Online Appendix to "Pandering in the Shadows: How Natural Disasters Affect Special Interest Politics"

Ethan Kaplan, *University of Maryland*
Jörg L. Spenkuch, *Northwestern University*
Haishan Yuan, *University of Queensland*

Contents

List of Figures

List of Tables

Appendix A: Measuring Politics and Disaster Reporting on the Evening News

A.1. *Politics Coverage*

As explained in the main text, Watson uses natural language processing and neural nets, among other methods, to extract concepts, entities, and sentiment from unstructured text. It also categorizes the content of the text according to an enhanced version of the IAB Quality Assurance Guidelines Taxonomy, which defines contextual categories that were originally designed to accurately and consistently describe the content of, say, a website or video clip, in order to facilitate better-targeted advertisements (Interactive Advertising Bureau 2013). We have used Watson to analyze the human-generated summaries of all news segments in the Vanderbilt Television News Archive from its inception in 1968 through July 2017. Critical for our purposes, Watson's taxonomy contains a category for content related to "law, government, and politics." This high-level category contains several subcategories for which Watson returns confidence scores. Since Watson's categorization is not mutually exclusive, we define a particular segment's overall "politics score" as the sum of the confidence scores for all subcategories, up to a maximum of one. In symbols,

$$PoliticsScore_s = \max\{\sum_{c \in C} Score_{s,c}, 1\}, \tag{5}$$

where C denotes the set of subcategories in "law, government, and politics."

Inspecting the distribution of politics scores, there is a very large mass point at exactly zero. Fewer than one in three segments receive a strictly positive score. Among these, however, we observe significant mass in the middle of the distribution. For segments with intermediate scores, it is *a priori* unclear whether or not they should be classified as "political." To make this decision in a principled way, we tasked a research assistant with manually coding a random subset of 1,000 segments. Taking the human judgment as the truth, we find that a cutoff score of .144 provides a good balance between sensitivity and selectivity. Using this cutoff, Appendix Table A.1 constructs the confusion matrix. Given an accuracy of 91.6% and a false positive (negative) rate of 7.7% (11.1%), our automated detection of political content appears to perform well—though it is certainly not perfect.

Given the classification of all news segments, we measure politics coverage by network n on day t as the fraction of total airtime the newscast devoted to political matters. In symbols:

$$(6) \qquad News_{n,t} \equiv (\sum_{s \in P_{n,t}} Duration_s)/(\sum_{s \in S_{n,t}} Duration_s),$$

where $P_{n,t}$ denotes the set of news segments that are deemed to contain political content and $S_{n,t}$ is the set of all segments, including commercials.

According to this measure, on an average day, the median network contained in VTNA spends about 29% of airtime reporting on political issues. Importantly, our measure appears to capture meaningful variation. Consider, for example, Figure A.1, which plots the nightly duration of politics coverage on the ABC evening news during 2012 (thick line), superimposing the start dates of natural disasters (dashed lines). Several patterns stand out. First, there is substantial high-frequency variation in politics coverage on the evening news. While some of that variation is undoubtedly measurement error, we find it reassuring that many of the local peaks occur around the same time as significant political events, such as the Republican National Convention (August 27–30), the (vice-)presidential debates (October 3, 11, 16, and 22), and Election Day (November 6).[1] Second, although we already restrict attention to nontrivial disasters, adverse events like floods, tornadoes, or hurricanes, are not terribly rare. Third, many, but by no means all, of the disasters in our data coincide with temporary lows in politics reporting. For instance, landfall of Superstorm Sandy on October 29 coincided with next-day politics coverage roughly 4.4 minutes, or about 62%, below normal—even though the presidential election was little more than a week away.

A.2. *Disaster Coverage*

As noted in the main text, we also use IBM Watson to detect disaster-related reporting on the evening news. In analogous fashion to our politics classifier, we define the $DisasterScore_s$ of news segment s as the sum of the relevant confidence scores that Watson returns. As was the case with respect to politics-related content, most news segments have a score of exactly zero, and a much smaller number has a score of one. Given that there is again nontrivial mass in the middle of the distribution, we proceed the same way as classifying political content. In the end, we say that news segment s is disaster-related if and only if $DisasterScore_s > .178$, where .178 corresponds to the chosen cutoff.

[1] Coincidentally, the RNC occurred during the landfall of hurricane Isaac.

Appendix Table A.2 presents the confusion matrix. Given an overall accuracy of 97.9% with a false positive (negative) rate of 1.5% (15.2%), we again conclude that our automated measurement of disaster-related content works reasonably well.

Appendix B: Correcting for Measurement Error in News Coverage

Since we use machine learning to detect political content on the evening news, our measure of politics reporting will inevitably contain measurement error. This measurement error is necessarily non-classical.[2] In fact, the errors are by construction "one-sided," meaning that they are correlated with the true outcome. In what follows, we show that this causes attenuation bias in linear probability models, and we provide estimates that correct for the bias.

While we are focused on the specifics of our setting, we note that similar problems arise in virtually all applications in which researchers use a machine-learning classifier to measure outcomes. The theoretical results below are, therefore, much more broadly applicable.

We also note that our derivations differ from prior work on measurement error, which, for the most part, assumes i.i.d. errors in either the dependent or independent variables. While we aware of models with non-independent measurement error in a right-hand side variable, we do not know of results pertaining to measurement error that is correlated with the realization of the left-hand side variable, as in our application.

B.1. *Derivation of Bias*

B.1.1. *Main Result*

We first study the simple case in which reporting on a particular day is either about politics ($Y = 1$) or not ($Y = 0$). After deriving a correction for measurement error in this setting, we extend our result to the case in which the outcome variable is a weighted average of segments that do and do not cover politics (i.e., $\bar{Y} = \sum_j \omega_j Y_j$).

Consider the following linear probability model

$$Y = X\beta + \epsilon, \tag{7}$$

where Y is the a outcome, X is a (de-meaned) vector of covariates, and ϵ denotes the

[2] To see this, note that the outcome is bounded by zero and one, which violates the assumptions in the classical measurement error model.

error term. The parameter of interest is β.

Let $\tilde{Y}$ denote the true outcome. If there were no measurement error, ϵ would be a binary random variable equal to $1 - X\beta$ with probability $Pr(\tilde{Y} = 1) = X\beta$ and equal to $-X\beta$ with the complementary probability. However, when using an automated classifier to measure the outcome, Y will generally contain some error. Assume that $\theta_0 = Pr(Y = 0|\tilde{Y} = 0)$ and $\theta_1 = Pr(Y = 1|\tilde{Y} = 1)$. In the language of machine learning, θ_0 denotes the specificity of the classifier (i.e., the probability of correctly identifying a "true negative"), whereas θ_1 corresponds to its sensitivity (i.e., the probability of detecting a "true positive"). With this notation in hand, the expectation of ϵ conditional on X is given by

$$\begin{aligned}
E[\epsilon|X] &= X\beta[\theta_1(1 - X\beta) + (1 - \theta_1)(-X\beta)] + (1 - X\beta)[\theta_0(-X\beta) + (1 - \theta_0)(1 - X'\beta)] \\
&= X\beta(\theta_1 + \theta_0 - 2) + 1 - \theta_0.
\end{aligned}$$

Thus, the expectation of the structural measurement error model in eq. (7) is

$$E[Y|X] = X\beta(\theta_0 + \theta_1 - 1) + 1 - \theta_0.$$

Further, using the standard formula, the expectation of the OLS estimator is given by

$$\begin{aligned}
E[\hat{\beta}_{OLS}|X] &= (X'X)^{-1}X'E[Y|X] \\
&= (X'X)^{-1}X'X\beta(\theta_0 + \theta_1 - 1) + (X'X)^{-1}X'(1 - \theta_0) \\
&= \beta(\theta_0 + \theta_1 - 1),
\end{aligned}$$

where the last equality uses the fact that $(X'X)^{-1}X'(1 - \theta_0)$ corresponds to regressing a constant on X, which returns zero. As a result, in the simple case, to correct for measurement error in the dependent variable, we must inflate the the OLS estimate by $(\theta_0 + \theta_1 - 1)$, i.e.,

$$\beta = \frac{E[\hat{\beta}_{OLS}|X]}{\theta_0 + \theta_1 - 1}. \tag{8}$$

Eq. (8) shows that unless $\theta_0 = \theta_1 = 1$—in which case there is no measurement error—the OLS estimate will be attenuated, and the bias depends on both the specificity and sensitivity of the classifier.

B.1.2. *Extension to Weighted Averages*

The result above is likely to be useful in a broad array of applications in which researchers use machine learning methods to measure outcomes. In our specific setting, however, it is not directly applicable because our measure of politics coverage on the evening news is a weighted sum of mismeasured binary variables. Nonetheless, it is straightforward to extend our result to this case.

In particular, our regression model is given by

$$\bar{Y} = X\beta + \bar{\epsilon},$$

with $\bar{Y} \equiv \sum_j \omega_j Y_j$, $\bar{\epsilon} \equiv \sum_j \omega_j \epsilon_j$, and weights $\sum_j \omega_j = 1$. In our application, ω_j corresponds to the length of news segment j relative to the entire broadcast. Note, X does not need to averaged because it varies only on the daily level and not across the different news segments within a given show.

Proceeding as above,

$$\begin{aligned} E[\hat{\beta}_{OLS}|X] &= (X'X)^{-1}X'E[\bar{Y}|X] \\ &= (X'X)^{-1}X'X\beta + (X'X)^{-1}X'\sum_j \omega_j (X\beta(\theta_1 + \theta_0 - 2) + 1 - \theta_0) \\ &= \beta + (X'X)^{-1}X'(X\beta(\theta_1 + \theta_0 - 2) + 1 - \theta_0) \\ &= \beta(\theta_0 + \theta_1 - 1). \end{aligned}$$

As a result, even when the outcome is a weighted average of mismeasured binary variables, as in our application, we can continue to adjust our regression estimates for attenuation bias through inflating them by $(\theta_0 + \theta_1 - 1)$.

B.2. *Corrected Estimates*

Appendix Table A.3 presents the measurement-error-corrected estimates and contrasts them to estimates without correction. Taking the judgement of the human coder as the ground truth allows us to estimate θ_0 and θ_1 by looking at the relevant confusion matrices (e.g., Appendix Table A.1 and Table 2). Relative to the their uncorrected counterparts in the main text, the point estimates estimates of disasters' impact on politics coverage need to inflated by approximately 25%.

Since $\hat{\theta}_0$ and $\hat{\theta}_1$ are themselves random variables, we can use the delta method to calculate standard errors for the adjusted point estimates. Under the assumption that classification errors are i.i.d., the estimated variance of a corrected coefficient is given by

$$(9) \qquad \widehat{Var}(\hat{\beta}) = [\widehat{Var}(\hat{\beta}_{OLS}) + (\hat{\beta}_{OLS}/(\hat{\theta}_0 + \hat{\theta}_1 - 1))^2(\widehat{Var}(\hat{\theta}_0) + \widehat{Var}(\hat{\theta}_1))]/(\hat{\theta}_0 + \hat{\theta}_1 - 1)^2.$$

Relative the standard errors in Figure 4 in the main text, the standard errors on the corrected estimates are slightly larger, but not enough to affect any of our qualitative conclusions.

For additional measurement-error-corrected estimates, see Appendix C.

Appendix C: Ancillary Results and Robustness Checks

Appendix Table A.3 probes the robustness of the effect of disasters on politics reporting with respect to: *(i)* standardizing the left-hand-side variable within each TV network (columns (1)–(6) and (13)–(18)); *(ii)* measuring coverage in raw minutes instead of relative shares (columns (7)–(12) and (19)–(24)); *(iii)* broadening the sample to include news shows on CNN and Fox News in addition to ABC, CBS, and NBC (columns (4)–(6), (9)–(12), (16)–(18), and (22)–(24)); and *(iv)* simultaneously correcting for the LHS measurement error introduced through machine learning, as explained in Appendix C (columns (13)–(24)). None of these changes materially affect our conclusions.

Appendix Tables A.4 and A.5 respectively replicate the results in Tables 2 and 6 in the main text, relying on either all domestic natural disasters reported in EM-DAT (upper panels) or large foreign as well domestic disasters (lower panels) instead of only the latter. As for the results in Appendix Figure A.4, when we refer to large foreign disasters, we mean the 178 foreign disasters that fall into the top-1% in terms of either the number of deaths, total number of people affected, or total damages. Again, the estimated effects decline somewhat in magnitude but are otherwise similar to their counterparts in the main text.

Appendix D: Data Description and Definitions

D.1. *MapLight*

As explained in the main text, information on connections between politicians and special interests, the positions of special interest groups on particular pieces of legislation, and congressmen's votes on the same measures comes from MapLight. MapLight is a nonpartisan, 501(c)(3) nonprofit organization whose goal it is to "reveal the influence of money in politics, inform and empower voters, and advance reforms that promote a more responsive democracy."

MapLight staff scour publicly available sources, like congressional testimony, news databases, and trade associations' websites, to compile lists of organizations and interest groups that either supported or opposed a particular piece of federal legislation, excluding bills and amendments that are purely ceremonial. Starting with legislation considered in the 109th Congress, MapLight provides data on interest group positions on more than 10,000 individual bills—most of which never receive a vote. MapLight also uses campaign contribution data provided by the Center for Responsive Politics in order to link interest groups' positions on a particular bill to their donations to individual congressmen, the relevant roll-call votes, and metadata on the bill. The linked records are then made publicly available at `http://classic.maplight.org/us-congress/bill`.[3]

Our analysis relies on the linked records for all 1,525 bills that (a) received a passage vote in the House of Representatives prior to October 2017, and (b) were supported or opposed by at least one special interest group.

D.2. *EM-DAT*

Data on natural disasters come from the Centre for Research on the Epidemiology of Disasters (CRED) at the Université Catholique de Louvain, which maintains the Emergency Events Database (EM-DAT). EM-DAT contains core information on the occurrence and effects of over 22,000 natural and man-made disasters worldwide. According to the CRED website "the main objective of the database is to serve the purposes of humanitarian action at national and international levels. The initiative aims to rationalize decision making for disaster preparedness, as well as provide an objective base for vulnerability assessment and priority setting."

For an adverse event to be recorded as a disaster in EM-DAT it must satisfy at least one of the following criteria: 10 or more people dead, 100 or more people affected, an officially declared state of emergency, or a call for international assistance. CRED staff assess these criteria based on various sources, including UN agencies, non-governmental organizations, insurance companies, press agencies, as well as other research institutes.

For our main analysis, we restrict attention to natural disasters that occurred within the United States. We further restrict attention to sudden-onset disasters and days that

[3]For additional information on MapLight and its methodology, see `http://classic.maplight.org/us-congress/guide/data`.

fall into the top tercile of the distribution in terms of either deaths, number of people affected, or damages. The latter restriction is intended to filter out relatively minor incidents that are unlikely to crowd out media attention, while the former one ensures that we only work with disasters for which the start date is precisely enough defined to obtain sharp identification. In practice, this means that we exclude epidemics, heat waves, and wildfires from our main analysis. In the robustness checks in Appendix C, we show that our findings remain qualitatively unchanged if we included all domestic disasters recorded in EM-DAT. After imposing these sample restrictions, we are left with 200 large domestic disasters that occurred between 2005 and the end of 2017.

D.3. *Vanderbilt Television News Archive*

Information on the content of TV news broadcasts comes from the Vanderbilt Television News Archive (VTNA). Starting in 1968, VTNA collects and archives daily recordings of the regularly scheduled evening news programs on ABC, CBS, and NBC. In 1995, coverage was expanded to include approximately one hour per day from CNN, and, in 2004, to also include Fox News. Originally, VTNA attempted to provide a short, human-generated summary of every story that aired, information on its duration, as well as its order of appearance. Unfortunately, in 2014 VTNA stopped producing human-generated summaries of stories from weekday newscasts on CBS, NBC, and Fox News. In private communication, representatives from VTNA indicated that they scaled down on human-generated content in order to experiment with automated techniques, which have not been as successful as they had hoped.

As explained in the main text, we use state-of-the-art machine learning as implemented by IBM Watson to classify each news story in VTNA based on the provided summary.[4] In particular, Watson categorizes the content of unstructured text according to an enhanced version of the IAB Quality Assurance Guidelines Taxonomy Interactive Advertising Bureau (2013), which defines contextual categories that were originally designed to consistently describe web content in order to facilitate more relevant advertising and allow for *ex post* analysis.

With the classification from Appendix B in hand, we measure politics coverage by network n on day t as the fraction of total airtime the newscast devoted to political matters. In symbols, $News_{n,t} \equiv (\sum_{s \in P_{n,t}} Duration_s)/(\sum_{s \in S_{n,t}} Duration_s)$, where $P_{n,t}$

[4]We access Watson remotely through an API. For a free demonstration of Watson's text-analytic capabilities see `https://natural-language-understanding-demo.ng.bluemix.net`.

denotes the set of news segments that are deemed to contain political content and $S_{n,t}$ is the set of all segments, including commercials.

To measure disaster-related news reporting we use Watson and the VTNA data in an analogous fashion (see Appendix A for details).

D.4. *NewsLibrary*

As explained in the main text, we complement our daily measure of attention to politics on the evenings news with a second one that focuses exclusively on individual representatives. To this end, we have searched the NewsLibrary database for newspaper articles that mention an in-state congressperson by name. Specifically, for each representative and each year she is in office, we limit our search to newspapers from her home state and submit the following query: "(Congressman AND *name*) OR (Congresswoman AND *name*) OR (Representative AND *name*)", where *name* denotes the person's last name. We then count, for each day, the number of articles returned, and use this information to construct a daily panel of newspaper reports on local congresspeople.

At the time of our searches, the NewsLibrary database was owned by NewsBank, Inc. and indexed more than 6,500 newspapers from all around the United States—though coverage varies considerably across space and time. For more information on NewsLibrary, see `https://newslibrary.com`.

D.5. *Web Searches*

We measure citizens' interest in Congress using Google searches for the following terms: "politics," "Congress," "Congressman," "Representative," "government," "House of Representatives," and "vote." The relevant data come from Google Trends, which we accessed via an API, and span the same frame as the MapLight data.[5]

Google Trends provides information on the daily search volume for arbitrary keywords. For each query the maximum of the time series that Google Trends returns is indexed by 100. Since it is not possible to download daily data for a period longer than three months, we proceeded by downloading, for each keyword, the daily data for any given month, which we then multiply by the monthly search volume index for the same keyword. In symbols,

$$v_{k,t} = \tilde{v}_{k,m,d}\bar{v}_m \tag{10}$$

[5] Google Trends is available at `https://trends.google.com/`.

where $v_{k,t}$ denotes the search volume for keyword k on date t, $\tilde{v}_{k,m,d}$ is the search volume for the same term on day d of month m, and $\bar{v}_t$ is the average volume during the same month. This adjustment follows Durante and Zhuravskaya (2018), and it ensures that the indexed daily search volume for a given keyword is comparable over time. We then standardize the entire time series for each keyword. The resulting variable serves as the outcome in the regression model in eq. (2) in the main text, i.e., $GS_{k,t}$.

To measure disaster-related searches we proceed in analogous fashion, focusing on the following set of keywords: "disaster," "volcano," "earthquake," "flood," "landslide," "storm," "hurricane," "blizzard" and "tornado."

D.6. *Other Data Sources*

D.6.1. *Congressional Speech*

Data on congressional speech come from Gentzkow et al. (2018). Gentzkow et al. (2018) obtained copies of the Congressional Record—which contains all text spoken on the floor of either the U.S. House or the U.S. Senate—for the 43rd to 114th Congresses from HeinOnline. They then used automated scripts to parse the text from each session in order to extract full-text speeches, metadata on speeches and their speakers, and counts of bigrams.

We use their data on full-text speeches in the House and the accompanying metadata for the 109th–114th Congresses. These restrictions are imposed to ensure that the setting for our analysis of congressional speech corresponds as closely as possible to the setting of our main analysis. We further process the full text of speeches by removing common stop words, such as "a," "about," "between," "because," etc., and by counting (i) the total number of remaining words spoken on a particular day, as well as (ii) the number of words that are plausibly related to natural disasters. To identify the latter we conduct a simple keyword search for the following terms: "disaster," "emergency," "relief," "help," "rebuild," "assistance," "victim," "storm," "hurricane," "tornado," "flood," "landslide," "earthquake," or "volcano." These daily counts then serve as outcome variables in our ancillary results (see above).

D.6.2. *Number of Votes, Roll-Call Types & Vote Issues*

Data on the type of a roll-call vote come from the PIPC House Roll Call Database Crespin and Rhode (2018). Coverage of PIPC begins with the 83rd Congress. Among other information, these data contain a variable classifying each roll call as one of

59 mutually exclusive types, such "quorum call," "final passage / adoption of a bill," "final passage / adoption of conference report," "passage / adoption of a bill under suspension of the rules," "passage / adoption of a joint resolution under suspension of the rules," "straight amendments," "amendments to amendments," "motion to discharge," "motion to reconsider," etc. Roll calls from the 83rd to 100th Congresses were manually assigned to one of these categories. Starting with the 101st Congress, PIPC began using a supervised machine-learning model to assign types based on the roll call-specific description and other information provided on the Clerk of the House's website. In training this model, the hand-coded votes from priors years served as examples.

We restrict attention to House votes during the 109th–115th Congresses and rely on the classification in the PIPC database in conjunction with ancillary information from *voteview.com* Lewis et al. (2018) to count the total number of roll calls of particular type that were held on a given day. Since our categories are broader than those in the PIPC database, we aggregate over related types.

PIPC also contains hand-coded issue codes for each roll call. PIPC obtains the relevant information comes from the Comparative Agendas Project (PAP), which collects and organizes data from archived sources to track policy outcomes across countries.[6] Again, we aggregate over different issue codes in the raw data to define the issue categories used in the main text.

References

Crespin, Michael H., and David Rohde. 2018. Political Institutions and Public Choice Roll-Call Database. available at `https://ou.edu/carlalbertcenter/research/pipc-votes/`

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

Frisch, Ragnar, and Frederick V. Waugh. 1933. "Partial Time Regressions as Compared with Individual Trends," *Econometrica*, 1(4), 387–401.

Gentzkow, Matthew, Jesse M. Shapiro, and Matt Taddy. 2018. Congressional Record for the 43rd–114th Congresses: Parsed Speeches and Phrase Counts. Palo Alto, CA: Stanford Libraries [distributor]. available at `https://data.stanford.edu/congress_text`.

Interactive Advertising Bureau. 2013. "IAB Quality Assurance Guidelines 2.0." retrieved from `https://www.iab.com/news/iab-releases-quality-assurance-guidelines-2-0` (July 1, 2017).

[6]See `https://www.comparativeagendas.net/`.

Lewis, Jeffrey B., Keith Poole, Howard Rosenthal, Adam Boche, Aaron Rudkin, and Luke Sonnet (2018). Voteview: Congressional Roll-Call Votes Database. available at `https://voteview.com/`.

Spenkuch, Jörg L. 2012. "Moral Hazard and Selection Among the Poor: Evidence from a Randomized Experiment." *Journal of Health Economics*, 31(1): 72–85.

Appendix Figure A.1: Politics Reporting on the ABC Evening News in 2012

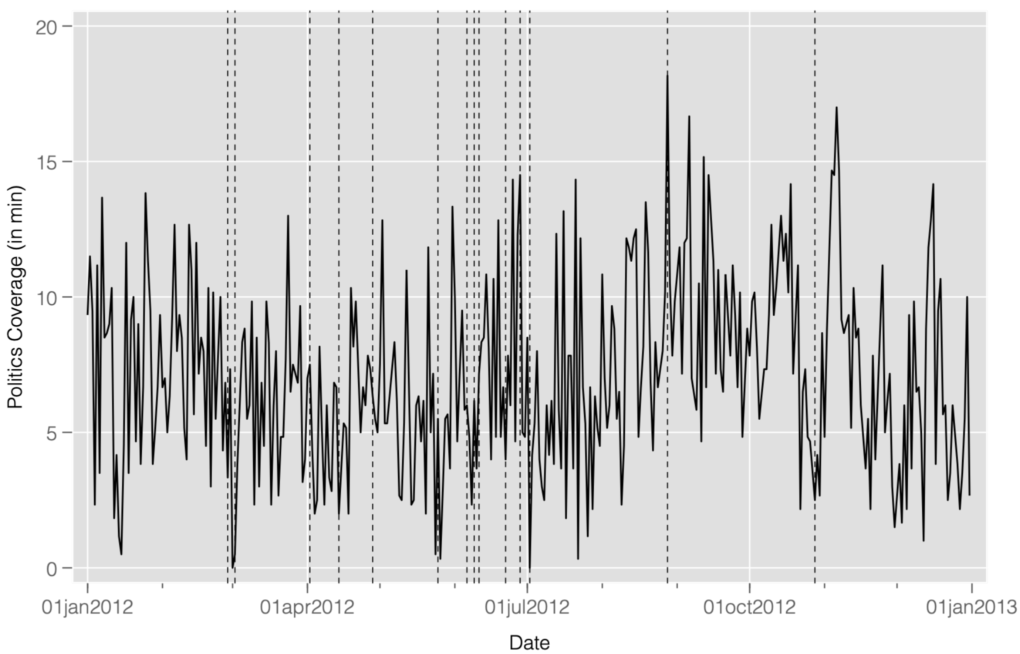

Notes: Figure shows the duration of politics coverage (in minutes) on the ABC evening news in 2012. Dashed vertical lines indicate the onset of natural disasters, as reported in EM-DAT. For a detailed description of the underlying data, see the Data Appendix.

Appendix Figure A.2: Dropping Days that Fall in Multiple Event Windows

(a) Effect of Disasters on Votes

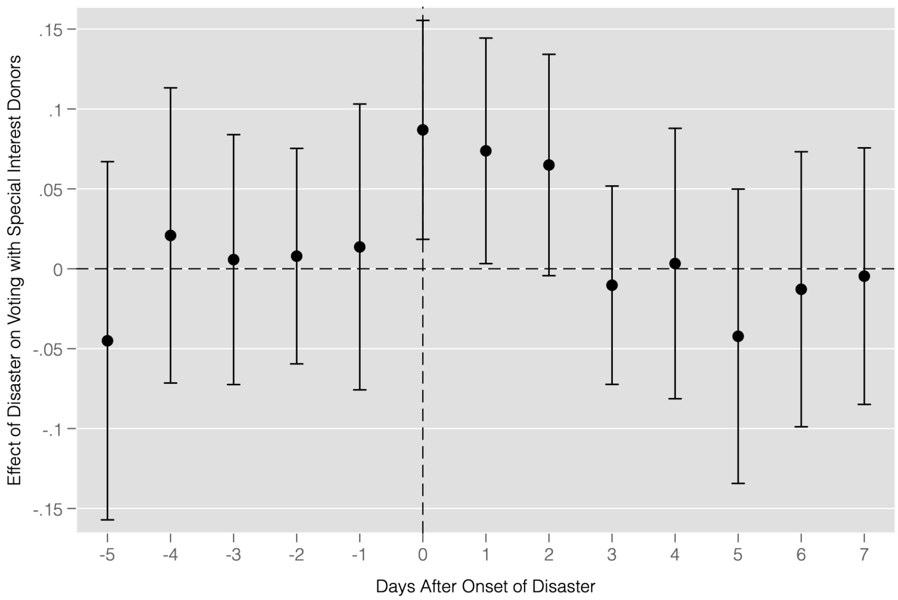

(b) Effect of Disasters on Politics Reporting

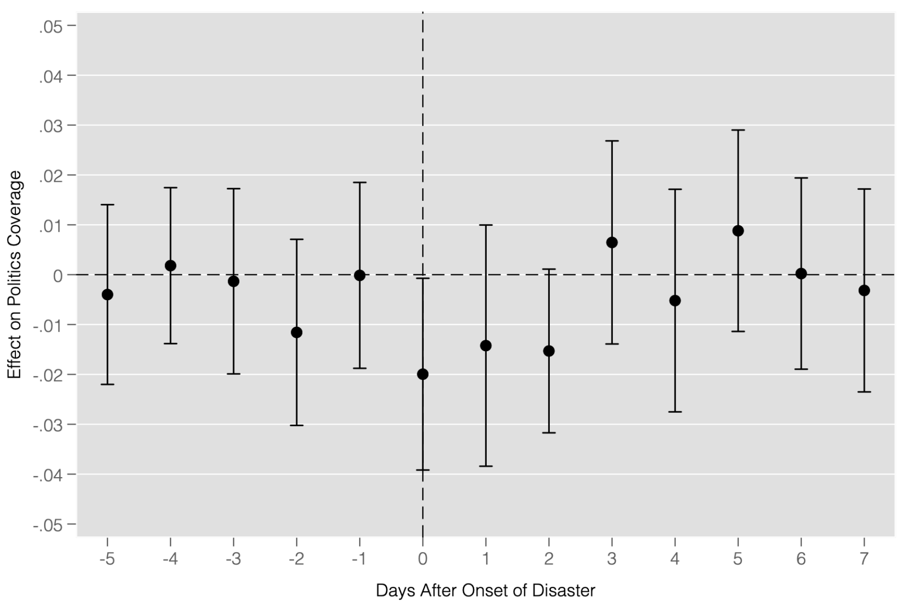

Notes: Figure replicates Figures 4 and 6 in the main text, dropping all votes (upper panel) and news reports (lower panel) panel that fall in the event window of more than one domestic natural disaster.

Appendix Figure A.3: Impact of Disasters on Newspaper Coverage of Local Representatives

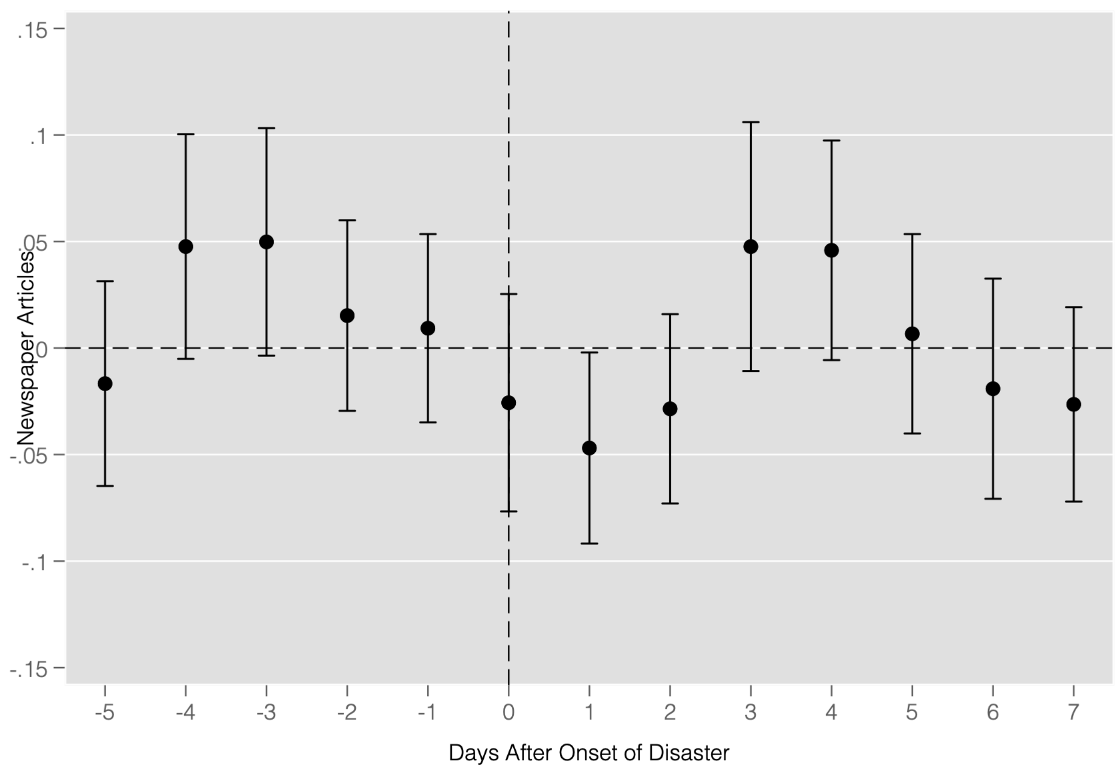

Appendix Figure A.4: Impact of Disasters on Disaster-Related Google Searches

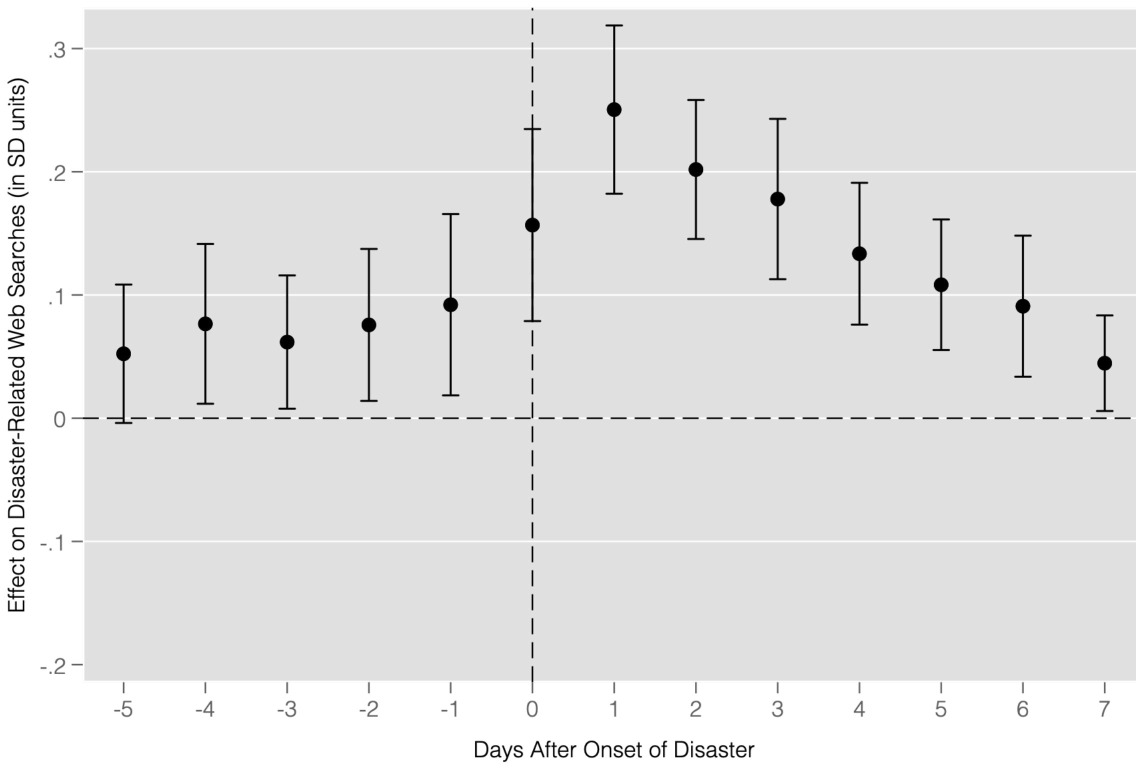

Appendix Figure A.5: Impact of Disasters on Congress-Related Google Searches

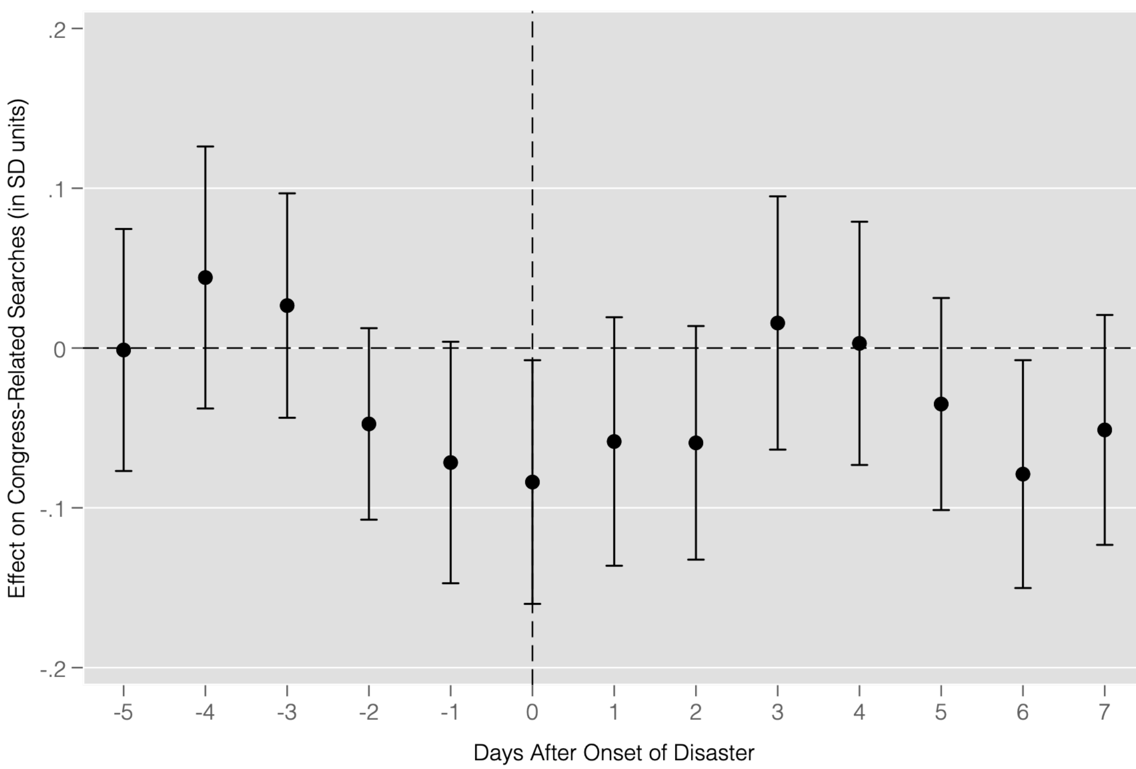

Appendix Figure A.6: Kolmogorov-Smirnov Test for Delay in Bills Voted Upon Before vs. After Disaster

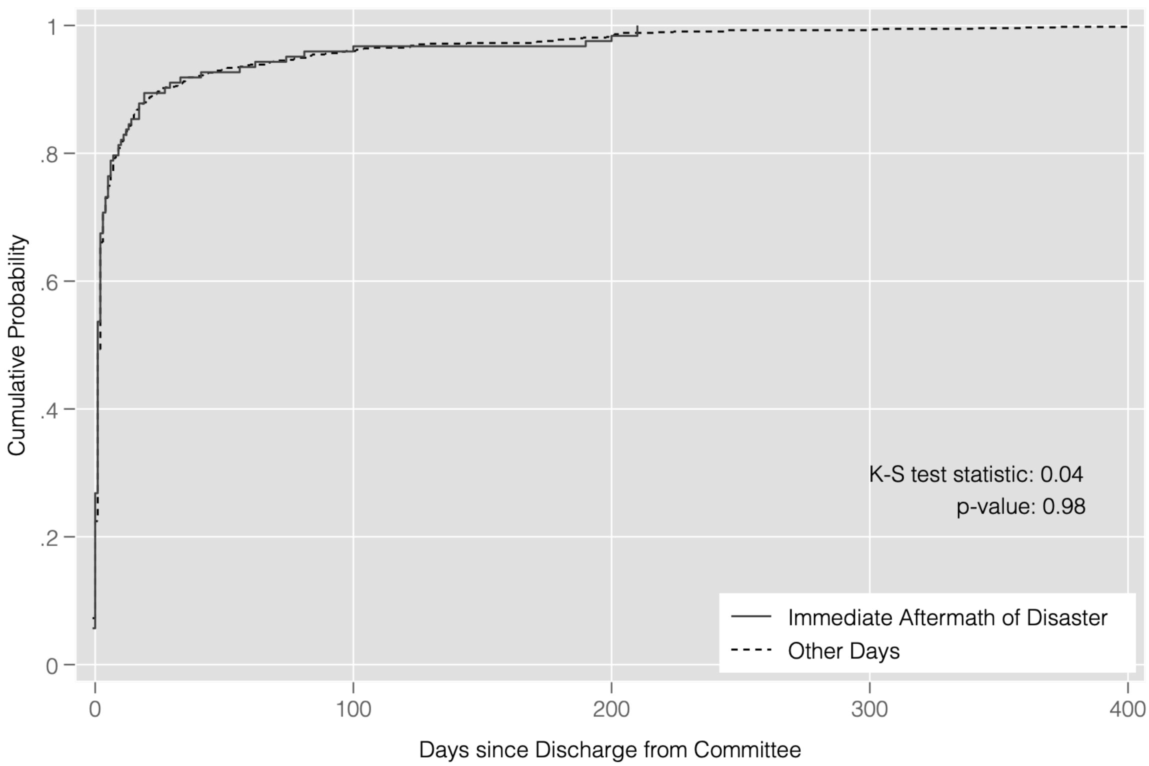

Appendix Table A.1: Performance of Politics-Reporting Classifier

A. Confusion Matrix

		Watson	
		Nonpolitical	Political
Human Coder	Nonpolitical	73.2%	6.1%
	Political	2.3%	18.4%

B. Performance Metrics

Correctly Classified:	91.60%
Sensitivity:	88.89%
Specificity:	92.31%
False-Positive Rate:	7.69%
False-Negative Rate:	11.11%

Notes: Entries in the upper panel are percentages comparing Watson's classification of 1,000 randomly drawn news segments as related to "law, government & politics" against the judgements of a human coder. Entries in the lower panel are descriptive statistics for the performance of the automated classification, taking the judgements of the human coder as ground truth.

Appendix Table A.2: Performance of Disaster Classifier

A. Confusion Matrix

		Watson: Not Disaster Related	Watson: Disaster Related
Human Coder	Not Disaster Related	94.0%	1.4%
	Disaster Related	.7%	3.9%

B. Performance Metrics

Correctly Classified:	97.90%
Sensitivity:	84.78%
Specificity:	98.53%
False-Positive Rate:	1.47%
False-Negative Rate:	15.22%

Notes: Entries in the upper panel are percentages comparing Watson's classification of 1,000 randomly drawn news segments as related to natural disasters against the judgements of a human coder. Entries in the lower panel are descriptive statistics for the performance of the automated classification, taking the judgements of the human coder as ground truth.

Appendix Table A.3: Regression Evidence on the Effect of Disasters on News Crowd-Out

A. OLS Estimates	Standardized Meaure of Politics Coverage						Minutes of Politics Reporting on a Particular Network					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Immediate Aftermath of Disaster	-.152***	-.147***	-.149***	-.114***	-.116***	-.117***	-.581***	-.552***	-.556**	-.566**	-.711***	-.568***
	(.046)	(.042)	(.041)	(.038)	(.033)	(.033)	(.191)	(.172)	(.170)	(.263)	(.220)	(.200)
Hypothesis Tests [p-values]:												
H0: No Effect of Disaster	.001	.001	.000	.003	.001	.000	.003	.002	.001	.033	.002	.005
Fixed Effects:												
Year × Month	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Network × Day of the Week	No	No	Yes	No	No	Yes	No	No	Yes	No	No	Yes
Sample	Big Three	Big Three	Big Three	All Channels	All Channels	All Channels	Big Three	Big Three	Big Three	All Channels	All Channels	All Channels
R-Squared	.003	.101	.110	.001	.102	.117	.002	.099	.120	.000	.073	.541
Number of Observations	10,932	10,932	10,932	17,207	17,207	17,207	10,932	10,932	10,932	17,207	17,207	17,207

B. Measurement-Error-Corrected Estimates	Standardized Meaure of Politics Coverage						Minutes of Politics Reporting on a Particular Network					
	(13)	(14)	(15)	(16)	(17)	(18)	(19)	(20)	(21)	(22)	(23)	(24)
Immediate Aftermath of Disaster	-.187***	-.182***	-.183***	-.140***	-.143***	-.144***	-.716***	-.680***	-.685***	-.697**	-.875***	-.699***
	(.058)	(.054)	(.053)	(.048)	(.043)	(.042)	(.242)	(.219)	(.217)	(.328)	(.280)	(.253)
Hypothesis Tests [p-values]:												
H0: No Effect of Disaster	.002	.001	.001	.004	.001	.001	.004	.002	.002	.035	.002	.006
Fixed Effects:												
Year × Month	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Network × Day of the Week	No	No	Yes	No	No	Yes	No	No	Yes	No	No	Yes
Sample	Big Three	Big Three	Big Three	All Channels	All Channels	All Channels	Big Three	Big Three	Big Three	All Channels	All Channels	All Channels
Number of Observations	10,932	10,932	10,932	17,207	17,207	17,207	10,932	10,932	10,932	17,207	17,207	17,207

Notes: Entries are coefficients and standard errors from regressing different measures of politics reporting on an indicator variable for the day of and the two days after the reported onset of a large domestic disaster, as explained in Appendix D. The outcome in columns (1)–(6) and (13)–(18) is the within-network standardized number of minutes of politics coverage on the evening news. The outcome in columns (7)–(12) and (19)–(24) is the raw duration of politics reporting (in minutes). The sample in columns (1)–(3), (7)–(9), (13)–(15), and (19)–(21) includes only news shows that aired on ABC, CBS, and NBC, whereas the sample in columns (4)–(6), (10)–(12), (16)–(18), and (22)–(24) also includes shows on CNN and FNC. Entries in the upper panel are OLS estimates and the associated standard errors, while those in the lower panel use the measurement-error-adjustment developed in Appendix C. All standard errors are clustered by year-month. ***, **, * denote statistical significance at the 1

Appendix Table A.4: Replication of Table 2

A. Using All Domestic Natural Disasters	Vote with Special-Interest Donors					
	(1)	(2)	(3)	(4)	(5)	(6)
Immediate Aftermath of Disaster (β_{post})	0.056***	0.060***	0.068***	0.061***	0.061***	0.060***
	(0.019)	(0.019)	(0.020)	(0.021)	(0.021)	(0.021)
Immediately Before Disaster (β_{pre})	-0.005	-0.004	0.006	0.010	0.010	0.010
	(0.018)	(0.019)	(0.022)	(0.018)	(0.018)	(0.018)
Constant	0.806***					
	(0.009)					
Hypothesis Tests [p-values]:						
H_0: $\beta_{\text{post}} = \beta_{\text{pre}}$	.027	.027	.032	.047	.046	.046
R-Squared	.002	.006	.041	.059	.067	.085
Number of Observations	478,946	478,946	478,946	478,946	478,946	478,946

B. Using Large Foreign and Domestic Natural Disasters	Vote with Special-Interest Donors					
	(1)	(2)	(3)	(4)	(5)	(6)
Immediate Aftermath of Disaster (β_{post})	0.041***	0.044***	0.045***	0.035**	0.035**	0.035**
	(0.015)	(0.015)	(0.017)	(0.017)	(0.017)	(0.017)
Immediately Before Disaster (β_{pre})	0.005	0.006	0.007	0.011	0.011	0.011
	(0.015)	(0.015)	(0.018)	(0.015)	(0.015)	(0.015)
Constant	0.801***					
	(0.009)					
Hypothesis Tests [p-values]:						
H_0: $\beta_{\text{post}} = \beta_{\text{pre}}$	.107	.093	.098	.245	.244	.252
R-Squared	0.002	0.006	0.040	0.058	0.066	0.084
Number of Observations	478,946	478,946	478,946	478,946	478,946	478,946
Fixed Effects (Panels A & B):						
Year × Month	No	Yes	Yes	Yes	Yes	Yes
Day of the Week	No	Yes	Yes	Yes	Yes	Yes
Legislator	No	No	Yes	No	No	No
Legislator × Congress	No	No	No	Yes	Yes	No
Bill	No	No	No	No	Yes	No
Legislator × Bill	No	No	No	No	No	Yes

Notes: Entries replicate Table 2 in the main text using all domestic natural disasters (upper panel) as well as large foreign and domestic disasters (lower panel), as explained in Appendix D.

Appendix Table A.5: Replication of Table 6

A. Using All Domestic Natural Disasters

	Vote "Yea" on Passage			
	(1)	(2)	(3)	(4)
Money from Supporting Interest Groups ($\beta^{(+)}$)	.020***	.018***	-.004	.008**
	(.004)	(.004)	(.003)	(.004)
Money from Opposed Interest Groups ($\beta^{(-)}$)	-.181***	-.174***	-.156***	-.126***
	(.025)	(.025)	(.022)	(.018)
Money from Supporting Interest Groups		.006	.011*	.011*
$\times$ Immediate Aftermath of Disaster ($\gamma^{(+)}$)		(.007)	(.006)	(.007)
Money from Opposing Interest Groups		-.052***	-.044**	-.027*
$\times$ Immediate Aftermath of Disaster ($\gamma^{(+)}$)		(.018)	(.018)	(.014)
Immediate Aftermath of Disaster (δ)	.031**	.033**	.029*	.020
	(.015)	(.016)	(.015)	(.016)
Hypothesis Tests [p-values]:				
$H_0 : \gamma^{(+)} \leq 0$	–	.181	.030	.049
$H_1 : \gamma^{(-)} \geq$	–	.002	.006	.031
$H_2 : \gamma^{(+)} = \gamma^{(-)} = 0$	–	.001	.009	.067
Fixed Effects:				
Legislator $\times$ Congress	No	No	Yes	Yes
Year $\times$ Month	No	No	No	Yes
Day of the Week	No	No	No	Yes
R-Squared	.047	.047	.239	.315
Number of Observations	674,726	674,726	674,726	674,726

B. Using Large Foreign and Domestic Natural Disasters

	Vote "Yea" on Passage			
	(1)	(2)	(3)	(4)
Money from Supporting Interest Groups ($\beta^{(+)}$)	.020***	.019***	-.004	.008**
	(.004)	(.004)	(.003)	(.004)
Money from Opposed Interest Groups ($\beta^{(-)}$)	-.180***	-.173***	-.155***	-.127***
	(.025)	(.025)	(.022)	(.018)
Money from Supporting Interest Groups		.005	.011*	.008
$\times$ Immediate Aftermath of Disaster ($\gamma^{(+)}$)		(.007)	(.006)	(.006)
Money from Opposing Interest Groups		-.052***	-.048***	-.023
$\times$ Immediate Aftermath of Disaster ($\gamma^{(-)}$)		(.019)	(.017)	(.015)
Immediate Aftermath of Disaster (δ)	.030*	.032*	.019	.013
	(.017)	(.018)	(.017)	(.017)
Hypothesis Tests [p-values]:				
$H_0 : \gamma^{(+)} \leq 0$	–	.224	.034	.097
$H_1 : \gamma^{(-)} \geq 0$	–	.004	.003	.056
$H_2 : \gamma^{(+)} = \gamma^{(-)} = 0$	–	.020	.004	.142
Fixed Effects:				
Legislator $\times$ Congress	No	No	Yes	Yes
Year $\times$ Month	No	No	No	Yes
Day of the Week	No	No	No	Yes
R-Squared	.047	.048	.238	.315
Number of Observations	674,726	674,726	674,726	674,726

Notes: Entries replicate Table 6 in the main text using all domestic natural disasters (upper panel) as well as large foreign and domestic disasters (lower panel), as explained in Appendix D.