Terror and Tourism: The Economic Consequences of Media Coverage

Timothy Besley
Thiemo Fetzer
Hannes Mueller

This version: May 2020

(December 2019)

Terror and Tourism:
The Economic Consequences of Media Coverage

Timothy Besley   Thiemo Fetzer   Hannes Mueller[*]

May 1, 2020

Abstract

We combine monthly aggregated and anonymized credit card data on tourism spending from 114 origin countries in 5 tourist destinations (Turkey, Egypt, Tunisia, Israel and Morocco) with a large corpus of news coverage of violent events. This allows us to separate the effect of events at the destination from the effect of news reporting at the origin of cardholders. We develop a model of tourist beliefs, actual violence and news coverage to quantify the effect of coverage. When a country is perceived as dangerous, aggregate spending falls by 55 percent. More than half of this effect comes from media coverage.

JEL Classification: O1, F5, F1, L8

1 Introduction

Media reporting can have a immense impact on how we perceive risks. In 2016, terrorism in the US caused less than 0.01% of all deaths but was covered by newspapers more than any other cause of death – including cancer and heart disease.[1] Often this focus on terrorism is not discounted by those who read the news and so we feel terrorized by something that is, in effect, costing relatively few lives. Hence, the economic consequences of terrorist incidents may depend on the way that they are reported and reactions to this. In psychology such biases have received a lot of attention. For example, individuals, might use so-called *availability heuristics*, whereby the ease at with which information on a phenomena is recalled, affects an individual's estimate of the likelihood of it recurring which can change key economic decisions.[2]

The macroeconomic effect of these heuristics could be substantial because media reports will be able to shape perceptions of how safe a country or region is. This is most obviously the case for developing countries for which international news is rare and tends to follow bad events. If this news selection is not discounted by those who read it, it can shape the international image of affected countries, and lead to greater economic isolation. Given the economic and social benefits of openness and economic integration in general (see Frankel and Romer, 2008; Melitz and Trefler, 2012) and tourism in particular (see Faber and Gaubert, 2016), the way media reports on countries could thereby have detrimental consequences for economic development.

But, while this account of the consequences is plausible, quantifying the independent effect of media coverage is difficult. In particular, it is challenging to separate the economic impact of an event from the impact that is due to the news coverage of the event. This paper makes headway on this in the case of tourism. Potential tourists are likely to rely predominantly on media reports in their home countries to learn about a potential travel destination and these reports need not be particularly good in informing potential tourists about the objective risks that they may face. Moreover, if origin countries report differently on the same event then tourists from these different

---

[1]Following Combs and Slovic (1979), Shen et al. (2018) compares causes of deaths in 2016 to news reporting on different causes of death by the New York Times and other news outlets. The webpage *Our World in Data* compiled a striking summary based on this effort.

[2]See Tversky and Kahneman (1973). Recent research in psychology is reviewed in Zhu et al. (2018).

countries will react differently to the same event.

To shed some light on this, we study violence and news coverage of the underlying violence exploiting two unique datasets. First, we use aggregated and anonymized spending data from Mastercard (hereafter "aggregate spending") for one hundred and fourteen origin countries in five destinations (Egypt, Israel, Morocco, Tunisia and Turkey) as a measure of tourist activity.[3] This proprietary data offers the advantages of being available monthly and being *dyadic*, i.e. it reflects aggregated spending from an origin country (where the cardholder accounts are registered) and a destination country.[4] This allows us to study the sharp declines of economic activity caused by events and media coverage.

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

(Figure 1)

The figure highlights that aggregate tourism spending falls immediately in the months following the attack. However, tourism spending from Britain dropped much more compared to tourism spending from Germany. One candidate explanation, which we explore here, is that this heterogeneity is down to the differential intensity of news reporting on the event in the two tourist origin countries.

To capture the news environment, we build a novel data set built on a corpus of nearly half-million news articles. This covers articles from fifty-seven origin countries covering our five destinations from 2009 to 2016. We auto-translate all articles to English and identify the subset of articles that report on fatal violence or violence against tourists. Given the size of the corpus we use methods from computational

---

[3]Mastercard made the anonymized and aggregated data available subject to robust privacy and data protection controls " and in line with their principles guiding the ethical collection, management and use of data.

[4]Since tourism is an export industry, this is similar to exercises that measure international trade flows at a dyadic level.

linguistics and supervised machine learning to develop an automated classifier of reports on violence. The classified articles then allow us to construct a measure of the *relative intensity of reporting* on violence at the dyad-level that can then be matched with the dyadic data on aggregate spending.

While the directional effect of news reporting makes intuitive sense, little is know about the nature of news reporting and the magnitude/duration of responses to it. Our key contribution is therefore to estimate the *quantitative* effect of news reporting on terror events in cases which had macroeconomic importance. We first show that coverage is short-lived and differs dramatically across dyad with differential media coverage in origin countries generating heterogenous responses in tourism spending around violent events in line with Figure 1. We then document a robust relationship between the intensity of reporting on violence and subsequent drops in tourism activity. These results are found even after fully accounting for destination- and origin-specific variation. Moreover, there are no discernible pre-trends in either news reporting or spending. Lastly, we also use an instrumental variable approach leveraging the distribution of nationalities of casualties around known events, yielding very similar results. The effects that we find are sizeable; we estimate that if media-reporting on a specific dyad switches from reporting on topics unrelated to violence to covering only stories about tourists being targeted, then tourism spending drops by about 42 percent a month later. The effect of news reporting that we find comes from the *share of news* coverage dedicated to violence. Furthermore, this effect is persistent, lasting for about nine months following the coverage.

Motivated by these reduced-form findings, we fit a statistical model of violence, reporting and spending to the data to help shed more light on the underlying mechanisms. We assume that tourists try to avoid destinations they perceive to be "dangerous" and we model the underlying risk as a latent state variable which is both related to reporting and violence.[5] We provide a model of how risk perceptions at different origins about this state are formed. In order to study the effect of media coverage we provide a model of belief formation based on events and media coverage.

---

[5]Our model here is motivated by our own previous work in Besley and Mueller (2012) and Besley et al. (2015). It also fits to the idea that humans are categorical thinkers as proposed by Tversky and Kahneman (1974).

As we fit the model, we differentiate between two kinds of tourist behavior. The first type we call "naïve" behaviour which uses news to form beliefs without adjusting beliefs to take account the fact that local media coverage does not capture risks at a destination perfectly. This has two dimensions (i) tourists rely exclusively on news sources that are available in their home country and (ii) ignore the fact that these news sources select news. Such naïve beliefs overreact to isolated news stories and ignore that intense reporting on other issues can "drown out" reporting about violent events - we interpret these features as the use of an availability heuristic in our context. We contrast this with "sophisticated" beliefs based on objective measures of violence.[6] This gives rise to two different sets of beliefs of how "dangerous" a destination is.

Based on this framework we first construct these different sets of beliefs from the data. We then show that, even if we control for sophisticated beliefs, changes in naïve beliefs are still strongly associated with aggregate tourism spending. Using a grid-search we then estimate the form of tourist behavior that best fits the observed patterns in aggregate spending in response to violence. We find that if a destination is regarded as dangerous by all tourists, spending drops by around 55 percent and about more than half of this effect comes from media coverage.

The separation into naive and sophisticated beliefs allows us to study the role of press freedom. We show that the effect of naive beliefs on spending is much larger in countries with a free press. This corroborates our view that naive beliefs capture the impact of signals transmitted through local media reporting. And since we model both reporting and risk perceptions together, we can also show explicitly how high levels of reporting of non-violent events can effectively "drown out" reporting about violent events. As a result, naïve beliefs in our model do not react at all to news on violence if the overwhelming majority of news is on other issues.

Finally, our model-based approach also gives insights into the pattern of persistence of media-coverage on spending. We hypothesize that booking vacations depends on beliefs held at the booking or cancellation date. We then estimate from the data what share of tourists appear to have booked or cancelled their travel at each date. This exercise suggests that the pattern of persistence in the spending response

---

[6]Both labels are obviously a modelling device as we only observe aggregate behavior by month and origin/destination dyad.

is "explained" by more than 40 percent of tourists booking their vacations more than three months in advance implying that there is some persistence in the response of spending to violence and news reporting.

This paper contributes to a large and growing literature which studies the impact of the media on economic and political outcomes (see Strömberg, 2015, Prat and Strömberg (2011) for reviews).[7] One of our contributions is to look at news events and coverage across multiple countries which complements a body of work that has focused in detail on US news coverage and its consequences. For example, Eisensee and Strömberg (2007), shows that news on droughts affect US disaster relief. Similarly, Durante and Zhuravskaya (2018), suggests that offensives in the Israel Palestine conflict are strategically aligned to minimize news coverage in the US, while Jetter (2017) suggests that Al Qaida activity may be endogenous to preceding US television news coverage.[8] We contribute to this literature by providing dyadic news data which we embed in a model of violence, reporting and beliefs.

The paper provides a link between empirical work on the media and a large literature in psychology which studies how individuals update their belies in response to news. Of particular relevance to our interpretation of the results is the large literature in psychology focusing on availability heuristics, the human tendency to judge probabilities by the ease with which information underpinning them can be recalled Tversky and Kahneman (1973).[9] We offer a specific theoretical model which we fit to the data where individuals are assumed to be categorical thinkers and who update their beliefs in a partially naive fashion based on the news in their own country.[10] Media coverage is an imperfect signal of the underlying categorical state and beliefs,

---

[7]The impact of the media on politics and economics has been the subject of many papers: Stromberg (2004) on redistributive spending, Besley and Burgess (2002) on government accountability, Gentzkow (2006); Bursztyn et al. (2017) on voter turnout, Snyder and Strömberg (2010) on citizen knowledge, DellaVigna and Kaplan (2007); Enikolopov et al. (2011); Adena et al. (2015) on voting patterns, Durante et al. (2019) on the proclivity towards populist rhetoric.

[8]We observe that news coverage sharply responds to violent events, but not in anticipation.

[9]Handel and Schwartzstein (2018) Zhu et al (2020) provide recent overviews of the relevant literature. This is only one among a range of heuristics and biases highlighted in Tversky and Kahneman (1974). Bordalo et al. (2016) provides a nice overview of different approaches.

[10]There is a wealth of evidence from psychology that people use crude categories when making sense of the world. Fiske (1998) provides a review of social psychology literature; Fryer and Jackson (2008) provide a model based on this idea and a discussion of the biases that this induces in decision-making.

and spending, "overreact" to it.

We show that the economic impact of violent events on tourism activity is significantly driven by the underlying news reporting. In this regard our work relates to an emerging literature on how news "shocks" affect economic behavior (see Ramey, 2011). Arezki et al. (2017) documents that news reports on resource discoveries have near immediate impacts on the current account and Eggers and Fouirnaies (2014) document that news about the technical declaration of a recession has a contractionary effect on the economy. This raises questions about how people collect and process news and whether they might overreact to particular forms of coverage (see Azeredo da Silveira and Woodford, 2019, Handel and Schwartzstein, 2018 and Bordalo et al., 2018). Given that we explain a significant amount of heterogeneity in the response to the same facts driven by news reporting, our findings also support the idea of "overreaction" and/or selectivity in access to information. This relates to discussions beyond economics. For example, Singer and Endreny (1993) suggests that hazards are distorted by mass media and that *"emotion, not reason, is likely to govern our response to those hazards for which we depend on the media for information."* (p.41) This is particularly relevant if social media amplify echo chambers (Mullainathan and Shleifer, 2005; Gentzkow and Shapiro, 2011), which can spread fake news and/or entrench extreme views (Barrera et al., 2020; Allcott and Gentzkow, 2017).[11] But our findings are also relevant in light of the Covid-19 crisis as they give a benchmark of how media coverage of death risks affects aggregate spending.[12]

The paper is related to a growing literature which uses text as data. Drawing on national newspapers and a simple dictionary, Baker et al. (2016) construct an economic policy uncertainty index for a host of countries. Hassan et al. (2019) make use of computational linguistics methods to develop a dictionary of terms for the measurement of political risk documenting that increased political risk induces more lobbying activity. Mueller and Rauh (2017), following a similar approach to Hansen et al. (2018), use unsupervised learning, so-called topic models, to assess to what extent Western

---

[11]See Zhuravskaya et al. (2020) for a review of the literature on the political effects of the internet and social media.

[12]In fact, the numbers from our model are close the effects of a complete lock down found by Carvalho et al. (2020).

(US and UK-based) media coverage can predict violence risk. In this paper, we instead deploy supervised machine learning methods to produce a measure of violence in the news. This approach allows us to test explicitly, through cross-validation, how good our method is in identifying reporting on violence in the news.[13] The good performance of our method suggests that this way of producing data from text, i.e. letting the machine figure out which parts of the text are most relevant, might also be useful in other applications.

We also contribute to an emerging literature on the consequences of violence and disorder on trade and economic integration. For example, Besley et al. (2015) measure the cost of piracy in the Gulf of Aden on shipping costs and hence on trade. They emphasize that violence can increase trade costs which the trade literature has shown can have a significant impact on trade flows (see, Feyrer, 2019, Donaldson, 2018). However, the mechanism we focus on here is closely related to Burchardi et al. (2019) who show that better information on a country, due to ancestry, is an important driver of foreign investment decisions. Tourism is an increasingly important sector that supports up to 313 million jobs across the globe.[14] We show that tourists with more information react less to idiosyncratic media reporting because they have a broader informational base and this can be an important factor in the overall effect of violence on economic ties. This is particularly relevant in the MENA region which is one of the least economically integrated regions (Rouis and Tabor, 2012) and where growing economic ties to Europe are important.

There is now a large literature on the economic costs of violence to which this paper contributes. For example, Abadie and Gardeazabal (2003) documents the sizable negative economic consequences of terrorism in Spain and Besley and Mueller (2012) looks at the impact of house prices using regional variations in violence in Northern Ireland. More recently this literature has shifted towards identifying the mechanisms through which lack of security affects the economy. For example, Amodio and Di

---

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

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

Maio (2017) shows how firms bear the direct and indirect costs of violence and political instability. Jha and Shayo (2019) explore how individuals re-evaluate the costs of conflict upon being exposed to financial assets whose prices may be vulnerable to the economic risks of conflict. Brodeur (2018) shows that successful terror attacks in the US reduce the number of jobs and total earnings in the affected counties. This paper adds to this literature by highlighting a new mechanism: news coverage itself can amplify the negative economic consequences of terrorism.

The remainder of the paper is organized as follows. In section 2 we discuss the data used in some detail. Here, we also discuss the supervised learning method through which we make the text data usable for the subsequent analysis. In section 3, we present reduced form results of the average effect of violence on tourism activity before incorporating the news data. In section 4, we propose a statistical model and fit this on both the news and the tourism activity data. Concluding comments are in section 5.

2 Data and Feature Extraction

This paper uses three main data sources: (i) aggregated spending data by origin and destination country, (ii) measures of terrorism and conflict 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.

2.1 Aggregated Spending Data

Mastercard provided us access to an anonymized and aggregated monthly data set which included the number of transactions and number of active cards based on spending in five different countries (Egypt, Israel, Morocco, Tunisia and Turkey) by the country of origin of the card.[15] Where Mastercard's controls may have resulted in, for example, in a blank monthly-dyad observation, we excluded all of those dyads where we have fewer than 60 months (5 years) of data and origin countries with fewer than 3 out of 5 destination dyads. The origin countries in our sample span all con-

---

[15]Mastercard made the anonymized and aggregated data available to us subject to robust privacy and data protection controls.

tinents but trend towards higher income countries and those that are geographically closer to the destination countries.[16] Figure 2 maps all of the origin (card-issuing) countries that we have in the sample.

(Figure 2)

There are notable differences between origin countries with low volume of cards active per month in tourism spending, such as countries like Haiti and Namibia compared to higher volume tourism spending from countries such as Germany and the United States.

The aggregated card data can be a proxy for annual patterns in travel flows: for a small set of countries annual data on travel flows is provided to the United Nations World Tourism Organisation (UNWTO). Appendix Figure A1 highlights that the annually-aggregated card data correlates with the travel flows data very closely. Appendix Table A1 presents further regression evidence highlighting the close fit.

2.2 Data on Violent Events

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

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

As supplementary human-coded sources of data, we also leverage the Georeferenced Event Dataset (GED) provided by the Uppsala Conflict Data Program (UCDP).[17]

---

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

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

Automated Data   We use the Integrated Crisis Early Warning System (ICEWS) database created for the Defense Advanced Research Projects Agency (DARPA) and Office of Naval Research (ONR). This event-level data comprises coded interactions between sociopolitical actors (i.e., cooperative or hostile actions between individuals, groups, sectors and nation states). Similar to the approach used in Fetzer (2019), events are identified in a fully automated way and are extracted from news articles, essentially consisting of triplets based on a subject (a source actor), an event type (indicated by a verb) and an object (a target actor). Geographical-temporal metadata are also extracted and associated with the relevant events. In this paper, we focus on events that have been coded as assaults, which include events like hijacking, suicide bombings and assassinations and along with data on fights or escalations, which includes the use of military force, fights with artillery and tanks and aerial bombing.

The second automated dataset is the GDELT platform which monitors the world's news media from nearly every corner of every country in print, broadcast, and web formats, in over 100 languages, every moment of the day stretching back to January 1, 1979 to produce data on events.[18] GDELT is more inclusive, yet it may also include also more false positives and it also has less stable source material over time and codes the news sources only from 2014 onwards.

Both of these data sources have in common that they aim to identify the "true" set objective violent events. Neither of them provides a measure of the likely salience of an event nor the intensity of news coverage about a violent event across different countries.[19] 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 tourists 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,

---

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

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

we develop a large scale corpus for 57 tourist-origin countries. For each tourist-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 2. 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.[20]

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 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 con-

---

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

sisted 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 D. This produces three different classifiers, indexed by k, which allow us to obtain for each document, denoted by D_i, three estimates of the probability that classifier k contains news coverage of the type that interests us, denoted by $\widehat{P}_k(Y_i = 1|D_i)$, where Y_i is an indicator denoting whether a given document D_i is either covering violent events with fatalities or violent events in which tourists were targeted.

Naïve Bayes methods belong to the class of generative linear classifiers, is known to perform well with textual data and sparse feature sets. Random Forests, on the other hand, are particularly suitable to allow for non-linearities using smaller feature

sets. The only difference between our two Random Forests, is that in one of them we first use Singular Value Decomposition (also referred to as Latent Semantic Analysis which has recently been used in Iaria et al. (2018)) to reduce the dimensionality of the feature space from (tens of) thousand of word counts features into a much lower 100-dimensional continuous score representation of individual documents D_i. These individual components are then used as numeric features in the construction of the classification trees using the random forest formulation.

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

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

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

To chose the cutoff c in (1), we count how often the indicator, $\mathbb{1}(D_i)$ would have been correct for different values of c within the training sample. The Bayes optimal decision rule that maximizes overall accuracy would be to pick a cut-off of $c = 0.5$. Yet, the cross-validation exercise highlights that the class imbalance may result in low precision under this rule. As we are concerned that we get too many false positives resulting in a very noisy monthly measure, we choose a higher cutoff which provides 90 percent precision within our training sample.[21] 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.

---

[21]Our results are robust to using alternative cut-offs. In Appendix Table A6 we use both a higher cut-off with 95 percent precision and the Bayes-optimal cut-off of 0.5.

To reduce measurement error we conducted some ex-post manual coding for the classification of articles indicating violence against tourists. While our results are robust to relying only on the machine-generated output, it is prudent to perform such a manual check and some amount of ex-post refinement. We considered all articles with an ensemble probability indicating violence against tourists above 0.75 along with the top 100 articles ranked by the ensemble estimate from in (1) all origin sources.[22] We set $\mathbb{1}(D_i) = 0$ by hand if we find false positive and set $\mathbb{1}(D_i) = 1$ if we find false negatives. Of 1,082 observations that were marked positive by the algorithm we recoded 103 to negatives, implying that our method did indeed achieve a precision of over 90 percent out-of-sample. In the almost 5,000 additional news items that were hand-coded we only found an additional 608 positives with a rapidly declining rate so that we suspect that the remaining articles will not contain a lot of actual positives. After hand-coding we therefore have 1,587 positives in over 450,000 negatives that feed into our media coverage-based measures of violence against tourists.

In the Online Appendix D we describe the classification approach in greater detail, while appendix Tables B2 and B3 provide some sample headlines of articles coded as covering violence and flagged up as capturing that tourists are targeted by a violent event. In the appendix we also discuss the "mistakes" made by the algorithm and why they are often capturing something indicating risks to tourists. It is therefore no surprise that our results, even in the most demanding specifications, are robust to using only the raw $\mathbb{1}(D_i)$ that come out of our automated procedure and using different cutoffs.

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

---

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

or dictionary-based method.

2.5 Patterns in the Reporting Data

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

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

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

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

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

(Figure 3)

In Figure 3 we plot the point estimates $\hat{\beta}_\tau$, which suggests that there is no anticipatory element in the reporting data. Panels A and B show the measures generated

from our method for classifying articles. Specifically, we construct the share of articles per day that are classified as reporting either fatal violence or tourists being attacked. The patterns suggest a sharp increase in the share right after the event date. This dissipates quite quickly with most reporting occurring on the day of the event and for around two days afterwards. It is important to note that this happens despite the fact that the total number of news stories increases slightly, i.e. we find this relative reporting effect despite increased reporting overall.[23]

Monthly Aggregates In our analysis we use aggregates of our news measures to the monthly level as the dyadic aggregate spend data is only available at the monthly level. Figure 4 reports the mean shares at a monthly frequency for the four countries most affected by violence against tourists in our sample period (Tunisia, Turkey, Israel and Egypt). It depicts the average share across all dyads of monthly events defined by (1) for violence against tourists – the red dashed line with the axis on the right hand side – and (1) for fatal violence on a monthly basis – the blue solid line with the axis on the left hand side.

(Figure 4)

Figure 4 shows a lot of variation across time in reporting for all countries. But there is also considerable variation in the intensity of reporting across destinations. Reporting on violence is often a considerable part of reporting on Tunisia. At the time of the Sousse attack, for example, violence against tourists occupied around 40 percent of all news. Reporting in Egypt, Turkey and Israel is more intense for fatal violence than it is for violence against tourists. However, this coverage never occupies more than 10 percent of the news. The most extreme example is Israel where news on violence never exceeds 7 percent of reporting and violence against tourists never more than 3 percent.

3   Reduced-form Evidence

In this section, we present reduced-form evidence on the relationship between spending and violence. We begin with country-level violence and then present the

---

[23]In Appendix A.1 we provide some further evidence shedding light on which event characteristics are associated with, on average, more extensive media coverage.

findings with dyad-specific news coverage. Both sets of results show that there is a highly robust relationship between aggregated spending and violence which depends not only on violent events but how intensively they are covered by news media.

3.1   Country-level Violence

We begin by looking at the relationship between the (log of) spending by origin country h in destination country d at date t, denoted by y_{hdt}, and violent events using the following specification:

$$y_{hdt} = \alpha_{hd} + \alpha_{ht} + \alpha_{dm(t)} + \xi v_{hd,t-1} + \varepsilon_{hdt} \tag{2}$$

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. The inclusion of origin country-time fixed effects capture a range of seasonal, political and economic effects in the origin country and the destination/month fixed effects capture seasonal patterns in spending in the different destinations. The first set of fixed effects implies that we are effectively looking at the rate of tourism spending from origin countries in the five different destinations. Any fluctuation in the overall flow of spending into the region will be absorbed by the home country/time fixed effect. In addition to aggregated spending as an outcome variable, we will also consider the number of active cards on which spending is happening as an outcome measure.

Our core violence measure, denoted by $v_{hc,t-1}$, is lagged by one month to capture the possibility that tourism reacts to past violence. We expect to find that $\xi < 0$ in (2), i.e. violence deters spending. We use four sources of data on violent events at the country level in different versions of (2). 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 2 reports regressions from the specification in (2) and shows compelling evidence of a negative link between violence and card spending. Columns (1) through (4) show that there is a significant correlation between all five measures of violence and the level of tourism spending in a country. The size of the coefficients is in the range of 4% to 7.6% decrease in 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 (2) with a four dimensional vector comprising the first four principal components, with the results reported in column (5). In line with the results in columns (1) through (4), we find a robust negative relationship between principal components 1, 2 and 4 and spending. In terms of magnitude, spending falls by about 7% with an increase in the first component and by about 4% with the second component. Since this summarizes information from a range of sources, we will use this representation of violence in the analysis that follows.

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

(Table 2)

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

To allay the concern that results based on (2) could be explained by different time trends between times/places which 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. It is comforting to see in Appendix Figure A4 that 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.

19

3.2 Exploiting Dyadic Variation in News Coverage

We now shift focus away from the violence data in an attempt to understand whether there is a differential effect of media reporting about violence on tourist activity. This allows us to explore dyad-specific responses to violence as mediated via news coverage.

Core results The core specification in this section extends (2) to:

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

where, as above, our independent variable is the (log of) tourism spending. In this instance, we measure v_{hct-1} using the first four principal components of violence across the seven violence measures used in the previous section; these capture more than 90 percent of the variation in the country-level violence variables. The variable n_{hdt-1} is our dyad-specific news variable.

In some specifications we will use α_{dt} instead of $\alpha_{dm(t)}$, i.e. we include destination by time fixed effects. These fixed effects capture *all* variation at the destination/time level including news events which are 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.

Define $B_{hdt} = \sum_{i \in hdt} \mathbb{1}(D_i)$ as the monthly *count* of negative news stories in dyad $(h.d)$ at date t, either about fatal violence or attacks on tourists, based on equation (1). Then our core variable to represent news coverage in a dyad is

$$n_{hdt-1} = \frac{B_{hdt-1}}{N_{hdt-1}} \tag{4}$$

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 (4) reflects the news coverage of violence as a *share* of all news. This captures the idea that news coverage of violence affects tourists more when they are important relative to other news. Thus, if bad news stories are swamped by other stories, they will have less

impact on tourism spending. We confirm this in section 3.2.

Before presenting the results, it is worth stressing that our analysis is only exploiting within-dyad variation which absorbs all factors like distance or cultural factors. In addition, we are including home country by time fixed effects, α_{ht}, de facto absorbing a host of factors that may drive the level of tourism activity that is explained by origin-country level idiosyncrasies (such as holiday periods, which may differ across countries). In this way we are modelling the rate of tourism activity for a given destination among our sample of five countries *relative to* the overall amount of tourism originating in country h. Hence all magnitudes are based on comparing the attractiveness of our destinations compared to the other five destinations in our data rather than other parts of the world. Thus, we are only able to say whether tourism to Turkey decreased after the terror attacks in the country relative to Egypt, Israel, Tunisia and Morocco. Arguably this is a conservative approach, since there could be reputational externalities whereby tourists shy away from the entire region due to the turmoil in one of our countries. Finally, it is important to keep in mind that there are no discernible pre-trends in coverage but that it responds sharply to terror events as shown in Figure 3.

The results from estimating (3) are in Table 3 where the top panel (Panel A) is based on news coverage of violence against tourists while the bottom panel (Panel B) is for reporting on fatal violence in general. In columns (1) through (3) the independent variable is the (log of) card spending. Column (1) of Panel (A) shows that if the share of stories about tourist violence were to go from zero to one then tourism spending would fall by 0.552 log points or 42 percentage points. This results holds up in column (2) when we add the controls for violent events and the coefficient stays roughly similar. Hence, news coverage is clearly containing valuable information over and above that in the underlying violent events themselves.

Column (3) in Table 3 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 macro-economic developments or violence at the destination. As a result, in this specification we rely only on the differential intensity in news reporting across different reporting countries. The coef-

ficient on the share of bad news (based on tourist violence) remains highly significant although the coefficient in this saturated specification falls to 0.205 log points. This continues to suggest that a significant part of the overall effect in columns (1) and (2) is driven by pure news reporting. Note, that 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 3 repeat the specifications in the first three columns but with the log of active cards as the dependent variable. The estimates are similar which is important as it indicates that the results on spending are not only due to changes at the intensive margin in which tourism spend is less, but rather on the extensive margin as fewer tourists travel to a country following media-coverage of violent events.

Panel B of Table 3 repeats the specifications in Panel A except for measuring reporting on all fatal violence rather than just attacks against tourists. The coefficients are only slightly smaller compared to those presented in panel A. We also find statistically and economically significant relationships throughout although we lose statistical significance in the most saturated specification in column (6). The results suggest that differential intensities in media coverage of violence may have an important independent direct effect on tourism travel and spending, which is particularly relevant given that violence more broadly, not necessarily directed at tourists, is a lot more common.[24]

(Table 3)

Taken together, the results in Table 3 suggest that the media have the capacity to influence tourism spending. The fact that news reporting has a negative correlation with tourism spending, even when controlling for an array of fixed effects and more objective measures of violence, shows that specific reporting matters over and above the violence itself. That said, it is hard to quantify the relative importance of violence

---

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

and news from the reduced form analysis. In section 4 we therefore provide a model of tourist perceptions, news and violence which allows us to distinguish the news effect much more carefully from the effect of violence.

Further Results   This section provides additional results which lend further support for the idea that intense reporting on violence can impact tourism spending. We also address concerns about reverse causality by using the composition of victims by nationality as an instrument for the intensity of news coverage.

News Coverage Intensity   So far, we have assumed that the effect of bad news, B_{hdt-1}, on tourism is weaker when a lot of other news about a country appear at the same time, i.e. that relative frequency $\frac{B_{hdt-1}}{N_{hdt-1}}$ is the right measure to use. One reason to do this is to lower measurement error given the large variation of N_{hdt-1} in the sources we have in our data. But it will also play a crucial role in the theory-based approach that we propose in section 4.

Table 4 presents a range of specifications testing for relative news effects. In Panel A we analyze news reporting on violence against tourists. We first explore the difference between using the count and the share of news coverage by including both in equation (3). The coefficient on the relative measure essentially unchanged compared to column (2) of Table 3. We get an additional effect from the number of bad news, B_{hdt-1}, but this it is small in size - bad news items lower spending by 0.7 percentage points the following month. In column (2) we control for violent events using the four principal violence components.

In columns (3) and (4) we explore how violence intensity matters by constructing a dummy variable for whether news reporting is in different quartiles of the distribution of dyad/months with $B_{hdt-1} > 0$. We find a striking pattern in which spending falls continuously with increasing relative reporting intensity. This is direct support for the idea that it is reporting intensity as captured by $\frac{B_{hdt-1}}{N_{hdt-1}}$ that impacts the tourism spending response. The top quartile experiences a drop of tourism spending of over 25 percentage points whereas the lowest quarter essentially shows no effect.

In Panel B of Table 4 we find the same pattern of coefficients for news reporting on fatal violence. Here, only the relative intensity measures reduces spending and spending falls with higher quartiles. The coefficients are also lower with an 11 percentage

point drop in the top quartile where there are now a lot more dyad/month observations. The overall damage of these news stories on spending remains economically important.

(Table 4)

Timing Remember that reporting on violence focuses on very recent events. To capture the dynamic effects of news we can therefore use simple lags and forwards on our main variable, n_{hdt-1}. We show such an analysis in Figure 5. Panel A shows that the coefficients on leads and lags on the tourism news count and Panel B shows the coefficients on leads and lags of the fatal stories news count.

Figure 5 suggests that there is a strong negative effect which starts with the news story and intensifies in the first few months afterwards.[25] In addition, we don't see significant effects on spending leading up to violent events in spite of the fact that such events do tend to be serially correlated. In Panel B we observe that lagged effects persist for some time and that spending does not appear to recover fully. This is consistent with long term effects of violence that have been remarked upon in the macroeconomic literature (Mueller, 2012).

(Figure 5)

These patterns suggest that either perceptions are extremely sticky or that cancellation and search costs imply that tourists who perceived a country as dangerous when booking their holiday are unlikely to switch back to that destination even when the situation has calmed down. To get a more quantifiable sense of the lagged effects of news reporting we model the share of tourists that take a decision regarding their holiday destination explicitly in section 4.

An IV Approach One possible concern is that there are common forces driving reporting and tourism spending if, as terrorists strategically consider their targets to give them maximum media exposure. Another concern is also measurement error since we proxy the news landscape at origin by only one or (at most) two newspaper sources.

---

[25]Note, however, that we often have different important events following one after another so not too much weight should be put on these coefficients.

We now propose a way of addressing this via an IV approach which uses the nationalities of victims of violence as an instrument. The idea is that demand for reporting on an event increases dramatically if those who are exposed to news stories empathize more with the victims. This may lead to more intensive reporting in the country of origin of the victims and in countries which are "similar" to the origin country of the victims e.g. share a common language or are in the same region.[26]

For a subset of events, we were able to identify the nationality of the individuals killed. If we assume that the distribution of casualties by nationality is random *conditional* on an attack occurring and our fixed effects, then we can use the composition of casualties by country as an instrument for news coverage on a specific dyad. The Sousse attack serves as motivating evidence as shown in Figure 1, since it is notable that while the bulk of casualties was British, 24 of reported fatalities were non-British.

For each attack where we have relevant data, we construct a variable, z_{hdt} to denote the nationalities of the casualties. To amplify power we also construct a version of the measure that attributes casualties from nations to their contiguous neighbors or to countries that share a common language. We then run a first stage regression to explain news reporting on a dyad with the structure of casualties

$$n_{hdt-1} = \alpha_{hd} + \alpha_{ht} + \alpha_{dt} + \upsilon z_{hdt-1} + \zeta_{hct}. \tag{5}$$

In Table 5 we present the results from the IV regression using our violence against tourists based news measure in columns (1) - (3) and the general news measure indicating violence with fatalities in columns (4)-(6). Panel A reports the first stage regression based on equation (5) while Panels B and C report second-stage IV regressions with either the log of spending or the log of active cards as the independent variables.

(Table 5)

In column (1) of Panel A we confirm the idea that having a causality from the country in question significantly predicts news coverage in a dyad. Column (2), also

---

[26]This is not a new idea. Adams (1986) already observed the dramatic differences in reporting depending on the nationality of the victims. Singer and Endreny (1993) contrast, for example, reporting in the US on the heart defect of an american infant compared to reporting on a famine in Ethiopia.

counts casualties for contiguous countries, the region of origin (e.g. West Europe, Northern Europe, etc.) and whether they are from a country speaking the same main language. Thus, for example, a British victim may also influence potential tourism in Ireland. We see that these are also strongly significant in predicting news coverage and the F-test on the instruments is, not surprisingly, much larger. Column (3) shows that an instrumentation based solely on the indirect exposure, i.e. contiguity, or sharing the same language together predict news coverage in a dyad quite well.[27]

In Panel B, we show that using the IV approach, there is a robust and negatively significant effect of the share of news stories on the log of card spending. In all cases, we include violence controls, origin/time and destination/month fixed effects. In all cases we find negative coefficients which are much larger than for the OLS. This could be explained by attenuation bias and/or by interpreting the coefficient as a LATE effect where our instrument is picking up something particular about events with identifiable casualties. There are arguments for both. First, we have some dyads for which we have no or little reporting and the news landscape at the origin might indeed look different than what we capture. Second, the coefficients increase much more in size for fatal violence which makes sense as we are now indirectly imposing fatal violence which affected foreigners, i.e. the LATE effect is driven by terrorism instead of military operations. Panel C shows that these results become even more robust when we look at the (log of) active cards as the independent variable.

Taken together, these results increase our confidence in our interpretation of the results as indicative of a causal effect of news coverage in tourist perceptions that depends on origin-specific reporting of violence. And it makes sense that idiosyncratic differences in reporting in dyads represent more than variation in objective risk factors faced by tourists.

Robustness In Panels A and B of Appendix Figure A5 we show that our results for violence against tourists are broadly carried by all the different destination countries we study. However, there does appear to be a larger effect for Tunisia, which saw

---

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

much more violence targeted at tourists.[28] In columns (1) - (4) of Appendix Table A6 we highlight that the results are robust to controlling for dyad-specific linear trends, in addition to destination specific non-linear time trends. We can also control for time-varying exchange rate movements at the dyad level, some of which may be triggered by changes in tourism-based revenues, and find that results are broadly the same. We further show that the results are robust to using alternative cutoffs for the classification of individual articles in the news corpus for which we did not deploy hand-coding. Results are in columns (5) - (8) and are robust except for the lax cutoff with destination time fixed effects. They are also robust to using our cutoff with 90 percent precision without hand-coding. This highlights that using the output of the machine learning method we deploy even without the hand-coding refinement would have led to very similar results.

4   Modelling the Impact of News Coverage on Tourist Spending

In this section we provide quantitative estimates of the impact of news coverage on tourism spending based on a model of belief formation which has rational and behavioral elements. The model allows us to separate the effect of events and media reporting and to formulate projections about extreme, out-of-sample events like the Covid-19 crisis. We also show that such a model can do a much better job than the reduced-form approach at capturing the timing and magnitude of expenditure fluctuations.

The core theoretical idea, drawing on the extensive psychological literature is that beliefs are partially formed using *availability heuristics.* Here this means relying exclusively on readily available news in an individual's home country without realizing that this may not be representative of all available information. This squares with the observation above that coverage reacts to terror events for a few days and is more intense if the victims were from origins that are linguistically or geographically close.

---

[28]In Figure A5 we do not include the event-controls. Once these are included, Turkey does not show an additional media effect which indicates that reporting on Turkey had relatively fewer outliers but happened more proportional to violent events.

It contrasts with a Bayesian approach which aggregates information from all available sources. We characterize rational decision-making as using sources such as GTD or GDELT which aggregate a much wider range of information. In the limiting case of there being only rational decision makers, this would imply that dyad-specific news coverage should not matter at all.

4.1 The Model

Core Features 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. At each date t, $\hat{P}_{hdt}$ is the belief that a destination country d is dangerous as perceived by potential tourists in country h. The different dates at which people book their travel spending is determined by a weighted average of past beliefs. Hence we suppose that equation (3) is replaced by:

$$y_{hdt} = \alpha_{hd} + \alpha_{ht} + \alpha_{dm} + \xi \sum_{\tau=0}^{-9} \omega_\tau \hat{P}_{hdt-\tau} + \varepsilon_{hct} \tag{6}$$

where ω_τ is the weight on each lagged value, i.e. at date τ. Equation (6) also has the same fixed effects as equation (3). In this framework, n_{dht} and v_{ht} affect spending through affecting beliefs $\hat{P}_{hdt}$. Since we do not directly observe beliefs, we posit that it 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, that only Ψ_{hdt} varies at the dyad level. We will specify $\Gamma(\cdot)$ below combining two approaches, one which refer to as "sophisticated" and the other which we refer to as "naïve".

Sophisticated Beliefs These are based on reading multiple information sources and, since this information is not confined to any specific country, such beliefs are common across all origin countries. This is "as if" the individual forming these beliefs can observe Ω_{dt}.

The empirical model that we use is based on Besley and Mueller (2012) and Besley

28

et al. (2015) who suppose that the underlying state can be modeled as a Markov process which switches between being dangerous and safe. Assume that our measure of violent events is distributed normally i.e., $v_{dt} \sim N(\mu_{sd}, \sigma_{sd}^2)$. This allows the mean, μ_{sd}, and the variance, σ_{sd}^2 to vary with the state s_{dt}. At each date, there can be a transition between states where p_d is the probability of transitioning from dangerous to safe and q_d is the probability of transitioning form safe to dangerous. This gives a parameter vector for the model with six elements for each destination country d, summarized as $\theta_d = \{\mu_{0d}, \sigma_{0d}^2, \mu_{1d}, \sigma_{1d}^2, p_d, q_d\}$.[29]

Sophisticated beliefs estimate that the probability that a destination is dangerous at time t by

$$\Pi_{dt} = \Pr(s_{dt} = 1 \mid \Omega_{dt}, \hat{\theta}_d), \tag{7}$$

$\hat{\theta}_d$ is an estimated parameter vector. This belief is updated using Bayes rule as information 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

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

is the prior from the previous period. For v_{dt}, we use the principal components across the data from different violence sources.[30] 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 $t - 1$. Estimates of $\hat{\theta}_d$ which go into constructing (7) are described in Appendix Table A7. They show strong persistence in the state in all five destination countries.

Figure 6 gives our estimates of Π_{dt} for Egypt, Tunisia, Turkey and Israel. This

---

[29]This vector can be estimated from the data on violent incidents using an EM algorithm (see Hamilton, 1990).

[30]To aggregate the different components into one number, v_{dt}, we use the point estimates on the first two components from Table 2.

approach permits a classification of whether a country is deemed to be dangerous or safe based on the level of violence but in a destination country-specific manner. Thus, unlike equation (2), the effect of a given change in v_{dt} is heterogeneous across different destinations depending the history and persistence of violence. This makes sense; what would be deemed to be violence pointing to a state of danger in, say, Israel (bottom right) is different from what would be considered danger in Tunisia (top right). These differences show up in the estimates of μ_{sd}. The probabilities based on equation (7) indicate that Tunisia was the first of our destinations to become dangerous based on the level of violence and this was followed by Egypt and Turkey. Danger in Israel based on this method is less persistent.[31] Crucially, sophisticated beliefs, as displayed in Figure 6 show dramatic differences to the news coverage we displayed in Figure 4.

(Figure 6)

Naïve Beliefs These use information based exclusively on news coverage in the country where the beliefs are formed, i.e. Ψ_{hdt}. We interpret this as a form of availability heuristic where news at home is more available and salient which we will show can lead to an overreaction to single events and is heterogenous across dyads. To model this, suppose that the belief that the state is dangerous is generated by a common distribution based on the fraction of news stories reporting violence, η_s. Specifically, the distribution of B_{hdt} and N_{hdt} follows a negative binomial distribution based on η_s given by:[32]

$$f\left(B_{hdt}, N_{hdt} \mid s\right) = \begin{pmatrix} N_{hdt} \\ B_{hdt} \end{pmatrix} (\eta_s)^{B_{hdt}} \left(1 - \eta_s\right)^{(N_{hdt} - B_{hdt})}. \tag{8}$$

This implies that beliefs will react strongly to local news in some dyads if B_{hdt} changes more dramatically than expected. To obtain an estimate of η_s from the data we use

---

[31]The model fits less well for Morocco since the differences in μ_{sd} are minimal between what the model picks as the two categorical states.

[32]We use the negative binomial distribution to capture the fact that news reporting has extremely fat tails. The distribution could be justified by a stopping rule for media consumption but we follow it purely for pragmatic reasons.

our estimate of Π_{dt} and the news data so that:

$$\hat{\eta}_s = \frac{\sum_{hdt} \Pi_{dt} \left[\frac{B_{hdt}}{N_{hdt}} \right]}{\sum_{hdt} \Pi_{dt}}, \tag{9}$$

where summation is over all dyads and time periods. Equation (9) is a weighted sum of the news share variable used in the reduced-form approach with dyad specific variation. Although beliefs are not based directly on Π_{dt}, they do not fully discount that newspapers will tend to report differently in different destination countries and type of events. Moreover, the further apart are $\hat{\eta}_1$ and $\hat{\eta}_0$, the stronger are naive beliefs updated based on news. We find that $\hat{\eta}_1 = 0.022$ and $\hat{\eta}_0 = 0.002$, i.e. beliefs update a lot by observing that more than about 1 percent of news is bad news.[33]

Naïve beliefs are updated using a likelihood ratio based on (8):

$$\lambda \left(B_{hdt}, N_{hdt} \right) = \frac{f(B_{hdt}, N_{hdt}|0)}{f(B_{hdt}, N_{hdt}|1)} \tag{10}$$

which depends only on news coverage in a dyad. The probability that a country is dangerous is therefore given by

$$\pi_{hdt} = \Pr(s_{dt} = 1 \mid \Psi_{hdt}, \hat{\theta}_d, \hat{\eta}_{d0}, \eta_{d1}).$$

which also evolves according to the Bayesian recursion:

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

where, as before, $E_{t-1}[\pi_{hdt}]$ is derived from the Markov chain governing the underlying categorical state.[34]

Two features of naive belief formation are worth highlighting. First, such beliefs do not use a model of the news cycle to discount news. For example, after a news

---

[33]This implies a subtle feature of the supposed nature of availability bias where naive beliefs do not take into account that reporting reacts differently to danger depending on the destination country. We have, for example, $\hat{\eta}_0 = 0.006$ for Tunisia but $\hat{\eta}_0 = 0.001$ for Israel.

[34]Appendix Figure A6 reports the posterior distribution of π_{dt} 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_{dt} has a much thicker tail during dangerous months.

event, reporting will spike and then calm down immediately and so will naïve be-liefs. A dyad in which there is no reporting interrupted by some bad news will jump back and forth between strong beliefs of safety and danger. Second, observing *non-violent news* reporting about a country leads to updating that a country is safe. Since $\hat{\eta}_1 > \hat{\eta}_0$, naive 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 (10) as a function of B_{hdt} relative to N_{hdt} at which naive beliefs attach a much larger probability to a country being dangerous. This feature of the model allows it to account for the sharp changes in beliefs in response to news that are needed to explain patterns in the spending data. However, these sharp reactions are not driven by a different model of danger being used in the formation of naïve beliefs but by the use of information when updating these beliefs.

4.2   Fitting the Model

We will estimate weights on naïve beliefs based on country-specific coverage with sophisticated beliefs in explaining the data. We also allow for a lagged effect of up to nine months, estimating the lag structure that gives the best fit to the data. We use the following specification based on (6):

$$y_{hdt} = \alpha_{hd} + \alpha_{ht} + \alpha_{dm} + \xi \sum_{\tau=0}^{-9} \omega_\tau \left(\chi \Pi_{dt-\tau} + (1-\chi) \pi_{hdt-\tau} \right) + \varepsilon_{hct}. \qquad (11)$$

where $\chi \in [0,1]$ reflects the extent to which beliefs are naïve or sophisticated on average. We interpret (11) as modelling the beliefs of a "representative" tourist. But χ could also reflect heterogeneity in the population in terms of naivety or sophistication among sub-groups of tourists.

The model-based approach differs from equations (2) and (3) in three ways. First, the belief estimates, (Π_{dt}, π_{hdt}), are both heterogeneous across countries and non-linear in their responsiveness to violence and news data. Second they depend on the entire history of violence rather one period lagged values and third, they allow a lag structure to reflect the timing of travel decisions. At the same time the weights

(ω, χ) reflect the importance of different kinds of tourist behavior for overall spending.

Table 6 explores how well our estimates of model-based beliefs $\{\Pi_{dt}, \pi_{hdt}\}$ explain variations in aggregated tourist spending. In column (1) of Table 6, we report the relationship between aggregated spending and Π_{dt-1}, i.e. supposing that $\chi = 1$ and $\omega_{-1} = 1$ in equation (11). 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. This magnitude is in the same ballpark as the reduced-form results reported in Table 2. But unlike the estimates based on (2), recall that Π_{dt-1} moves as in Figure 6 and therefore reacts much more strongly to some changes in violence than others. Moreover, effects coming through Π_{dt-1} are also naturally heterogeneous across destination countries.

(Table 6)

Column (2) focuses on responsiveness of spending to naive beliefs, π_{hdt-1}, i.e. imposing $\chi = 0$ and $\omega_{-1} = 1$. We now get a fall of 0.364 log points or 30 percent if the news-based belief that a destination country is dangerous within a dyad moves from 0 to 1. This is true even though most dyads with high values of π_{hdt} also have high values of Π_{dt} so that these two effects are complements with π_{hdt} increasing the intensity of the impact of violence on tourism spending. In column (3) we estimate π_{hdt} based on news reporting on fatalities rather than 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.

An important take-away from columns (1) to (3) is that there is a precisely estimated association between aggregated card spending and our measures of sophisticated and naive beliefs about danger: Π_{dt} and π_{hdt}. This holds when we combine news-based and violence-based beliefs in the same regression in column (4). In other words, our model of beliefs based solely on media coverage captures perceptions that are almost unrelated to the underlying risk displayed in Figure 6. The reason is that aggregate spending seems to follow media coverage even if this coverage is not closely related to underlying risks. We see this as empirical support for the idea of *availability heuristics* as captured by our model.

In column (4) we show the estimate of after fitting the entire model to the spending data. For this we first find the weights that give the best fit to the spending data using a grid search over the weights in equation (11) to maximize goodness of fit.[35] We will use these estimated weights and $\hat{\hat{\zeta}}$ to quantify the separate impact of news on spending in the next section. We find that $\hat{\omega}_0 = \hat{\omega}_1 = 0.2$ so that 20 percent of spending is driven by contemporaneous beliefs and 20 percent are coming from the first lag. After that, the weight based on the best fit falls (the weight sequence is 0.15, 0.1, 0.1, 0.1, 0.05, 0.05, 0.05). We estimate that the weight on sophisticated beliefs is $\hat{\chi} = 0.4$. This suggests that the availability heuristic implicit in news-based belief formation might be a powerful force in understanding the economic impact of new coverage on spending.

Putting this together, we find that "optimal" weighted average reported in column (4) suggests that if all tourists switched their categorical beliefs that a destination is dangerous from zero to one, then spending would fall by about 55 percent (0.75 log points). Of course, this is quite an extreme thought experiment but relevant for destinations in which news coverage and casualty data lead to maintained beliefs of danger over several month as is would be the case in ongoing terror campaigns in Afghanistan or Iraq or as it is the case for most locations in the Covid-19 crisis.

4.3 Information Suppression

One interesting issue raised by these findings is whether suppressing media coverage could affect behavior in the face of terrorist incidents and more widely. Our results suggest that up to sixty percent of the economic effect of terrorism would not arise in the absence of news coverage. The fact that many countries around the world suppress media provides an interesting natural experiment and a further test of the importance of the role of news coverage on spending.

Specifically, we hypothesize that suppression of information through media censorship would affect updating of beliefs π_{hdt}. Hence our estimate of π_{hdt} should have less predictive power in countries without free media. In order to test this, we collected data from reporters without borders (RSF) who score countries according to their press freedom. We use this to split countries depending on whether they

---

[35]See Appendix B for more details.

are above or below the median score in our sample of countries. This permits us to test whether the variable π_{hdt} enters heterogeneously across the free and censored countries.

(Table 7)

The results are presented in table 7. 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 data set from six countries including Russia and China. Also striking is that this amplification of the effect is not found for sophisticated beliefs, Π_{dt}, as shown in column (3). This supports the view that such beliefs are relevant only for those who are able to access broader news sources downweighting news media. Columns (4) to (6) present similar findings when we use the total number of accounts as the dependent variable.

These findings corroborate our view that naive beliefs, π_{hdt}, capture the impact of signals transmitted through local media reporting. The naive behavior we model treats free press coverage as informative and can lead to local "overreactions" due to availability bias. Of course we would not use this finding to argue that suppression of information is normatively desirable even when there is a propensity for overreaction. There are many beneficial aspects of societies with free media that transcend concerns of the kind that we have focused on. And there is also the question of whether there is the possibility that citizens in free media countries are able to learn to become more sophisticated users of news media over time. This would be an interesting topic to investigate in future.

4.4 Simulation of a News Shock

A crucial difference between sophisticated and naïve approaches to belief formation is that the probability that a destination is more likely to be regarded as safe on the basis of naïve beliefs if there is more non-violent (other) news coverage, i.e. $N_{hdt} - B_{hdt}$. Figure 7 illustrates the role of other news by considering an outbreak of violence on tourists (at date 0) which changes sophisticated beliefs from $\Pi_{d,-1} = 0$ to $\Pi_{d,0} = 1$ for one month. This yields the dashed line in both panels of Figure 7. The

immediate effect is that the 20 percent of tourists which react immediately to danger ($\omega_0 = 0.2$) do not travel to the destination and tourism spending falls by around six percentage points. The dashed line shows the persistent effect on tourism spending, i.e. the effect of $\Pi_{d,0} = 1$ falls only slowly because most tourists book their travel in advance.

(Figure 7)

We then show the *additional* impact of news reporting working via a change in naive beliefs in a scenario in which the switch in sophisticated beliefs is accompanied by bad news $B_{hd0} = 1$. To illustrate the importance of other news, $N_{hdt} - B_{hdt}$, we contrast two levels of background reporting $N_{hd0} \in \{0, 100\}$. Panel A shows the effect of bad news on tourism spending with no other news, $N_{hdt} = 0$, is a further eight percentage points in spending. Again, because tourists book their holidays in advance this effect persists. However, with $N_{hd0} = 100$, the news about tourist violence is "drowned out" by other news and the impact of $B_{hd0} = 1$ is reduced dramatically. This kind of scenario is not atypical of the reporting landscape for tourist violence in a country like Israel where there is a lots of other news coverage.

This vividly illustrates the importance of background news in "distracting" tourists or "putting things into perspective" for tourists when they rely exclusively on news coverage to form their beliefs - even when such tourists update using Bayes rule. The behavior of beliefs in our model is explained by the fact that, to the extent that they are naive, then tourists fail to learn from a wider set of news sources. By restricting their updating to domestic news sources, their beliefs are biased in a similar way 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 even without any other background news coverage.

4.5 The Aggregate Economic Impact of Reporting

To quantify the news effect we can contrast the average effect that is coming from sophisticated beliefs,

$$\hat{\tilde{\varsigma}} \sum_{\tau=0}^{-9} \hat{\omega}_\tau \hat{\chi} \hat{\Pi}_{dt-\tau} \tag{12}$$

36

, and the overall effect,

$$\hat{\xi} \sum_{\tau=0}^{-9} \hat{\omega}_\tau \left(\chi \Pi_{dt-\tau} + (1-\chi) \pi_{hdt-\tau} \right) \qquad (13)$$

. We can think of the difference between these two, $\hat{\xi} \sum_{\tau=0}^{-9} \hat{\omega}_\tau (1-\hat{\chi}) \hat{\pi}_{dt-\tau}$, as the effect of news coverage. In Figure 8 we show the effect that we would expect based on sophisticated beliefs using a grey line and the overall effect as a black line.

The left-hand panel in Figure 8 shows the results for Tunisia which shows that a large part of the variability in tourism spending comes from the change in sophisticated beliefs, i.e. the part that is not dyad-specific. However, there is a visible news effect and, in 2015, it alone accounts for a spending decline of about 15 percent. This illustrates the influence that reporting may have had a dramatic impact on tourism for Tunisia in that year. In the right-hand Panel we show the same data for Egypt. Again, we get some news effect but it is nowhere nearly as large as the main effect from sophisticated tourists. This is a dramatic illustration of the crowding out effect that we illustrated in Figure 7.

(Figure 8)

The estimates in Figure 8 suggest material losses to the economy in all four countries that we study. The World Bank reports that tourism revenues 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 in this section indicate losses between 2011 and 2016 of over 35 billion USD due to violence with in excess of 10 billion USD being due to negative news reporting.[36] That said, Egypt and Tunisia are predicted to have recovered the losses towards the end of the sample period.

Figure 9 explores how well the model-based approach fits the data. The left-hand panel illustrates the predicted effect, averaged over all origin countries, for Tunisia based on (11) and compares it to the spending residuals. The model captures both the early decline and recovery at the beginning of the Arab spring. However, the

---

[36]For calculations see the Appendix C.

most striking observation is for 2015 where it accurately captures both the decline and recovery in spending.

(Figure 9)

The model-based prediction compares favorably with that from the reduced-form model as shown in the right-hand panel of Figure 9. For this we use the estimated effects, $\xi_1 n_{hdt-1} + \xi_2 v_{hct-1}$ averaged over origin countries based on Table 3. Although some of the broad patterns are visible, the fit to the data is notably inferior to that in the left-hand panel, failing to reproduce the timing and magnitude of the fluctuations in spending. This is particularly noticeable during the Arab spring and getting the timing right in sudden expenditure declines in 2015. Thus Figure 9 supports the utility of an approach which models patterns of belief formation alongside a reduced-form approach. Moreover, we it is an interesting insight that a behavioral model where beliefs are at least partially naive, using an availability heuristic, can provide such a good account of the aggregate spending patterns in the data.

5   Conclusions

The core finding of the paper is that there is a robust negative relationship between violence, news coverage and tourism spending based on a sample of five destination and fifty-seven origin countries. The quantitative magnitude of this relationship is established by leveraging novel data sources on both aggregated spending from credit card data and news coverage. We show that the effect of violence and news coverage varies markedly across both destination and origin countries. In an attempt to explain this, we have fitted a specific model of belief formation to study patterns in the data. In the best fit to the data, we find that more than half of the weight in terms of updating beliefs comes from country-specific news reporting but there is significant heterogeneity across countries and events.

The paper contributes to an emerging literature which explores the power of the media in influencing economic decisions. Citizens rely heavily on the media as a source of information when deciding where and when to travel. Our supervised machine learning method of identifying specific negative news linked to violent events

allowed us to study this. We find that there is significant country-level heterogeneity in responses to violence are attributable to negative news reporting. In particular, we find that sudden slumps in spending with gradual recovery are best explained by a media effect. The media is indeed powerful in affecting international flows of tourists.

It is often claimed that perceptions matter for more than tourism. Negative coverage of the prospects for African countries create a climate of opinion among corporate boards and shareholders which 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 and for a region which is economically not well integrated (Rouis and Tabor, 2012). How far news biases have real aggregate economic consequences has yet to be studied but is ripe for investigation.

It is tempting to see our main results as an endorsement of the restriction of press freedom in some circumstances. When countries want to overcome the economic effects of a continuous death risk like a Covid-19, for example, it is tempting to suggest that media reports of the case load should be suppressed. However, as our results on press freedom show such behavior could have broader repercussions in which local economic agents do not form their opinion with the local press any more.

The basic ideas and methods developed here can be extended to a number of additional contexts. We have cautiously identified the effects of violence and its reporting *within* a set of five countries. But this leaves open the question of whether tourists divert their spending to other countries or even choose to stay at home. So there can be gainers and losers from violence. To study this would require aggregated spending data from a broader range of destination countries. We have also looked at quite extreme violent events because their news effect is easier to identify. It would be interesting to look at other forms of crime perpetrated on tourists. This too could have an impact on tourism via a news reporting channel. For this, destinations such as the Caribbean and Latin America would be particularly interesting.

Finally, our paper contributes to wider debates about forces that could lead to a reversal in globalization. Given that increased international travel has been a significant component of global integration, determinants of tourism are important not just in terms of generating traditional economic gains from trade but also in creating greater

cross-cultural understanding. To the extent that security concerns increase the perceived costs of travel, politics may therefore have an important impact on this aspect of international integration. And the results in this paper suggest that how the media chooses to report these risks has a role to play in this process.

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.

Adams, W. C. (1986). Whose Lives Count? TV Coverage of Natural Disasters. *Journal of Communication 36*(2), 113–122.

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.

Allcott, H. and M. Gentzkow (2017). Social Media and Fake News in the 2016 Election. *Journal of Economic Perspectives 31*(2), 211–236.

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

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

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

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

Barrera, O., S. Guriev, E. Henry, and E. Zhuravskaya (2020). Facts, alternative facts, and fact checking in times of post-truth politics. *Journal of Public Economics 182*, 104123.

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. and R. Burgess (2002). The Political Economy of Government Responsive-ness: 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.

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

Bordalo, P., N. Gennaioli, Y. Ma, and A. Shleifer (2018). Over-reaction in Macroeco-nomic 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: Ap-plied Economics 10*(4), 246–282.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Gentzkow, M. and J. M. Shapiro (2011). Ideological Segregation Online and Offline. *The Quarterly Journal of Economics 126*(4), 1799–1839.

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). Terrorism and the Media : The Effect of US Television Coverage on Al-Qaeda Attacks. *IZA Discussion Paper* (10708).

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

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

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

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

Mueller, H. (2012). Growth Dynamics: The Myth of Economic Recovery: Comment. *American Economic Review 102*(2008), 3774–3777.

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.

Mullainathan, S. and A. Shleifer (2005). The Market for News. *American Economic Review 95*(4), 1031–1053.

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.

Singer, E. and P. M. Endreny (1993). *Reporting on Risk: How the Mass Media Portray Accidents, Diseases, Other Hazards*. Russell Sage Foundation.

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

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

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

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

Tversky, A. and D. Kahneman (1974). Judgment under Uncertainty: Heuristics and Biases. *Science 185*(4157), 1124–1131.

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

Zhu, J., A. N. Sanborn, and N. Chater (2018). The Bayesian Sampler: Generic Bayesian Inference Causes Incoherence in Human Probability Judgments. *PsyArXiv*, 1–26.

Zhuravskaya, E., M. Petrova, and R. Enikolopov (2020). Political Effects of the Internet and Social Media. *forthcoming, Annual Review of Economics*, 1–32.

Figure 1: Overall aggregated spending patterns of British- and German-issued cards in Tunisia in the wake of the Sousse attack)
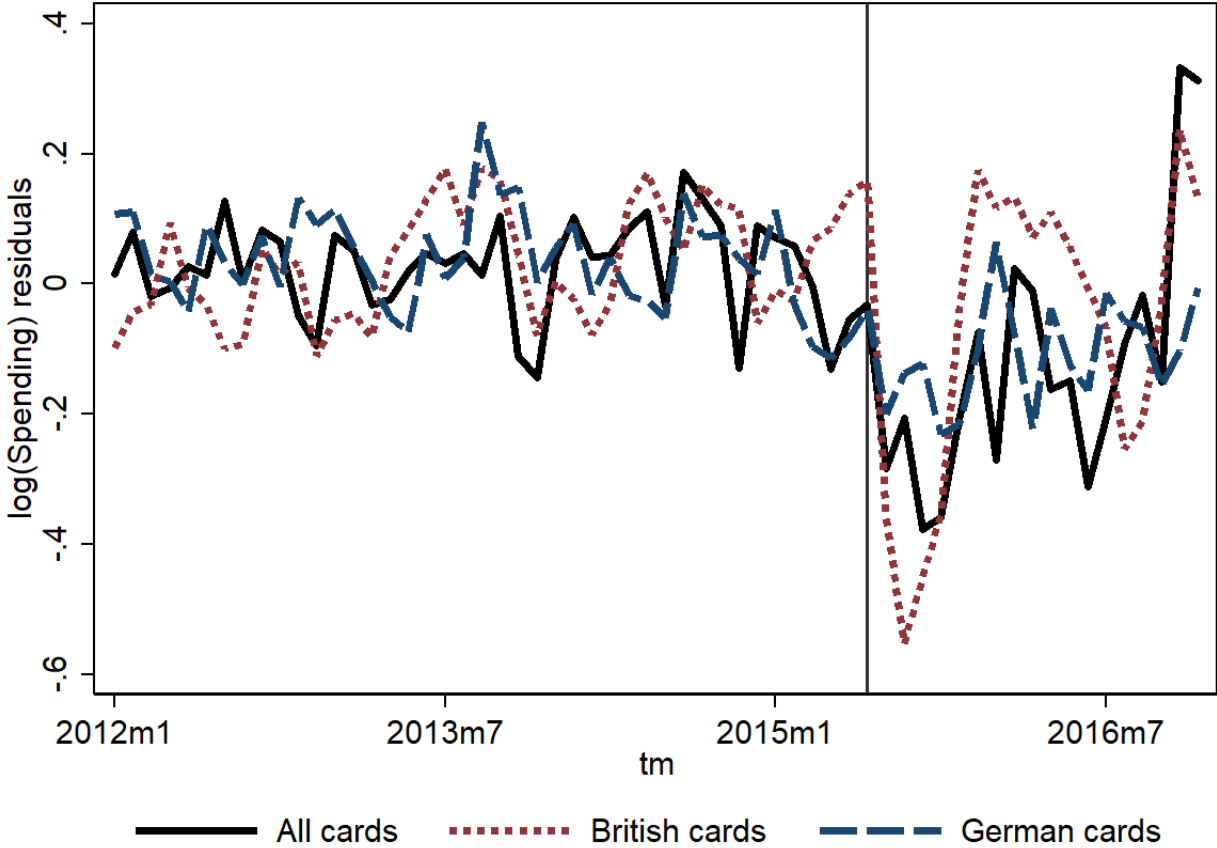
Notes: The solid black line presents residuals of a regression of the log of aggregated card spending removing destination country fixed effects as well as destination-specific seasonality. The other two lines plot residuals of a regression of the log card spend for German- and British-issued cards in Tunisia over time, having removed dyad fixed effects, issuing country by time fixed effects and destination by month fixed effects. The drop in tourism spending is markedly larger for British-issued cards.

Figure 2: Map of countries included in our estimation samples across the exercises
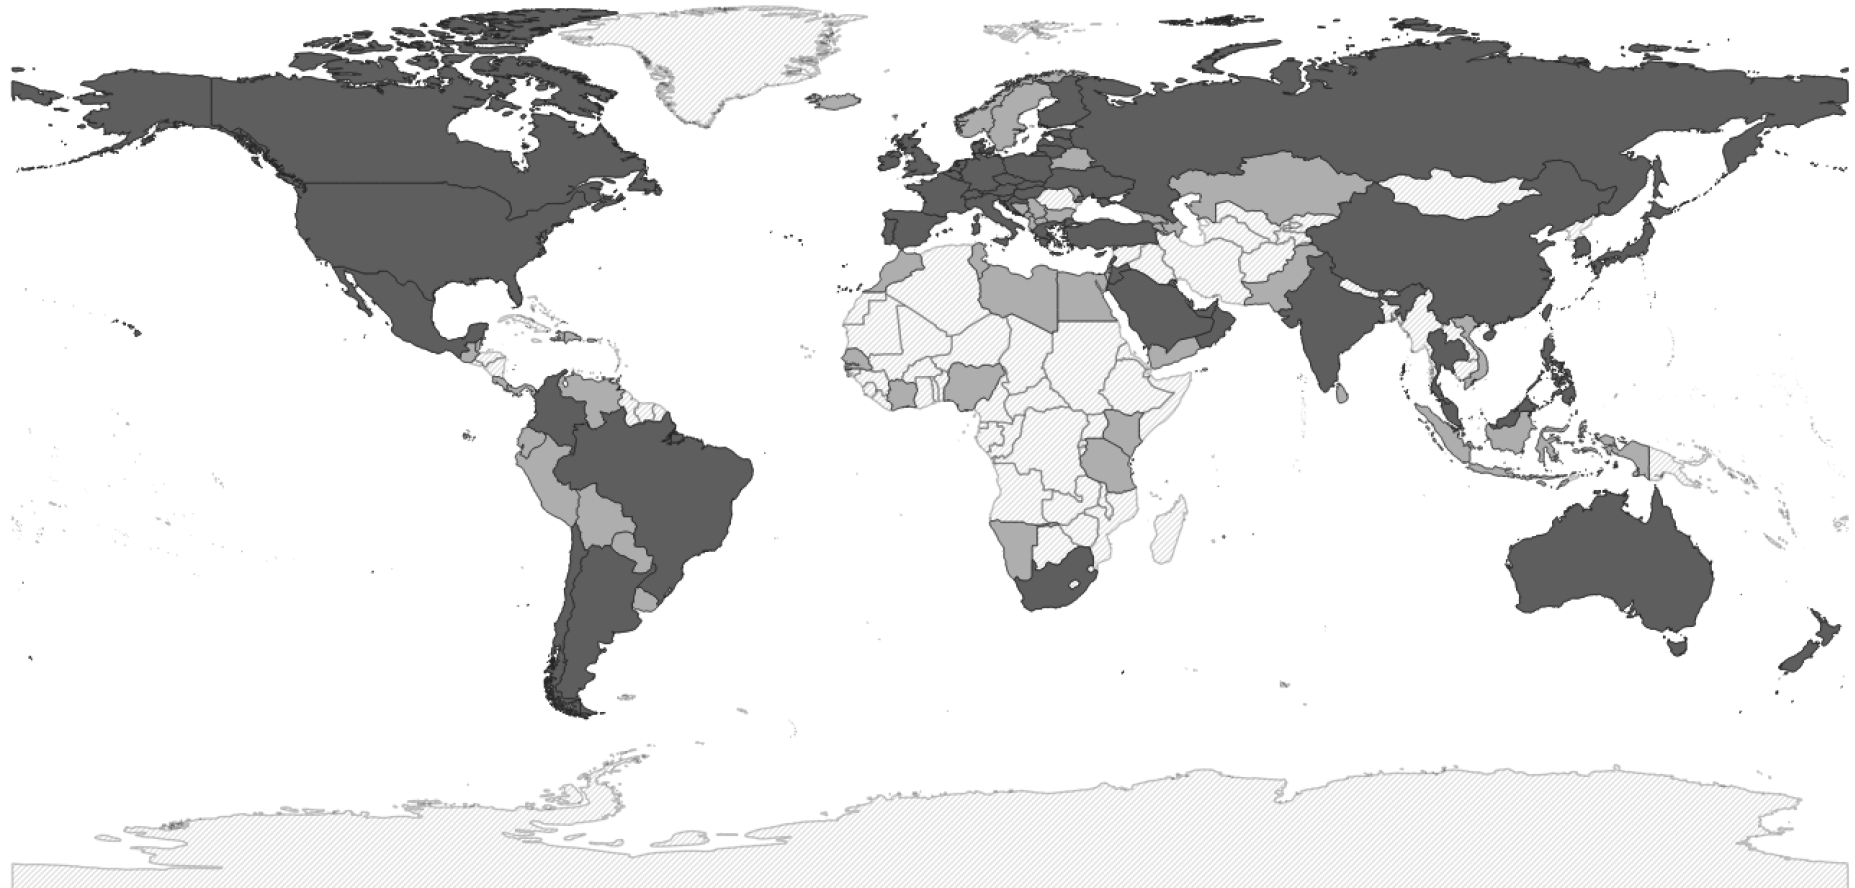
Notes: Figure indicates the origin countries included in our estimating sample. Dark-shaded are countries for which both newspaper and aggregated spending data is available, lightly shaded areas are countries for which only aggregated spending data is available.

Figure 3: 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
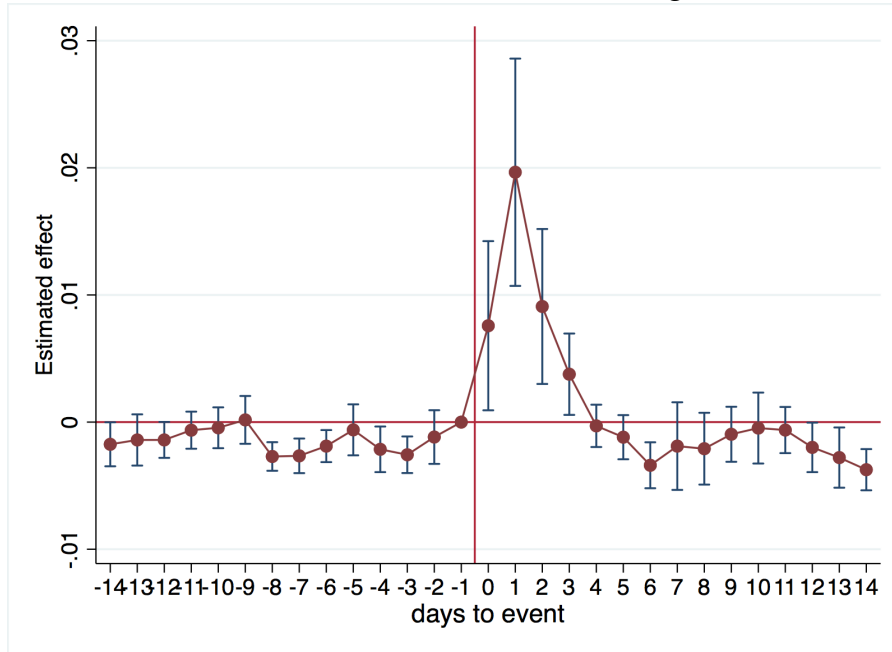
Panel B: Share of articles indicating tourist targeted
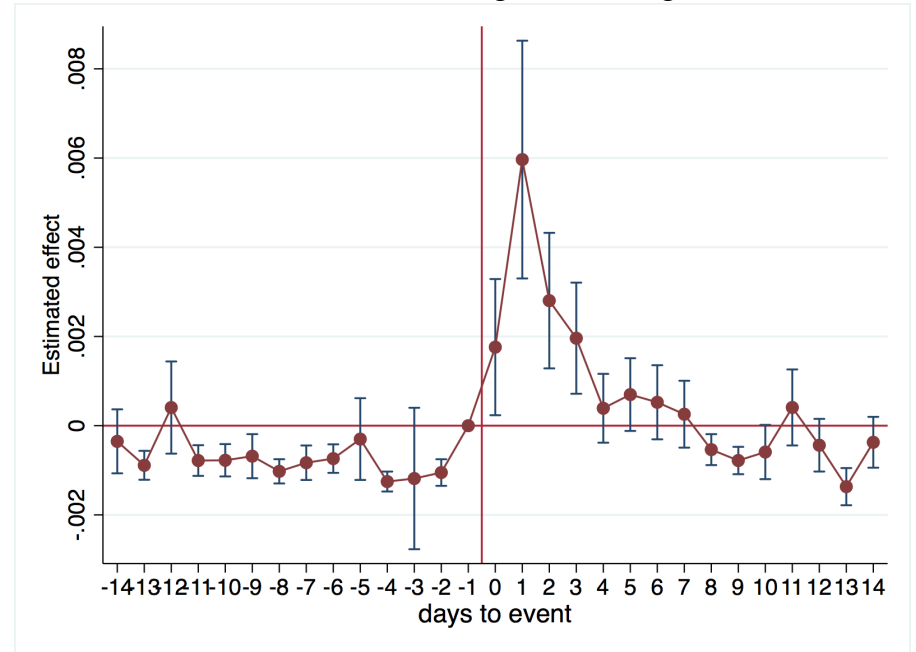
Notes: Figure plots point estimates from a regression that absorbs event, reporting dyad and day fixed effects. The dependent variable in Panel A measures the share of articles on a day and dyad that are classified as reporting on violence that involved fatalities. The dependent variable in Panel B measuring 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 4: Monthly share of articles classified as reporting on violent events across four destination countries averaged across the set of 57 tourist origin countries for which news reporting data is available
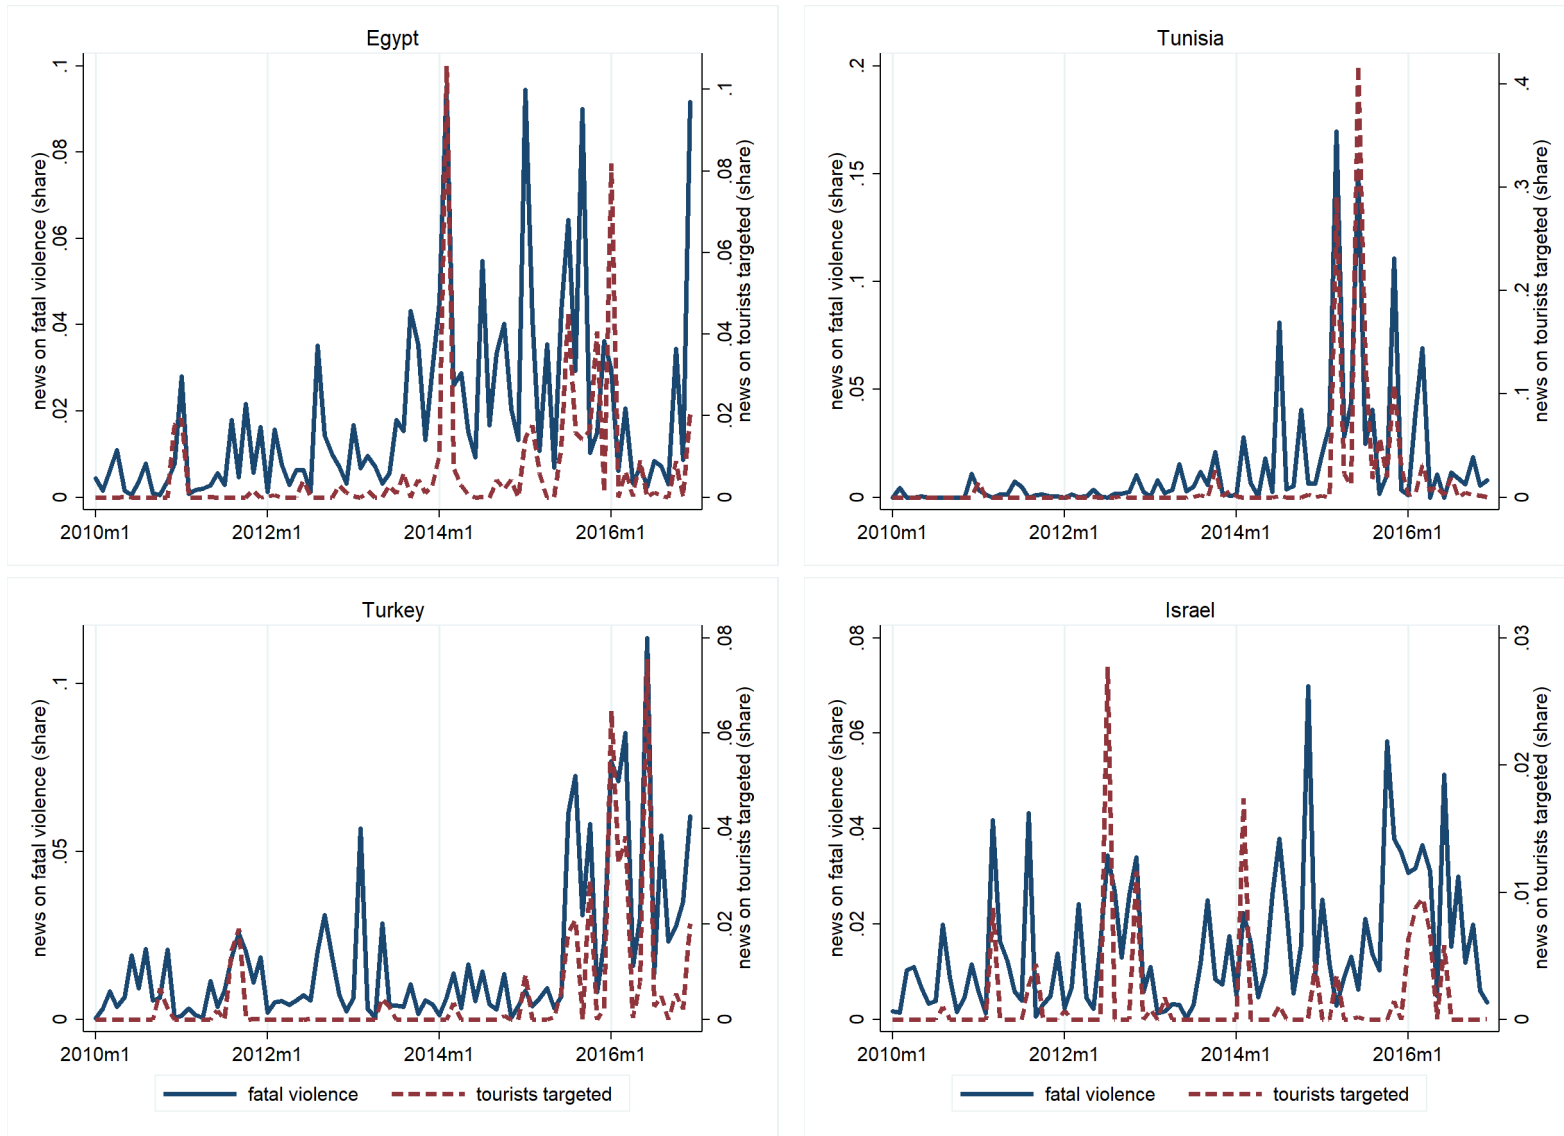
Notes: Figure plots the average reporting on violence across the 57 tourist origin countries for which data is available covering four main destination in our sample. The figure provide for each of destination the average share of news articles across the tourist-origin countries that are classified as reporting on violent events (left axis) or that are classified as covering violence against tourists (right axis). Note that the data was aggregated from daily- to monthly level as the this coincides with the temporal resolution at which aggregated and anonymized credit card data has been made available.

Figure 5: Lead- and lagged effect of violence targeted against tourist's on tourism spending

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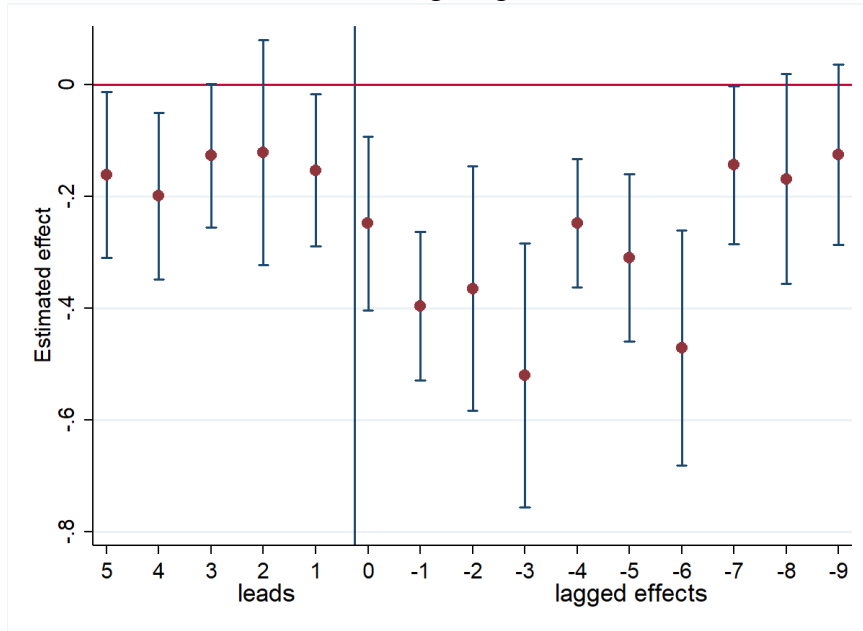
Panel B: News on any fatal violence
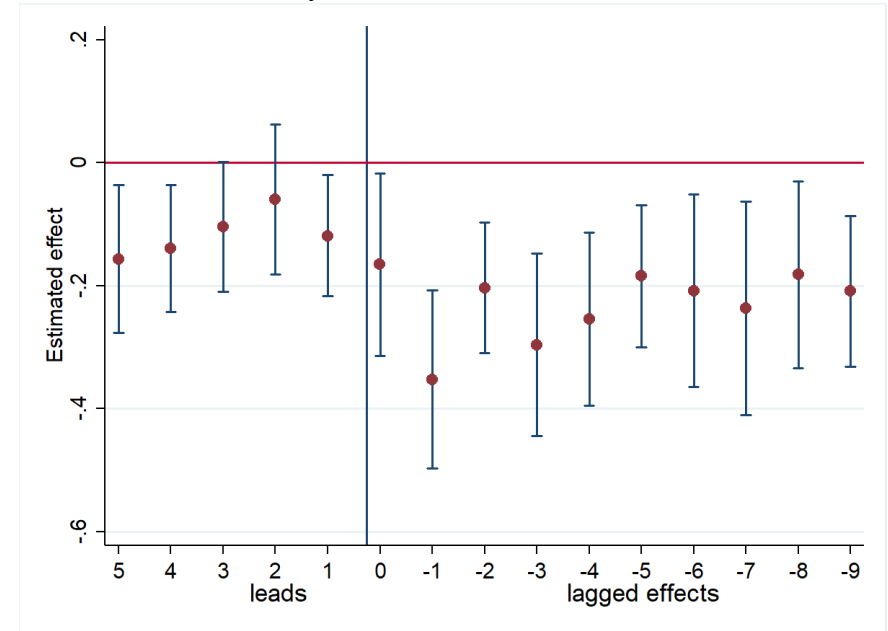
Notes: Figure plots the results from estimating a linear regression controlling for dyad fixed effect, issuing-country by time fixed effect and destination by month seasonality. The coefficients that are reported are the estimates on different leads and lags of the violence reporting measure on the log-value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

Figure 6: Markov Chain Fitted Probability of Danger Across Sample Countries
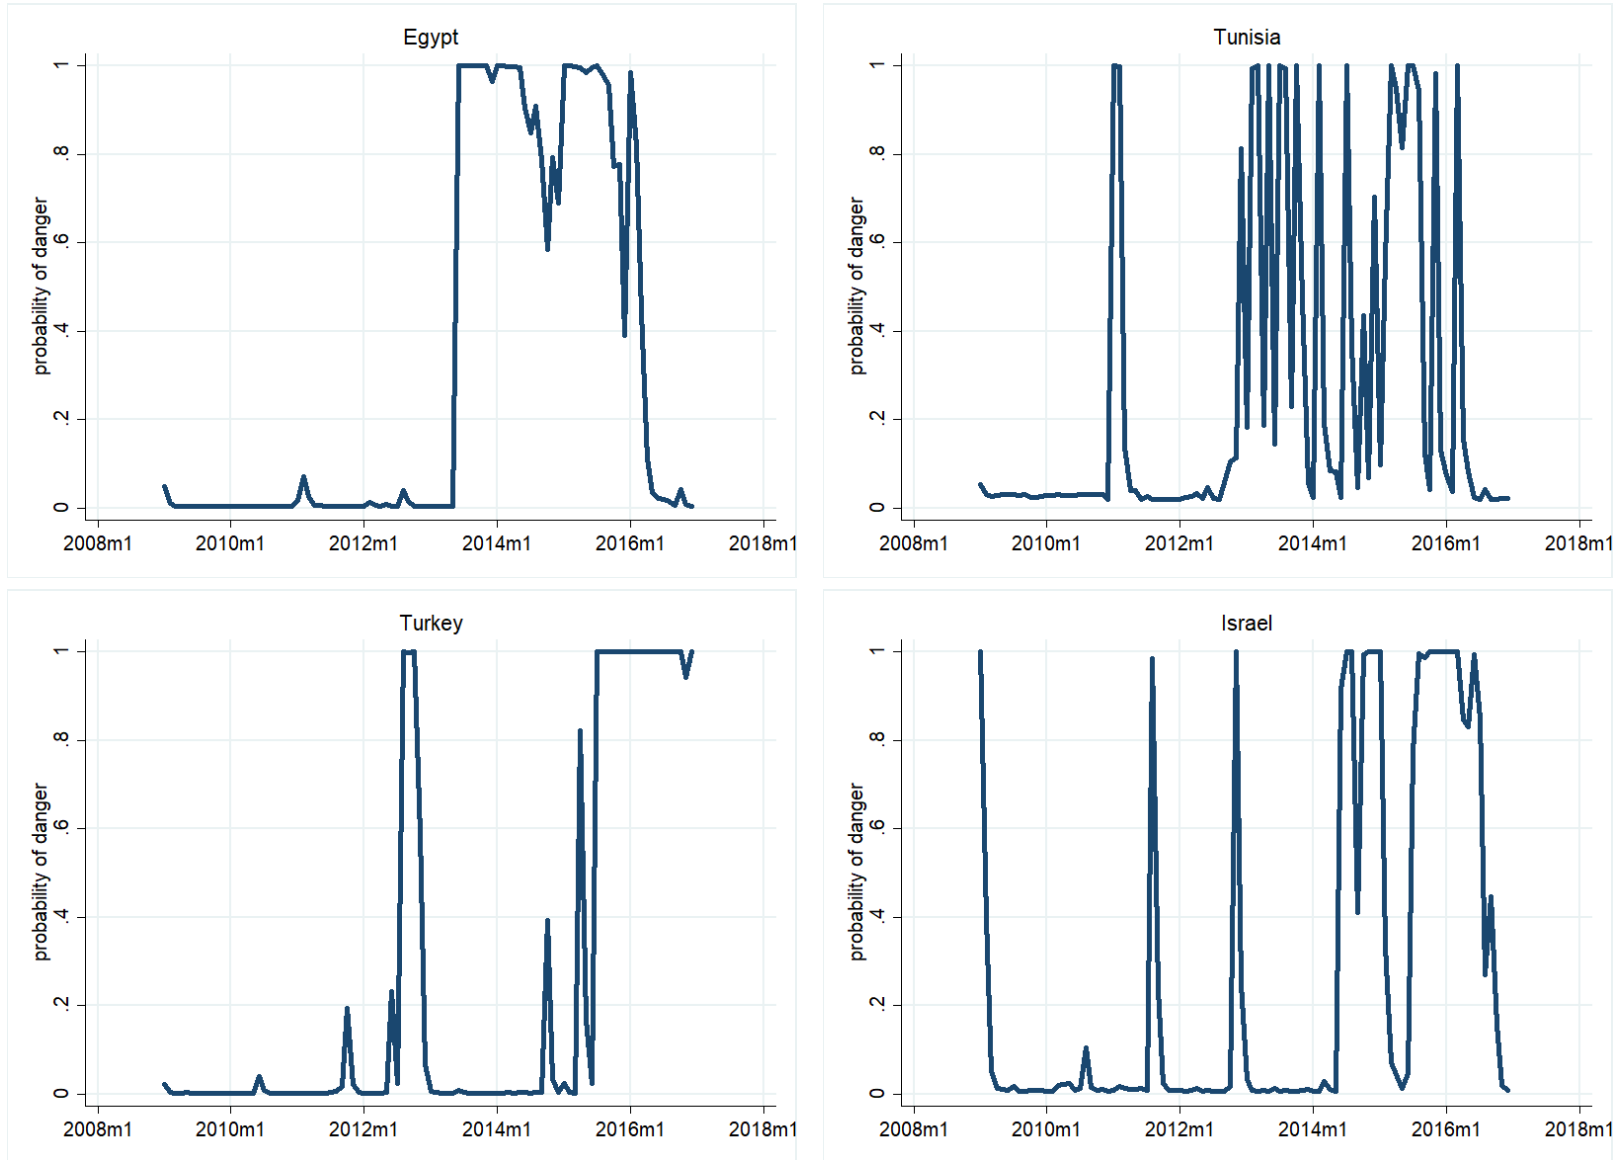
Notes: Figures plot out the probability of danger Π_{dt} as inferred by "sophisticated" tourists from the time-series data on violent events. The figure plots the fitted Markov chain estimates across the four main tourist destinations in our data. The Markov switching model and model fitting exercise is described in more detail in the main text in Section 4.

Figure 7: Simulated effect of violent news shock triggered by a violent event and moderating impact of other news coverage
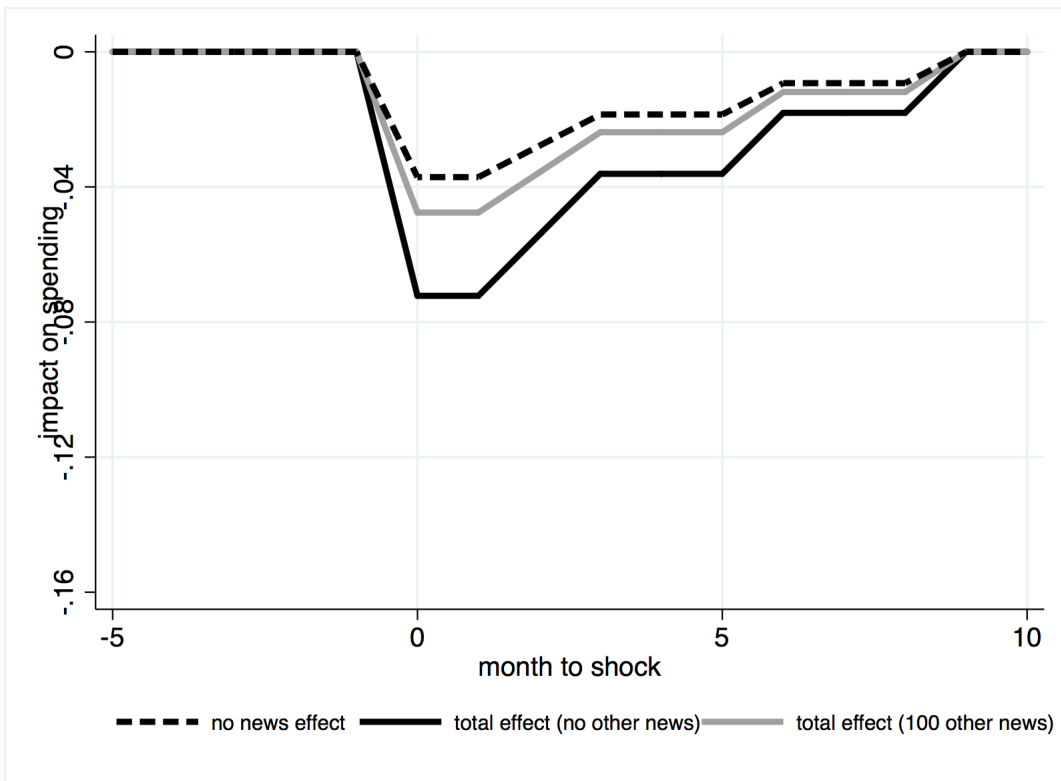
Notes: The figure plots out the model-implied impact of a violent event occurring at month $t = -1$ on tourism spending across subsequent months under different scenarios concerning the news environment. The effect of the violent event itself without any news coverage around it is provided as the dashed line. This is a reference point capturing the impact that is attributable to "sophisticated" tourists updating their beliefs. The other lines present the total impact on tourism spending that incorporates the independent effect that media coverage has. This effect manifests itself through the impact the reporting shock has on "naive" tourists beliefs. We contrast two information environments: one where there are, on average, 100 other notable news items on a dyad and month covering news items 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 or drown out the impact that the news shock has on "naive" tourists beliefs.

Figure 8: Model-based effects in the case of Tunisia and Egypt

Panel A: Tunisia
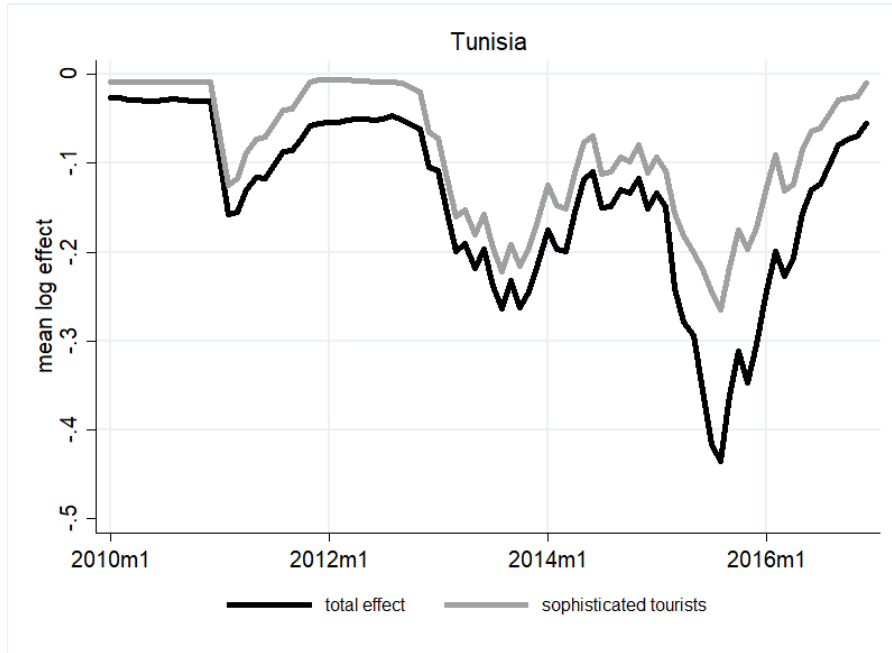
Panel B: Egypt
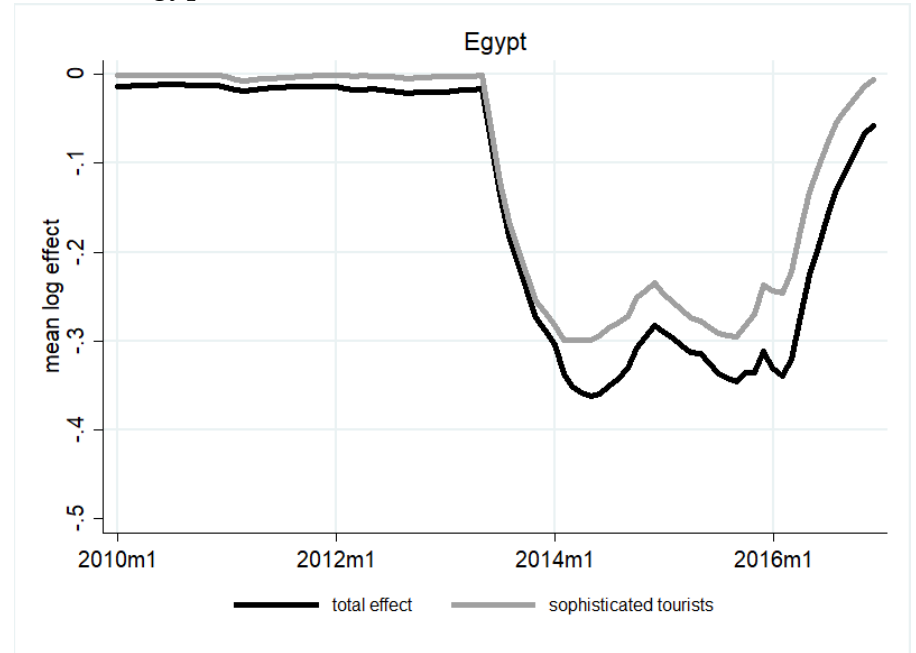
Notes: Figure plots the overall model-based impact of violent news shocks on tourism spending for Tunisia (left) and Egypt (right). The solid black line presents the total effect while the grey line presents the effect that would be implied by the sophisticated tourists.

Figure 9: Comparison of model- and reduced-form implied effects in capturing variation in tourism spending in Tunisia over time

Panel A: Model based
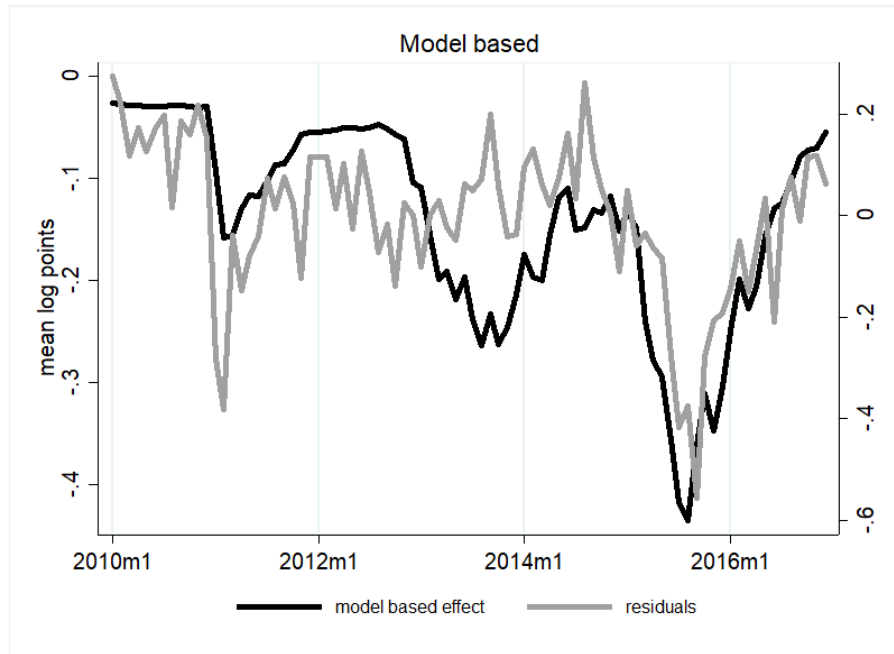
Panel B: Reduced form
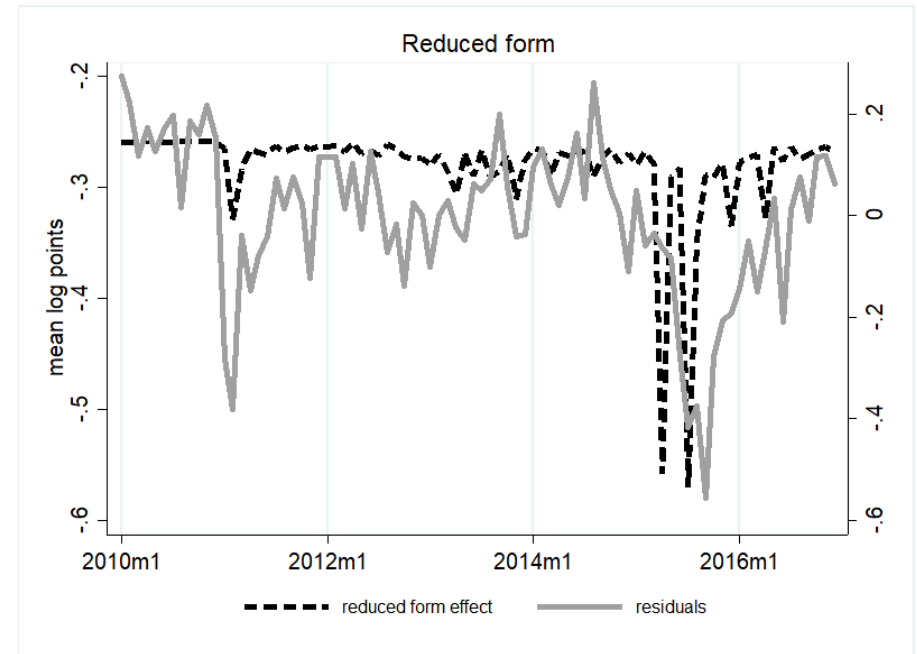
Notes: The figure illustrates the model fit vis-a-vis the fit that would be implied in reduced form exercises commonly studied using the case of Tunisia as a destination country. The left figure plots the mean of the model-based estimated effect of violent events and their ensuing media coverage on tourism spending across the 57 tourist origin countries for which news reporting data is available. The figure also plots the average values of the residualized data after removing dyad-, origin-by-time and destination by month fixed effects for comparison. The right figure plots these same residuals against the implied average reduced effects across the 57 tourist origin countries for which media reporting is available. This represents the effect that would be captured in a purely reduced-form regression exercise as e.g. studied in Table 3.

Table 1: 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 2: Effect of Country-level Violence measured by different event data sets on tourism spending and active cards

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

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

Table 3: Reduced Form Results: Impact of News Reporting and Tourism activity

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

Notes: Table presents regression capturing reduced form effect of dyadic (tourist-origin by destination) specific news coverage on the dyadic log values of card spend in columns (1)-(3) and the number of cards in a month in columns (4)-(6). Panel A uses as news measure the share of articles in a month on a dyad that is classified as capturing tourists being targeted by violent events. In Panel B the news measure captures the share of news in a month on a dyad that is classified as covering violent events. Explanatory variables are lagged by one month to account for the lagged response of tourism to events and news. Event controls are the first four principal components combining the main country-level violence data series studied in Table 2 and Appendix Table A3. Robust standard errors clustered at the dyad level in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table 4: Reduced form Results: Relative Nature of Variation in Reporting Measure

	(1)	(2)	(3)	(4)
		log(Spending)		

Panel A: News on tourist being targeted

	(1)	(2)	(3)	(4)
News on tourists targeted (count of articles)	-0.007**	-0.006**		
	(0.003)	(0.002)		
News on tourists targeted (share of all articles)	-0.512***	-0.494***		
	(0.095)	(0.101)		
News on tourists targeted (share of all articles) - first quartile			-0.005	0.051
			(0.051)	(0.048)
News on tourists targeted (share of all articles) - second quartile			-0.087*	-0.051
			(0.049)	(0.047)
News on tourists targeted (share of all articles) - third quartile			-0.160***	-0.123***
			(0.041)	(0.042)
News on tourists targeted (share of all articles) - fourth quartile			-0.304***	-0.297***
			(0.079)	(0.081)
Observations	23859	23859	23859	23859
R2	.966	.967	.966	.967

Panel B: News on any fatal violence

	(1)	(2)	(3)	(4)
News on violence with fatalities (count of articles)	-0.003	0.001		
	(0.002)	(0.003)		
News on violence with fatalities (share of all articles)	-0.422***	-0.351***		
	(0.097)	(0.102)		
News on violence with fatalities (share of all articles) - first quartile			-0.065	-0.019
			(0.047)	(0.044)
News on violence with fatalities (share of all articles) - second quartile			-0.063**	-0.018
			(0.025)	(0.025)
News on violence with fatalities (share of all articles) - third quartile			-0.085***	-0.035
			(0.024)	(0.025)
News on violence with fatalities (share of all articles) - fourth quartile			-0.113***	-0.069**
			(0.032)	(0.032)
Observations	23859	23859	23859	23859
R2	.966	.967	.966	.966
Dyad FE	YES	YES	YES	YES
Origin/Time FE	YES	YES	YES	YES
Dest./Month FE	YES	YES	YES	YES
Event controls	NO	YES	NO	YES

Notes: Table presents regression capturing reduced form effect of dyadic (tourist-origin by destination) specific news coverage on the dyadic log values of card spend. The table illustrates that a relative measure of news coverage is adequate in capturing the underlying relationship in the reduced form. Panel A uses as news measure the share of articles in a month on a dyad that is classified as capturing tourists being targeted by violent events. In Panel B the news measure captures the share of news in a month on a dyad that is classified as covering violent events. Explanatory variables are lagged by one month to account for the lagged response of tourism to events and news. Event controls are the first four principal components combining the main country-level violence data series studied in Table 2 and Appendix Table A3. Robust standard errors clustered at the dyad level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 5: Instrumental Variable Regression of Reduced Form Effect: Instrumenting news coverage intensity with known casualty-distribution for select events

	(1)	(2)	(3)	(4)	(5)	(6)
	News on tourists targeted			News on any fatal violence		
Panel A: first stage with reporting spillovers						
Any tourist killed	0.181***	0.191***		0.063***	0.069***	
	(0.050)	(0.051)		(0.023)	(0.023)	
Contiguous country x Any Casualties		0.149***	0.147***		0.065**	0.064**
		(0.038)	(0.038)		(0.029)	(0.029)
Same region x Any Casualties		0.059**	0.058**		0.034*	0.033*
		(0.024)	(0.024)		(0.019)	(0.019)
Common language x Any Casualties		0.034**	0.033**		0.030***	0.030***
		(0.015)	(0.015)		(0.010)	(0.010)
R2	0.260	0.303	0.291	0.307	0.317	0.315
Panel B: Second stage: tourism spending						
News measure (share of articles)	-0.811	-0.790***	-0.787**	-3.032	-1.492**	-1.350*
	(0.601)	(0.291)	(0.333)	(2.880)	(0.723)	(0.738)
R2	0.967	0.967	0.967	0.963	0.966	0.966
Weak IV	11.265	15.611	15.331	3.121	7.792	8.516
Panel C: Second stage: number of active cards						
News measure (share of articles)	-1.188*	-1.109***	-1.088***	-4.443	-2.327***	-2.133***
	(0.663)	(0.297)	(0.323)	(3.989)	(0.771)	(0.758)
R2	.971	.971	.971	.961	.969	.969
Weak IV	11.3	15.6	15.3	3.12	7.79	8.52
Observations	23859	23859	23859	23859	23859	23859
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	YES	YES	YES	YES	YES	YES

Notes: Tables presents first-stage (Panel A) and second stages (Panel B and C) of an instrumental variable regression that attributes variation in dyad-specific violent news coverage on tourism activity measured as the log value of tourism spending in Panel B and the log value of the number of active credit cards in Panel C. Columns (1) - (3) use as news measure the share of articles in a month on a dyad that is classified as capturing tourists being targeted by violent events. Columns (4) - (6) use as news measure the share of news in a month on a dyad that is classified as covering violent events. The instrument exploited in column (1) and (4) 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) and (5) 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 a German casualty in a tourist destination and moth would also count as a casualty for Austria and Switzerland. Columns (3) and (6) highlight that the IV is also estimable if we only exploit these spillovers: that means we identify the impact on tourism activity of a German casualty only through its impact on media reporting in Austria and Switzerland. 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 2 and Appendix Table A3. Robust standard errors clustered at destination/month level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

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

Notes: Table presents results from a regression explaining variation in tourist activity measured either as the log value of card spending in columns (1)-(5) and the log number of active cards in columns (6) - (10) on a dyad over time with the probability of a country being in the latent state of being "dangerous". Columns (1) and (6) explore the relationship between the dependent variables the probability of a country being "dangerous" as inferred by "sophisticated 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 "naive tourists" π_{hdt} where the beliefs are either learned through the news reporting on violence targeted against tourists (columns 2 and 7) or through the general news reporting on any violent events with fatalities (columns 3 and 8). Column (5) and (10) explore the weighted average of the two. Please refer to section 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 7: Role of Freedom of Press in Shaping Impact

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

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

Online Appendix

"Terror and Tourism: The Economic Consequences of Media Coverage"

For Online Publication

May 1, 2020

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

A   Further results and robustness checks

A.1   Event Regression Evidence for News

Complementary to Figure 3 we run the following specification

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

where we have defined a dummy variable $Post_{k,t} = 1$ for $\tau = 0, 1, 2$ for up to two days following a violent event. Estimating equation (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 American's are among the casualties and attacks involving tourists.

We present results from specification (14) in Table A2. In columns (1) through (5) the dependent variable is our news coverage covering the share of articles on a day is classified as indicating violence with fatalities. In columns (6) to (10), the dependent

1

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.2   Event Study Evidence for Spending

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

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

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

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

where, as above, y_{khdt} is the log of tourism spending in an event-month k from home country h in country d at date t. This specification includes event fixed effects α_k,

dyad fixed effects α_{hd} and issuing country by time effects α_{ht}. As before, we adjust standard errors two way at the level of the dyad and event.

Estimating (15) permits us to trace out the patterns of aggregated spending around an event month. The results are depicted in Figure A4 for both log of spending and log of active card accounts. In both cases, there is no evidence of any anticipation of the event. Moreover, the observed pattern suggest a sharp contraction in card spending and the number of active cards with effects manifesting a month after an event occurred. Moreover, this occurs with a one month lag as in core specification. That said, it is clear that recovery from an event is quite slow.

In Appendix Table A4 we show that results are robust to dropping each country in turn, highlighting that the results are not an artefact of any of the five destination countries in our sample.

B Grid Search

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

$$y_{hdt} = \alpha_{hd} + \alpha_{ht} + \alpha_{dm} + \xi \sum_{\tau=0}^{9} \omega_\tau \left(\chi \Pi_{dt-\tau} + (1-\chi) \pi_{hdt-\tau} \right) + \varepsilon_{hct}. \tag{16}$$

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

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

which we then use as a regressor in equation 16. We pick the parameter values that yield the highest within R-squared. From this it should already be clear that assuming two different sets of weights on $\Pi_{dt-\tau}$ and $\pi_{hdt-\tau}$ would lead to an explosion of the complexity of the grid search. We therefore focus on one set of weights.

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

0.2, 0.2, 0.15, 0.1, 0.1, 0.1, 0.05, 0.05, 0.05.

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

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

C   Calculations of Total Loss

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

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

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

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

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

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

Figure A1: Validation of aggregated spending data as a proxy for tourist arrivals: comparing subsets of data from the UN World Tourism Organisation

Panel A: Egypt
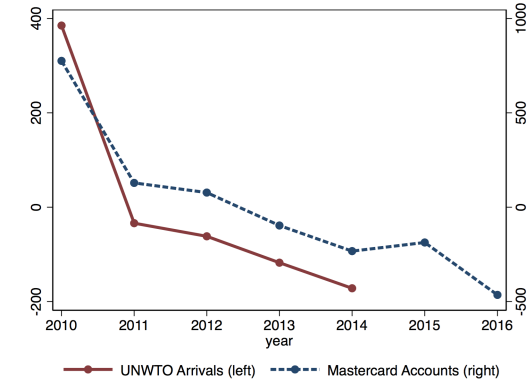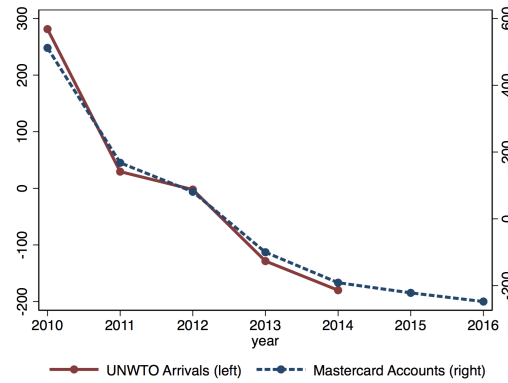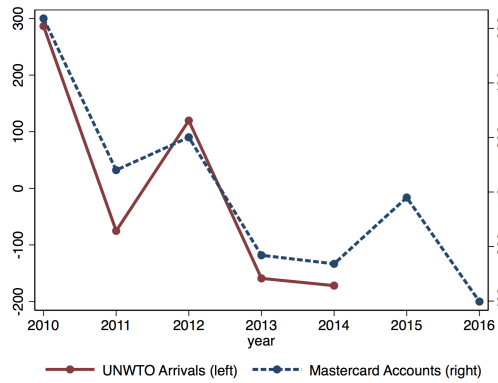
German   French   British

Panel B: Tunisia
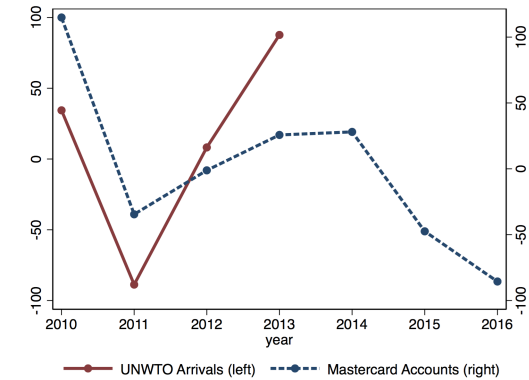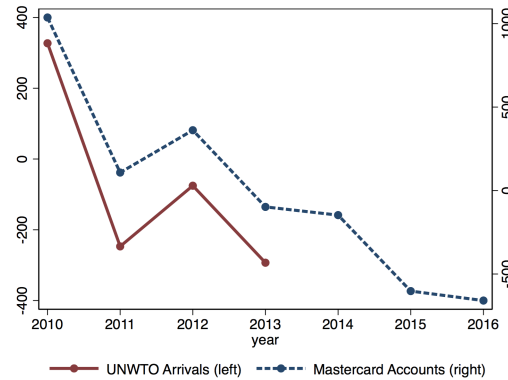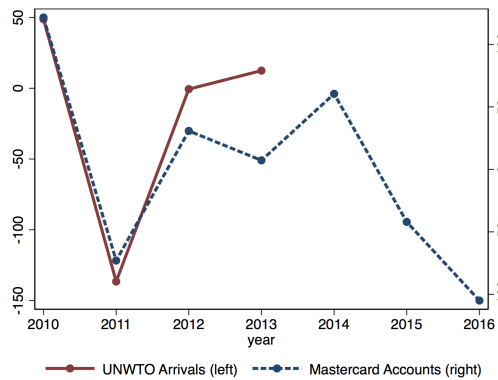
German   French   British

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.

5

Figure A2: Distribution of Share of Articles on Fatal Violence and Violence Against Tourists
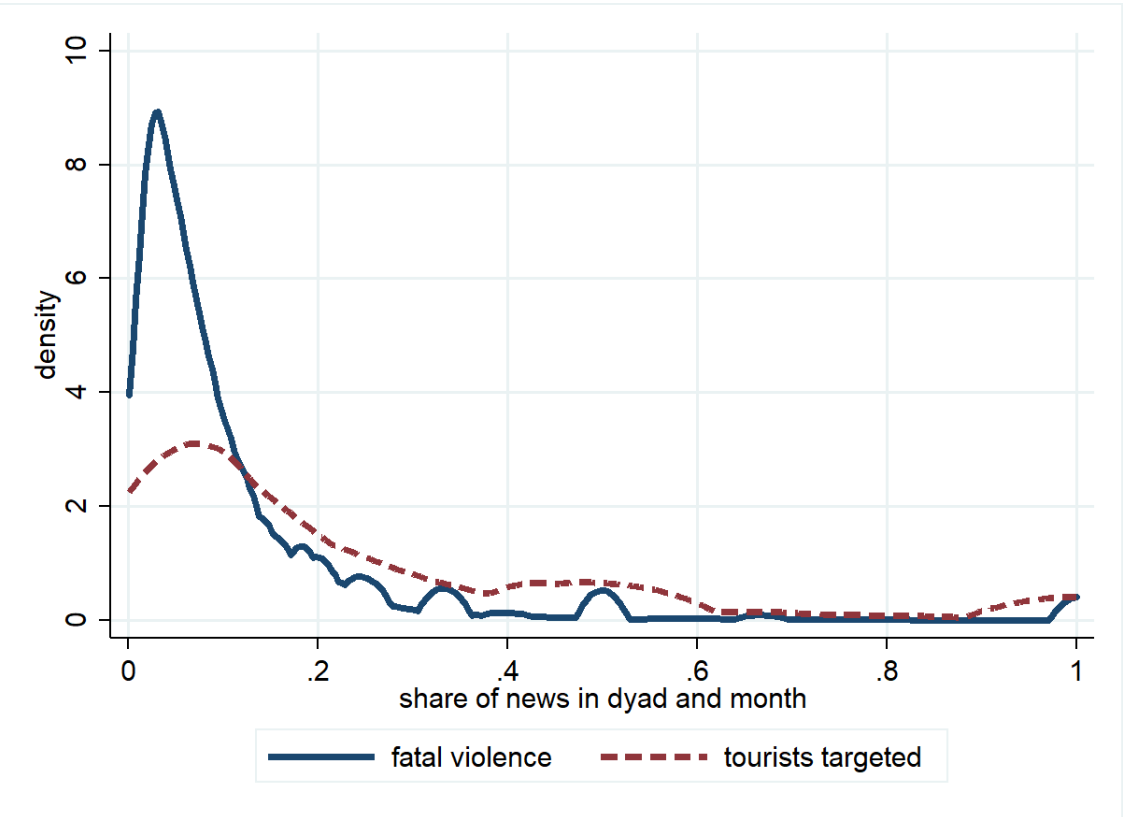
Notes: Figure plots kernel density plotting the distribution of share of newspaper reporting on (any) fatal violence or on violence directed towards tourists.

Figure A3: GTD Events and Reporting Activity: Noisy level effect on number of articles around violent events
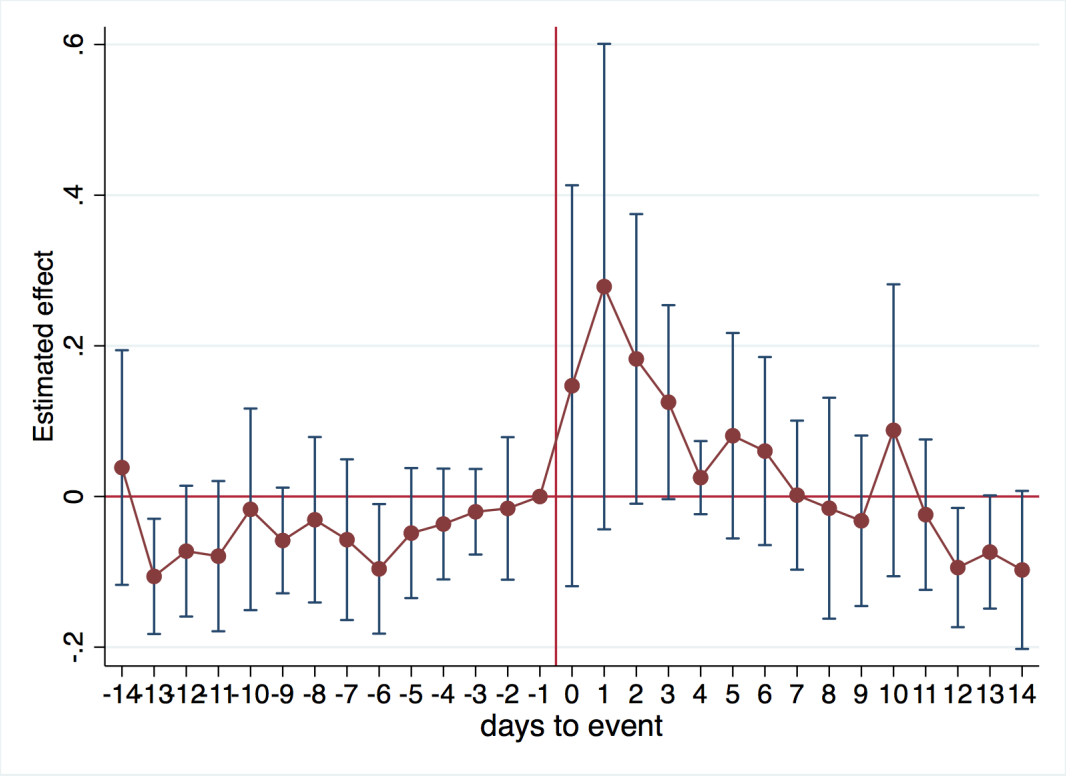
Notes: Figure plots point estimates from a regression that absorbs event, reporting dyad and day fixed effects. The plotted point estimates capture the timing of reporting on a dyad specific to the timing of an individual event recorded in the GTD dataset. 90% confidence intervals obtained from clustering standard errors at the dyad level are indicated.

Figure A4: Event study evidence of the average effect of violent events on tourist activity

Panel A: log(Spending)
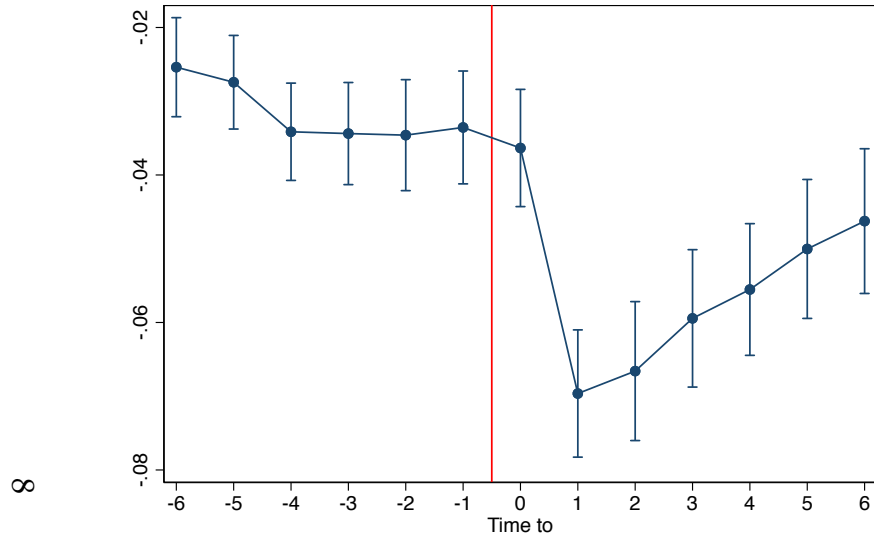
Panel B: log(Number of Cards)
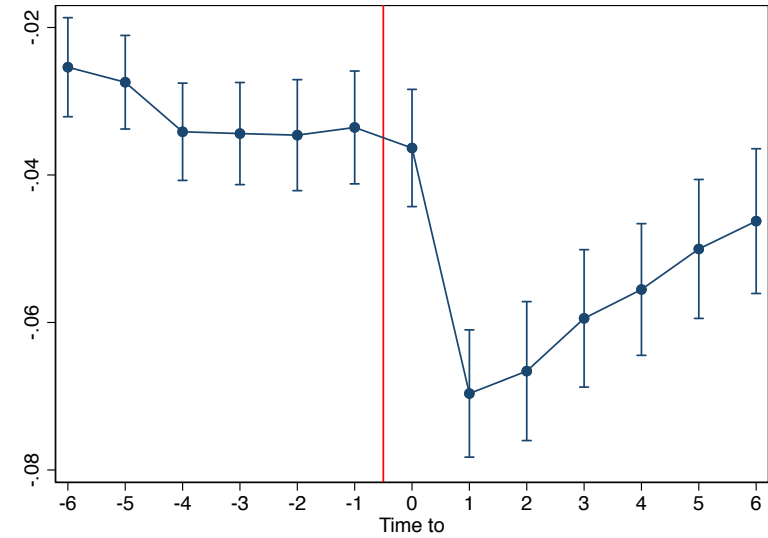
Notes: Figure plots results from an event study design exploring the effect of time series variation in the share of (any) fatal violence or on violence directed towards tourists across four main countries.

8

Figure A5: Heterogeneity of the Effect of News reporting on Aggregated Spending by Card Issuing Country

Panel A: News on Tourists Targeted (Share)
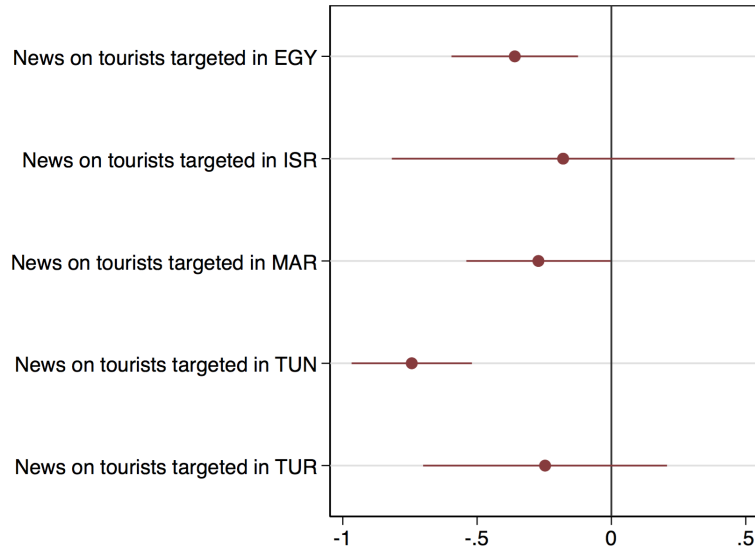
Panel B: News on Fatal Violence (Share)
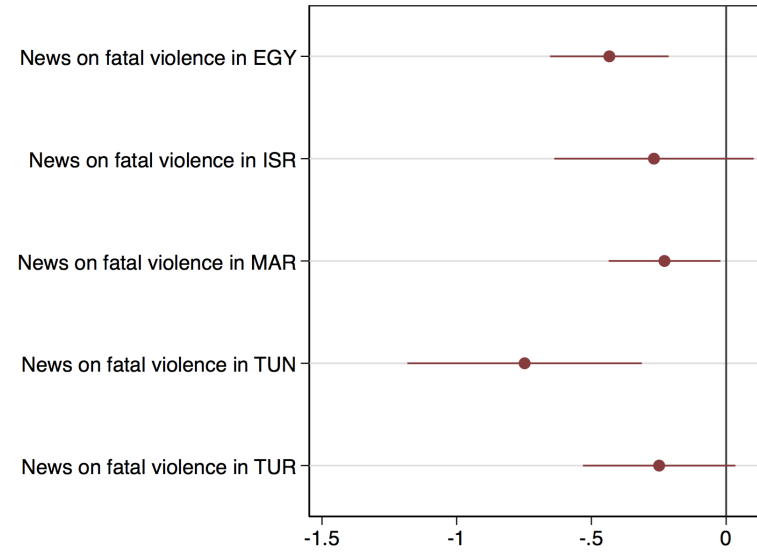
Notes: Figure plots point estimates from a regression that absorbs dyad, issuing-country by time fixed effects and destination by month fixed effects. 90% confidence intervals obtained from clustering standard errors at the dyad level are indicated.

Figure A6: Distribution of Beliefs about Safety or Danger
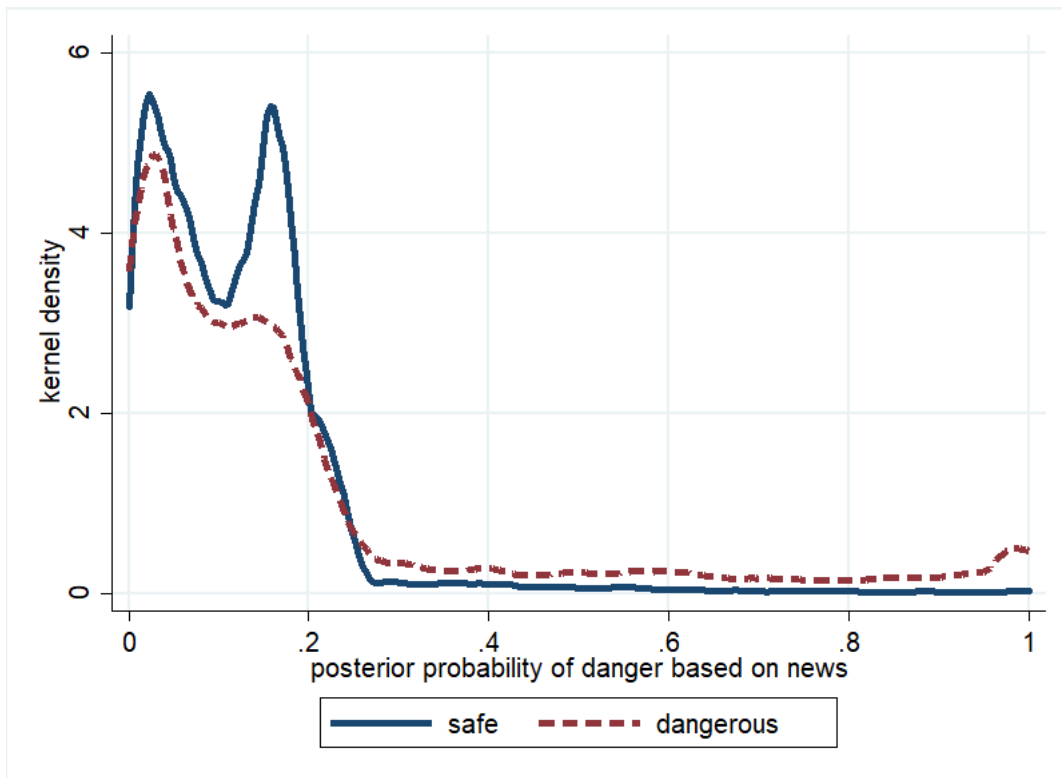
Notes: Figure plots kernel density plotting the distribution of share of newspaper reporting on (any) fatal violence or on violence directed towards tourists conditional on any reporting on fatal violence or tourists being targeted.

Figure A7: Evidence on the correlation between the level of news reporting and average beliefs about violence in months with violent events
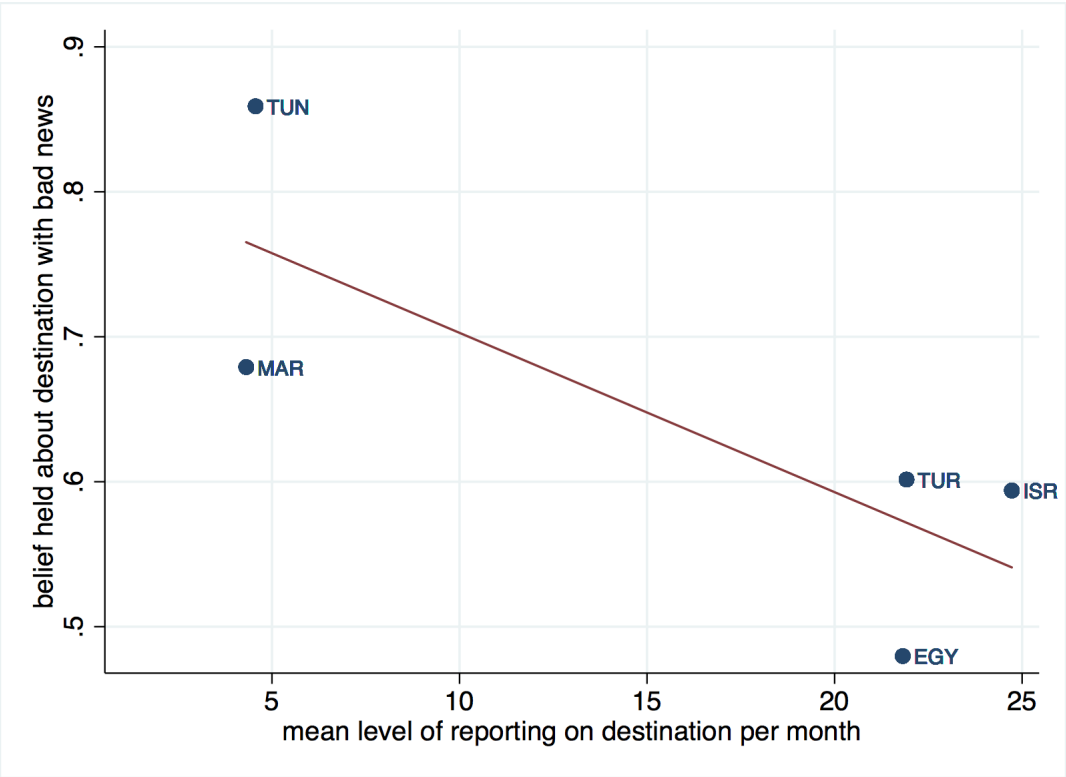
Notes: Figure highlights the interactive relationship between the mean-level reporting on the x-axis about a destination country and the average estimated belief about the state of the world being dangerous conditional on there being news on tourists being targeted on the y-axis. The figure shows that lower levels of mean reporting are associated with higher levels of beliefs about the underlying state of the world indicating danger.

Table A1: Validation of aggregate spending data and official annual dyadic tourist arrival data available from UNWTO for a small subset of countries

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

Notes: The table reports regressions to validate that the anonymized and aggregated credit card data is a good proxy measure of tourism activity. The dependent variable measures the number of active cards in a destination by year in column (1), the number of transactions in column (2) and the total spend 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 A2: Event Characteristics and Reporting Intensity *across dyads*

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

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

Table A3: Principal component-based violence measures and effect on aggregate spending

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	log(Spending)					log(Number of Cards)				
Armed violence component 1	-0.081***			-0.067***	-0.068***			-0.053***		
	(0.017)			(0.017)	(0.016)			(0.017)		
Armed violence component 2		-0.043*		-0.040***		-0.034		-0.034**		
		(0.025)		(0.014)		(0.021)		(0.015)		
Armed violence component 3			0.021	0.017			0.013	0.010		
			(0.014)	(0.015)			(0.013)	(0.015)		
Armed violence component 4				-0.031**				-0.039***		
				(0.013)				(0.014)		
Observations	42254	42254	42254	42254	42299	42299	42299	42299		
R2	.947	.947	.947	.947	.97	.969	.969	.97		
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	NO	YES	YES	YES	YES	NO
Dest. x Time FE	NO	NO	NO	NO	YES	NO	NO	NO	NO	YES

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

Table A4: Robustness of Effect of Violence and Aggregate Spending: Dropping each country in turn

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

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

Table A5: Relationship between Reporting and Aggregate Spending

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

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

Table A6: Robustness to additional control variables and different violent news coding cutoffs

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

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

Table A7: Markov Chain Estimates of Parameters

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

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

Table A8: Effect of Markov Chain Fitted Probability of Dangerous State on Spending across destination countries

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

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

D Description of Machine Learning Method and Validation

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

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

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

Figure B2 shows the performance of the three models, two random forest models and one NB model, on three different folds and on two different samples with sampling rate of 1:1 and 1:10. The y-axis shows the AUC and the x-axis the average

19

precision the respective classifier reached in the sample. The green dots show the result of the ensemble classifier, the simple average of the other three scores. Three things are clear from this. First, the NB performs a lot worse than the random forest - both in terms of AUC and precision. Second, the ensemble is performing better than the random forest, despite the fact that it uses the NB. Thirdly, the precision of the ensemble is less affected by the imbalance and therefore a lot more stable across the different samples. This is an important reason for us to adopt the ensemble method.

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

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

in the direct context of an attack.

We also used the hand-coding to get an impression of the error rate that we imposed through our cutoffs. Between the cutoffs of $c = 0.8$ and $c = 0.75$ we hand-code 626 negatives and 192 positives, i.e. the proportions change considerably compared to the sample above $c = 0.8$. Also we find a rapidly declining rate of false negatives in the remaining recodings. In 4,000 additional observations below $c = 0.75$ we found only an additional 416 positives. This is a rate decline of positives per article of 0.904 to 0.235 and 0.104 so that we suspect that the remaining articles will not contain a lot of actual positives. The resulting distribution of coded attacks is displayed in Figure B5 which shows two kernel densities. The first kernel density displays the overall distribution of predictions for violence against tourists coming out of the ensemble. Clearly, the predictions indicate that attacks against tourists is a rare event with most mass at low predictions. The red curve then shows the distribution of probabilities for the articles we identified through the hand coding. This provides a good confirmation of the decreased rate at which positives could be in the sample.[1]

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

---

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

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

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

Figure B2: AUC and Precision across a set of different models
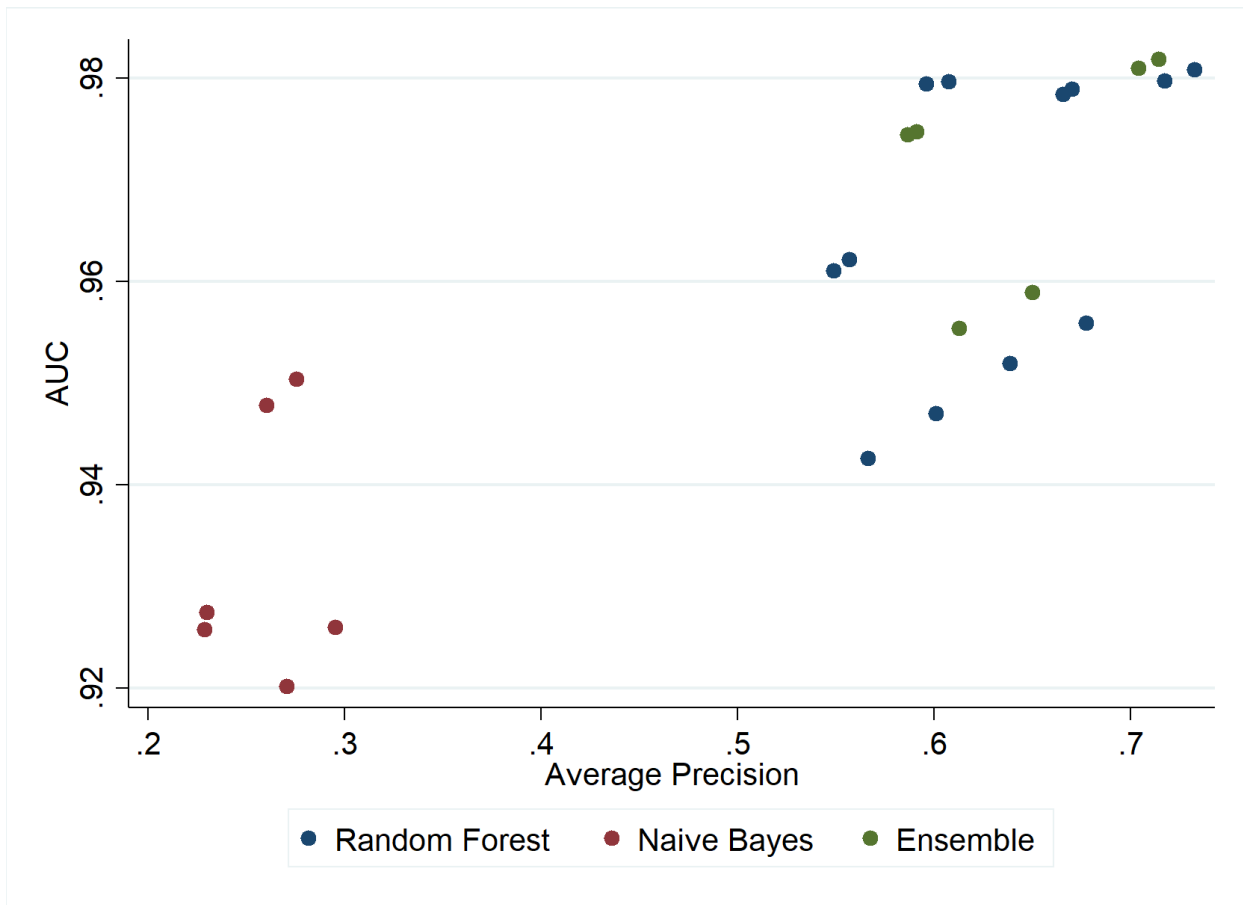
Notes: Figure plots the AUC scores obtained on the testing sample. The AUC is the area under the curve of the ROC curve. For each different maximum tree depth, the AUC is recalculated the three folds. The figure highlights the bias-variance trade-off as the AUC first rises as models are allowed to become more flexible but then stagnates and starts falling as models become too flexible and start overfitting the training data, resulting in worse accuracy in the testing samples.

Figure B3: Classification of articles capturing fatal violence: precision and recall across folds
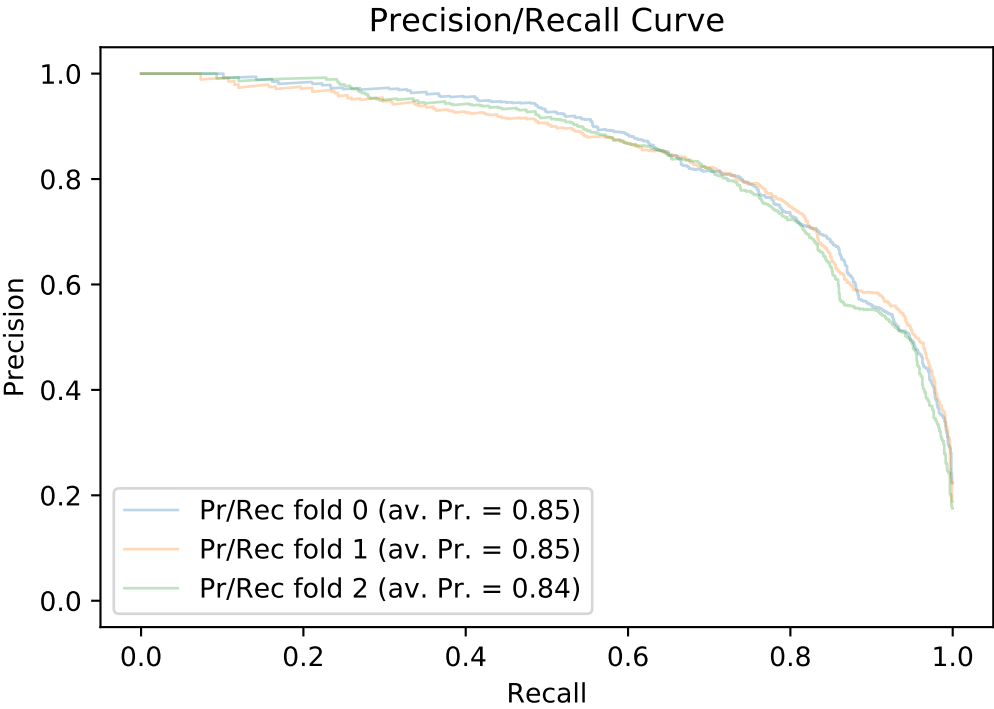
Notes: Figure plots the AUC and precision measured obtained on different folds of data highlighting that Naive Bayes performs worse compared to Random Forests, but that the ensemble of models outperforms the two. AUC is the area under the curve of the ROC curve. Precision of the algorithm captures the share of all correctly classified articles indicating violent events with fatalities among all articles that the algorithm classifies as indicating violence with fatalities.

Figure B4: Classification of articles capturing violence against tourists: precision and recall across folds
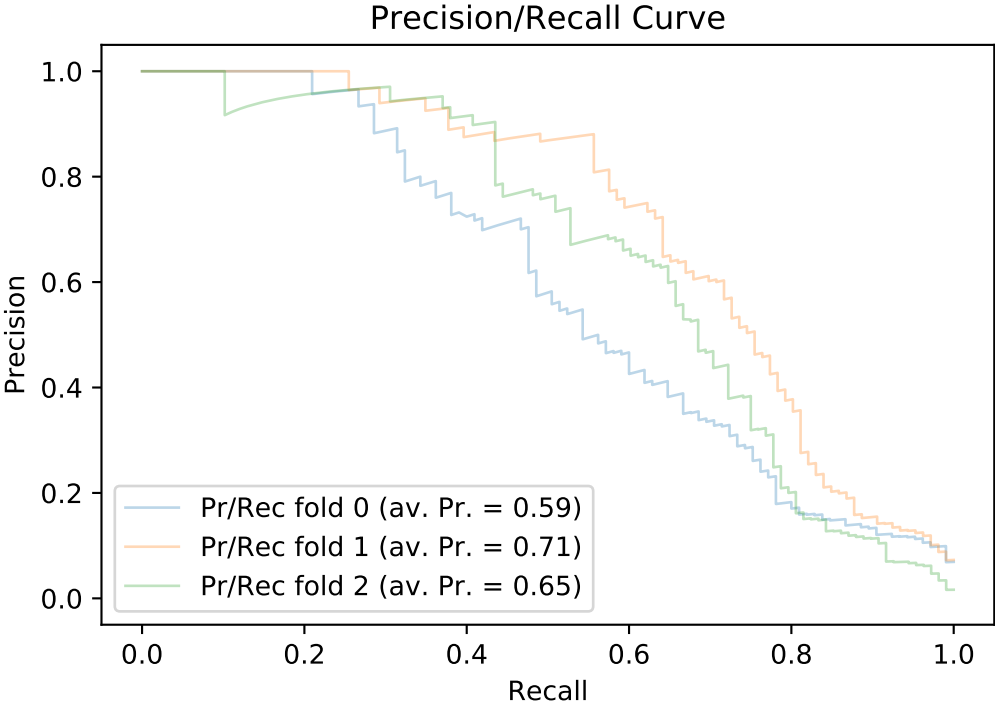
Notes: Figure plots recall on the x-axis. Recall is the true positive rate capturing the share of all actual articles with violence that are correctly classified by the algorithm. The vertical axis plots out precision, which is the share of all correctly classified articles indicating violence against tourists among all articles that the algorithm classifies as indicating violence against tourists.

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

Table B1: Newspaper coverage sample and sources included

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

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

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

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

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

Table B3: Example News Headlines coded as covering violence

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 ¡¿	0.826
TUR	2016	6	A suicide attack causes at least 36 deaths at the Istanbul airport	0.825
TUN	2015	3	Slovak gov't sends plane to evacuate Children's Folk Group from Tunisia	0.818
TUR	2016	1	The jihadist attack on the hotel in Burkina causes 23 dead	0.817
EGY	2012	2	Three South Korean tourists held by locals in Egypt's Sinai	0.814
TUN	2015	3	Third French tourist probably killed in Tunis attack: Hollande	0.814
TUR	2016	6	A suicide attack causes at least 28 deaths at the Istanbul airport	0.808
TUN	2015	6	Thousands of visitors are evacuated from Tunisia after the attack	0.806
TUN	2015	3	Two Spanish pensioners die in the attack against the Bardo Museum	0.805
TUN	2015	6	Bloody Friday Jihadism shows its cruelty in the attacks in Tunisia Lyon and Kuwa	0.803
TUN	2015	3	Feature: Italy mourns four victims in Tunisia's museum attack	0.798
TUN	2015	6	Tunisia's transitional priority target of terror	0.794
TUN	2015	3	We thought we were going to die, we've had a terrible time	0.793
TUN	2015	3	Roundup: Tunisia tries to restore national image after deadly museum attack	0.792
TUN	2015	3	Tunisia ... Hostages taken after attack at Bardo museum	0.791
EGY	2012	2	Urgent: Egypt's Bedouins release three South Korean tourists	0.791
TUN	2015	6	Gunman Focused on Tourists in Slaughter at a Tunisian Beach Hotel	0.786
TUN	2015	6	I could hear the bullets whining Gary Pine English tourist on the beach in Souss	0.786
TUN	2015	3	2nd LD: 21 killed, including 17 foreigners, in Tunisia's museum attack: PM	0.783
TUN	2015	3	Militants hold tourists hostages inside Tunisia museum	0.780
TUN	2015	6	Irish woman among fatalities in Tunisia attack	0.762
TUN	2015	6	Scores Die in Attack at Tunisian Beach Hotel	0.758
TUN	2015	6	5th LD: Death toll rises to 37 in catastrophic hotel attack in Tunisia	0.756
TUN	2015	6	Norway condemns attacks in Tunisia, France, Kuwait	0.753
TUN	2015	3	Belgium to open own investigation into Tunisia attacks	0.752
TUN	2015	3	Bardo museum reopens a week after killings; Tunisia sends out message country s	0.744
TUN	2015	6	4th LD: 28 killed, 36 injured in terror attack on Tunisia hotel	0.739
TUN	2015	6	Thousands of visitors are evacuated from Tunisia after the attack ¡¿	0.735
TUN	2015	3	The attack of the Bardo museum in Tunisia .. What do we know about the nationali	0.734
TUN	2015	6	Germany condemns deadly hotel attack in Tunisia	0.732
TUN	2015	6	After Tunisia attack, UK ups Wimbledon security	0.729
TUN	2015	3	Hollande expresses solidarity with Tunisia after deadly attack	0.726
TUN	2015	6	Deaths of British nationals in Friday's attack in Tunisia rise to 15: FCO	0.725
EGY	2012	2	1st LD Egypt's Bedouins release three South Korean tourists	0.717
TUN	2015	6	3rd LD: Terrorist suspect in Tunisia's hotel attack arrested: official	0.711

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