (0.101)
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 4.

Table A12: Heterogeneous effects: role of geographic distance and other similarities

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	log(Spending)				log(Number of cards)			
probability of danger (tourist news-based)	-0.284*** (0.104)	-0.411*** (0.048)	-0.355*** (0.070)	-0.318*** (0.080)	-0.300*** (0.098)	-0.378*** (0.050)	-0.372*** (0.060)	-0.343*** (0.072)
probability of danger (tourist news-based) * christian share	-0.001 (0.002)				-0.001 (0.001)			
probability of danger (tourist news-based) * muslim share		0.005 (0.003)				0.003 (0.003)		
probability of danger (tourist news-based) * far away			-0.009 (0.102)				0.059 (0.099)	
probability of danger (tourist news-based) * distance				-0.154 (0.193)				-0.005 (0.198)
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)-(4) or the log number of active cards in columns (5) - (8) 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 4.1 for how we leverage the violence and news reporting data to estimate these. The table explores whether the impact of violence shocks is heterogeneous 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.

Table A13: Relationship between fitted beliefs and leads of the violent event data (principal component)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	One month lead			Two months lead			Three months lead		
probability of danger (based on violence data)	0.077*** (0.006)		0.084*** (0.007)	0.056*** (0.006)		0.063*** (0.007)	0.055*** (0.006)		0.062*** (0.007)
probability of danger (based on news data)		0.024** (0.010)	-0.024** (0.010)		0.013 (0.009)	-0.022** (0.010)		0.009 (0.009)	-0.026*** (0.009)
Observations	27075	27075	27075	26790	26790	26790	26505	26505	26505
R2	.476	.458	.477	.468	.457	.468	.466	.457	.467
Dyad FE	YES	YES	YES	YES	YES	YES	YES	YES	YES
Origin/Time FE	YES	YES	YES	YES	YES	YES	YES	YES	YES
Dest./Month FE	YES	YES	YES	YES	YES	YES	YES	YES	YES

Notes: Robust standard errors clustered at destination/month level are in parentheses. *** p<0.01, ** p<0.05, * p<0.1. The left hand side is are the first two principal components of the violence event data from GTD, GDETLT and UCDP.

Table A14: Relationship between fitted value from markov chain model, the violence data and card activity

	(1)	(2)	(3)	(4)	(5)	(6)
	log(Spend)			log(Cards)		
Probability	-0.199*** (0.019)		-0.165*** (0.020)	-0.182*** (0.019)		-0.170*** (0.019)
PCA Violence		-0.704*** (0.077)	-0.275*** (0.085)		-0.543*** (0.080)	-0.102 (0.082)
Observations	23859	23859	23859	23869	23869	23869
R2	.966	.966	.966	.971	.971	.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: Robust standard errors clustered at destination/month level are in parentheses. *** p<0.01, ** p<0.05, * p<0.1. The table shows the fitted value of the Markov chain model as the probability of danger and compares this to the first two principal components of the event-data that the Markov chain is estimated from.

E Description of Machine Learning Method and Validation

This section describes the machine learning method used to identify the presence of text describing a fatal attack or an attack on a tourist in a corpus of close to 500,000 news articles. We deploy a simple supervised machine learning method which exploits 35,000 labelled articles. However, we also produce results using a simple dictionary method as it is the most common NLP method in economics.

It is important to stress that the goal of the method we develop is not to be general in that we design it to detect news text covering tourist harm or fatal attacks generally in all countries. What we claim here is that it will detect the events in the countries and time period of our sample. What this means, specifically, is that we will *overfit* the algorithm to pick up locations of known events like Sousse in Tunisia.

However, for reasons of intellectual curiosity, we also conducted a true out-of-sample test of the algorithm which we discuss at the end of this appendix section.

E.1 Overview and Goal

We want to discover three different types of violent events. The human coders were asked to label articles in three binary categories:

1. *label 0: the article indicates that there was violence*
2. *label 1: the article indicates that there were fatalities as a result of violence*
3. *label 2: the article indicates that tourists were harmed due to a violent event*

The first category (label 0) turned out to be extremely broad - we therefore focused on the latter two (labels 1 and 2) and only used the first category for noise reduction. Note that, while fatal violence and violence against tourists are related, they are not

perfectly overlapping sets. Fatalities capture a lot of violence that is directed at local civilians, insurgents, soldiers and police forces. In Egypt, Israel and Turkey there is substantial militarized violence to drive this category to be relevant. Our definition of harm of tourists includes hijackings and injuries so there are events that are labelled as attacks on tourists without being flagged as fatal attacks.

The difference in the sets becomes clear from a simple cross-tabulation. In our 35,000 labelled articles, we have over 500 articles labelled as covering violence against tourists (label 2) and close to 5000 articles labelled as covering fatal violence (label 1). This means there is a large number of articles covering fatal violence that do not mention tourists. But there are also over 130 articles that mention violence against tourists that do not mention fatal violence.

From this, it should be clear that we have extremely imbalanced classes. In the case of violence against tourists, this is particularly extreme where we have only 1.4 percent of the labelled sample being assigned a 1 and 98.6 percent being assigned a 0. This means that precision will be an important concern, especially as we are relying on the monthly fluctuations in our reporting measures to identify the effect of news shocks.

The machine learning task is further complicated by additional downloads from Factiva which we only conducted once the hand-coding exercise had been completed. This meant that we had training data from articles downloaded from LexisNexis but we did not have human labels for Factiva content.

Our strategy, therefore, integrates human coding and re-coding after machine learning for tourist harm. This is to ensure a reasonable out-of-sample performance in the Factiva and LexisNexis samples, but also to make sure that our human coders spotted the needle in the haystack, i.e. having the highest possible true positive rate while maintaining high levels of precision. We deployed two strategies for tourist harm:

1. We check the human coding using the machine learning output before training the

final model. After the first run of machine learning, we re-visit all negative labels identified as positives by the classifier. Here we found 20 articles that the human coders wrongly classified and that the machine had given high fitted values to. We change these labels and re-estimate the tourist harm model.

2. After our final classification we revisit all positives (above the threshold) as indicated by the classifier and close to 5,000 articles that were ranked highest by the classifier but below the threshold. This brings down error and allows us to check performance once more and compare them to the cross-validation statistics we calculate during training. The exercise suggests that performance measures are realistic.

The final round of human coding gives us the chance to check performance in the Factiva sample which we will turn to below.

E.2 Performance of a Simple Dictionary Method

We first start with a very simple dictionary method to detect fatal violence and violence against tourists. We developed the dictionaries by iterating on terms to include and even tried interaction with various sub-dictionaries following the methodology behind the Economic Policy Uncertainty index (EPU) developed by [Baker et al. \(2016\)](#). The results are two dictionaries reported in the replication code which allow us to rank all the articles in the data according to how likely we think they are to contain reports on fatal violence or violence against tourists.

These rankings can then be easily held against our 35,000 coded observations to check how such a method would perform. We get an AUC of 0.89 for fatal harm and an AUC of 0.90 for tourist harm. This suggests similar performance for our two labels. But [Figure B1](#) shows the precision curve for fatal harm whereas [Figure B2](#) shows the performance for violence against tourists. The much larger imbalance problem

in the articles indicating tourist harm becomes very obvious from this where average precision for fatal harm is 59 percent, but for tourist harm it is only 14 percent. This low performance in average precision for tourist harm is driven by a fall in precision to values of close 20 percent at low recall rates.

Still, it is clear that even such simple dictionary methods can yield some signal from the data for very high cutoffs, i.e. if we require that articles mention tourist harm or fatal violence many times. Changing to a relative frequency model for our dictionary terms did not improve performance.

E.3 Machine Learning Method

The development of our method relies on cross-validation with typically 5 train and test splits. Cross-validation works as follows. The sample is split into 5 random parts. The model is trained on 4 parts and then applied to the data of the 5th part, i.e. we calculate fitted values from the model that was trained on another part of the data to analyze performance. Cross-validation then rotates the parts used for training and testing, but it never uses data to test the model that was also used to train the model. This ensures that the actual performance on the rest of the data is as close as possible to the performance we report below. We rely on cross-validation grid search (GridSearchCV in sklearn) for fixing hyperparameters and on cross-validation fits for calculating the performance of the algorithm overall.

Our classification pipeline has three elements – a count vectorizer, a tf-idf transformer and a classifier. A special feature of our hyperparameter search is that we optimize hyperparameters throughout this pre-processing pipeline and fit three classifiers: a naive Bayes classifier, a random forest and a boosted tree classifier. We use the implementations by sklearn for the first naive Bayes and the random forest classifiers. We use the XGB python package for the latter one.³

³The entire code and pre-processed texts will be made available on GitHub.

The final hyperparameter settings for fatal harm labels and the tourist harm labels are similar with some variation in pre-processing. Details are provided in the replication codes. Strikingly, both tree methods choose very deep trees, *depth20*, which is compensated by high minimum samples of around 10. This means that the trees will fit to complex language patterns to distinguish 0s from 1s. In the pre-processing, we get a mix of unigrams and trigrams being optimal. This means that the ensemble we build can rely on different types of pre-processing and classifiers. Also, upsampling of the positive class is optimal in training – this is a standard finding in the ML literature. Note, this does not affect the performance statistics we produce in any other way than by affecting the trained model as we do not permanently change the balancedness of the classes in the data but only during the training procedure.

The cross-fitted AUCs and the best average precision is reached by the XGB classifier in both cases. The precision-recall curve for this classifier for fatal harm is shown in figure B3 and for violence against tourists it is in figure B4.⁴ The other curves for the other models are all reported in the jupyter notebook on the GitHub repository but they are similar. The XGB classifier reaches an AUC of 0.93 for fatal violence and 0.98 for tourist harm.

Importantly, we are showing performance statistics with the original balance of the data. This implies that the precision-recall curve for fatal violence will converge to over 10 percent for a recall of 1 whereas the curve for tourist harm will converge towards just over 1 percent. It is therefore no surprise that the precision-recall curve for harm to tourists is showing lower performance despite a higher AUC for tourist harm, where average precision is 60 percent. The precision-recall curve indicates fairly high precision scores at low true positive rates. We reach a precision of close to 90 percent for a recall rate of 10 percent and a precision of 60 percent for a recall rate of 50 percent. In words,

⁴For tourist harm we report the performance after manual recoding of coding errors.

we are right 60 percent of the time if we want to filter out half the articles covering violence against tourists.

To reduce error and improve precision, we build an ensemble across our three classifiers. We average across fitted values for our three classifiers and use the resulting average. The final cross-validated performance, after manual recoding for tourist harm, for our three labels (any harm, fatal harm and tourist harm) is shown in ROC curves in Figure B5 to Figure B7 and precision-recall curves in Figures B8 to Figure B10. Again, performance is shown on the original sample with its heavy label imbalance.

Clearly, the ensemble improves the fit of the model. For tourist harm, this yields the cross-fitted ROC curve in B7 and the precision curve in B10. It is clear from this that the ensemble helps to lift precision exactly in the section of observations that are most important for our detection algorithm – for low recall values. However, we still face a real trade-off between precision and recall. Deciding on a cutoff is extremely important in this context. As a default, we follow the common procedure and classify an article as capturing violence if the fitted values are larger than 0.5. In the code we reposit, we show that this leads to a confusion matrix in which we produce 200 true positives, 60 false positives and 305 false negatives. This implies a precision of $200/260 = 0.76$ and a true positive rate (recall) of $200/505 = 0.4$ – we sacrifice high recall for high precision. In other words, we want to make sure that the articles we flag as indicating tourist harm really are related to tourist harm and we sacrifice not detecting all articles that indicate tourist harm for this high precision.

Finally, we train the model with the entire sample and optimal hyperparameters and produce fitted values and predictions for the entire sample of articles. We are left with 52,760 articles as coded containing any violence, 33,127 articles containing fatal violence and 1,361 articles containing tourist harm.

E.4 Post-processing

The precision-recall trade off makes some hand-coding important. As a first step in post-processing, we impose positive human labels over the machine learning labels if the machine proposes a negative – this affects only 30 observations. We then add 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. One of the most common mistakes was driven by condemnations by governments of violence in other countries, for example, when the Egyptian government condemned violence in Tunisia. We excluded these through blanket conditions excluding articles with the word *condemn* in the title.

Other 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 additional human coding to get an impression of the error rate that we imposed through our cutoffs. Because of the revision process, we changed the machine learning part so that we lost some of the targeting in this exercise. We therefore focus on the fact that over 5000 articles were checked by hand – of this, the large majority were in the Factiva sample, so we will assume here that at least 2500

Factiva articles were checked by hand. In the overall sample that we confirmed by hand, the human coder found 114 false positives and 493 false negatives. However, because of the fact that these are fitted to the data for the LexisNexis sample, we focus on the performance in the Factiva sample. Here we get 45 false positives and 401 false negatives. True positives in this sample are 501, so the precision that we get is over 90 percent. The true positive rate is $504/905 = 0.56$. This is better than the cross-validation exercise. One explanation for this is that we reduced noise in the earlier post-processing steps by, for example, excluding shark attacks and condemnations.

The resulting distribution of coded attacks is displayed in Figure B11, 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 an attack 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 our machine learning method plus the human coding in the post-processing step. This provides a good confirmation of the decreased rate at which positives could be in the sample.⁵ Figure B12 shows aggregates of our coding at the dyad/month level. This is the information we are exploiting in the article.

Importantly, this true out-of-sample test makes us confident that the around 1000 observations we have coded as tourist attacks really constitute reports on tourist harm with a very high probability.

E.5 True Out-of-sample Test

Given the supervised learning set-up here, we do not expect the trained algorithm to detect violence against tourists well in other countries.⁶ However, it is interesting to

⁵The few positives for low scores were identified from sources with very few articles as we sampled the 100 top articles from all sources.


⁶If we were to train a generalizing algorithm we would, for example, filter out country- and location-specific vocabulary first. A simple Tf-idf at the country level would be an attractive and simple way to do this.

see whether it does generalize somewhat.

Figure B13 shows the evaluation of a corpus of articles on Sri-Lanka from different sources than our sample sources. Sri-Lanka is a good country to check our algorithm because it is far away from the training sample and had a vicious, out-of-the-blue, terror attack on tourists in April 2019. The figure shows the average evaluation of the ensemble by month. The algorithm clearly picks up the attack remarkably well with scores shooting up in the month of the attack and then decreasing again as the discussion of the attack declines again.


The second finding from this exercise is that the “overfitting” is also noticeable in the absolute level of predicted values which rarely cross the 0.5 line. This is why we show the overall average of the continuous score in Figure B13. If one wanted to use the algorithm in true out-of-sample one would therefore need to adjust the threshold downward from 0.5 or de-mean at the country level to produce a useful time series. We still think it is absolutely remarkable that our supervised learning generated something that is able to do a domain transfer without even training for it specifically.

Figure B1: Precision-recall curve for dictionary predicting fatal violence




Notes: Figure displays the precision-recall curve of a simple dictionary method using a dictionary indicating fatal harm. Simple counts are used to rank articles. No cross-validation is needed.

Figure B2: Precision-recall curve dictionary predicting tourist harm




Notes: Figure displays the precision-recall curve of a simple dictionary method using a dictionary indicating tourists were harmed. Simple counts are used to rank articles. No cross-validation is needed.

Figure B3: Precision-recall curve XGBoosted trees predicting fatal violence




Notes: Figure displays the precision-recall curve of a XGBoosted trees classifier from a 5-fold cross-validation with the optimal hyperparameters.

Figure B4: Precision-recall curve XGBoosted trees predicting tourist harm




Notes: Figure displays the precision-recall curve of a XGBoosted trees classifier from a 5-fold cross-validation with the optimal hyperparameters.

Figure B5: ROC curve ensemble predicting any violence




Notes: Figure displays the ROC curve of an ensemble between naive Bayes, random forest and XGB classifier from a 5-fold cross-validation with the optimal hyperparameters.

Figure B6: ROC curve ensemble predicting fatal violence




Notes: Figure displays the ROC curve of an ensemble between naive Bayes, random forest and XGB classifier from a 5-fold cross-validation with the optimal hyperparameters.

Figure B7: ROC curve ensemble predicting tourist harm




Notes: Figure displays the ROC curve of an ensemble between naive Bayes, random forest and XGB classifier from a 5-fold cross-validation with the optimal hyperparameters.

Figure B8: Precision-recall curve ensemble predicting any violence




Notes: Figure displays the precision-recall curve of an ensemble between naive Bayes, random forest and XGB classifier from a 5-fold cross-validation with the optimal hyperparameters.

Figure B9: Precision-recall curve ensemble predicting fatal violence




Notes: Figure displays the precision-recall curve of an ensemble between naive Bayes, random forest and XGB classifier from a 5-fold cross-validation with the optimal hyperparameters.

Figure B10: Precision-recall curve ensemble predicting tourist harm



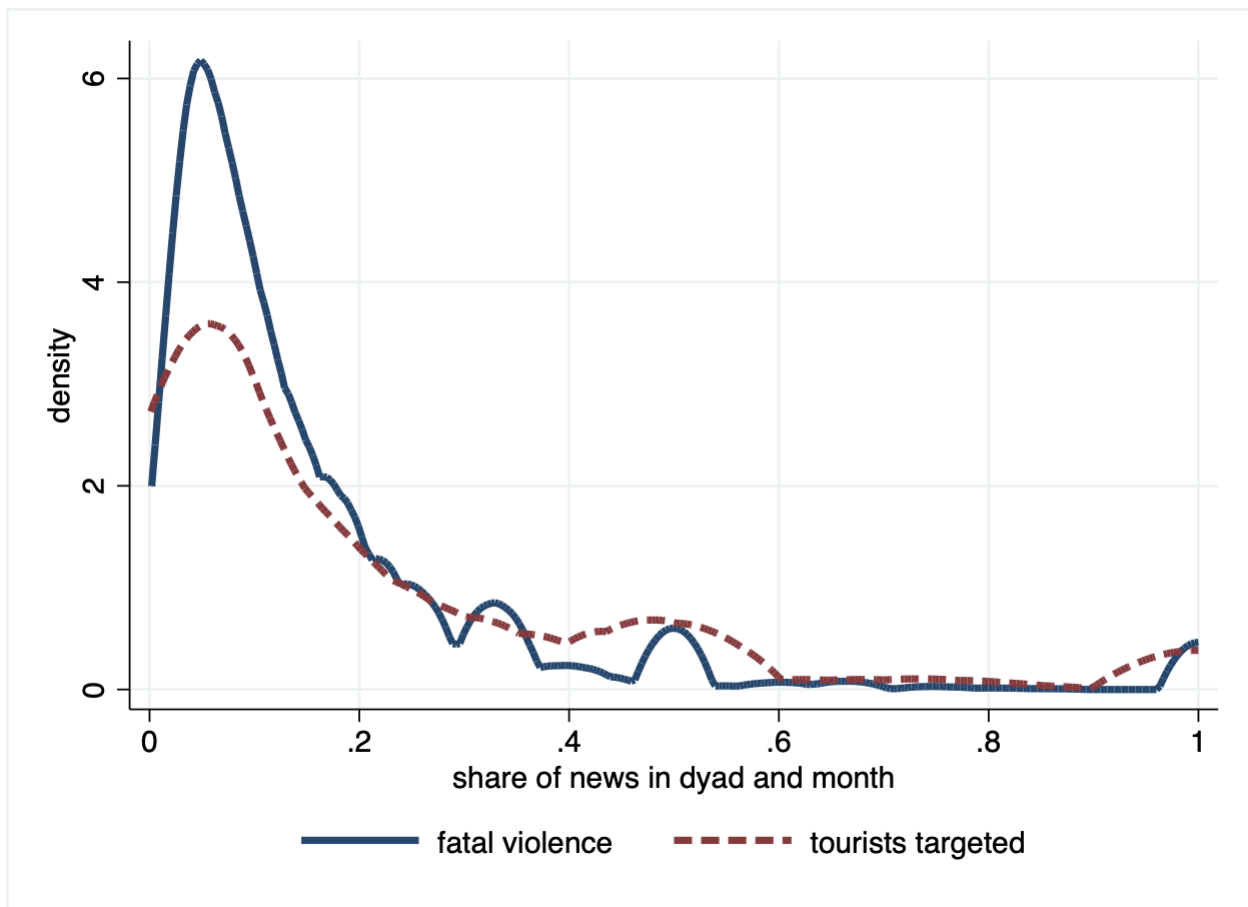
Notes: Figure displays the precision-recall curve of an ensemble between naive Bayes, random forest and XGB classifier from a 5-fold cross-validation with the optimal hyperparameters.

Figure B11: Predicted conditional probabilities of class labels after computing ensemble: clear separation of data is achieved



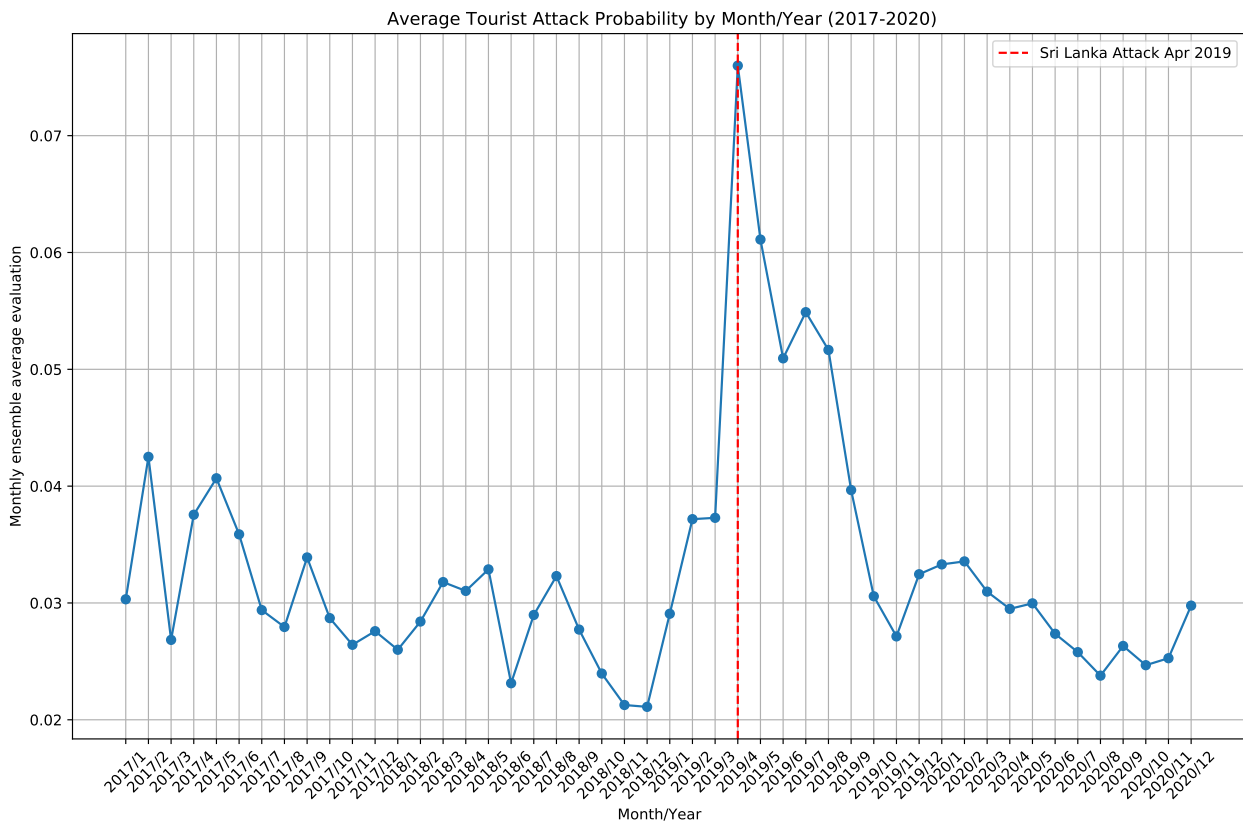
Notes: Figure displays the kernel densities associated with the empirical distributions of the predicted class labels after the ensemble. The data are clearly separated between the two classes.

Figure B12: Resulting distribution of share of articles on fatal violence and violence against tourists for dyad/months



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

Figure B13: Domain-transfer: terror attack on tourists in Sri-Lanka April 2019



Notes: Figure shows the average continuous score for tourist harm provided by the ensemble.

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	Sydney 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 Star	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	Corriere 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 each 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	$\hat{P}_k(Y_i = 1 D_i)$
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 covering violence with fatalities along with the estimated $\hat{P}_k(Y_i = 1 | D_i)$.

Table B3: Example news headlines coded as covering violence against tourists

Country	Year	Month	Headline	$\hat{P}_k(Y_i = 1 D_i)$
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 iç	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 iç	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 covering violence against tourists along with the estimated $\hat{P}_k(Y_i = 1 | D_i)$.