MULTIDIMENSIONAL MEDIA SLANT: COMPLEMENTARITIES IN NEWS REPORTING BY US NEWSPAPERS

2018

Sandra García-Uribe

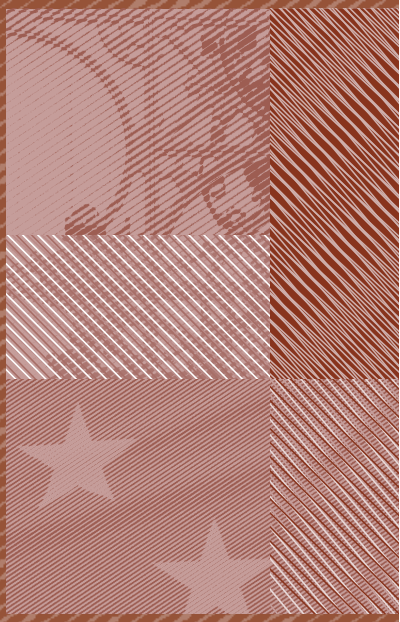

Documentos de Trabajo
N.º 1817

BANCO DE ESPAÑA
Eurosistema

MULTIDIMENSIONAL MEDIA SLANT: COMPLEMENTARITIES IN NEWS REPORTING BY US NEWSPAPERS

MULTIDIMENSIONAL MEDIA SLANT: COMPLEMENTARITIES IN NEWS REPORTING BY US NEWSPAPERS

Sandra García-Uribe (*)

BANCO DE ESPAÑA

(*) I am very grateful to Manuel Arellano for his guidance and encouragement. I am also very grateful to Guillermo Caruana, Rubén Durante, Diego García, Camilo García-Jimeno, Gerard Llobet, Mónica Martínez-Bravo, Pedro Mira, Diego Puga, Enrique Sentana and Pinar Yildirim for their many comments and suggestions. I am also grateful to comments and suggestions by seminar participants at CEMFI, University of Pennsylvania, 41th Simposio de Análisis Económico at Bilbao, Banco de España, Universitat de les Illes Balears and the 3rd Economics of Media Bias Workshop at University of Cologne. I gratefully acknowledge Sabadell INVES 2014 and a Spanish Ministry of Economy and Competition FPI Grant BES-2014-070531 for financial support. The views expressed here are those of the author and do not necessarily reflect those of the Banco de España. All remaining errors are mine. Email: sandra.garcia.uribe@bde.es. Banco de España. Alcalá, 48, 28014 Madrid, Spain. Personal website: http://sites.google.com/view/garciauribe.

Documentos de Trabajo. N.º 1817

2018

The Working Paper Series seeks to disseminate original research in economics and finance. All papers have been anonymously refereed. By publishing these papers, the Banco de España aims to contribute to economic analysis and, in particular, to knowledge of the Spanish economy and its international environment.

The opinions and analyses in the Working Paper Series are the responsibility of the authors and, therefore, do not necessarily coincide with those of the Banco de España or the Eurosystem.

The Banco de España disseminates its main reports and most of its publications via the Internet at the following website: http://www.bde.es.

Reproduction for educational and non-commercial purposes is permitted provided that the source is acknowledged.

© BANCO DE ESPAÑA, Madrid, 2018

ISSN: 1579-8666 (on line)

Abstract

Are editors' choices of front page news based on the potential complementarities between the news items? This paper studies front page choices made by editors of major newspapers in the US. I document that newspapers front pages are biased to certain combinations of news on top of biased to certain news. To identify my measures of bias, I exploit the variation in news relevance across different topics and days. To measure the news relevance I use lead news choices of other US mass media. As a consequence, my measures of bias are relative to the overall media bias. I also provide a reader-maximization model for front page decisions that I use to interpret the empirical biases of the newspaper as preferences of its population of target readers. From my estimation, I recover maps of complementarities among pairs of topics for each of the major US newspapers. I find that complementarities between news contribute in a large portion to the probability that news on a topic appears in the front page.

Keywords: Media bias, discrete choice, complementarities, multiple products.

JEL classification: D22, C25, C55, L82.

Resumen

¿Están las decisiones de portada de los periódicos basadas en las posibles complementariedades entre las noticias? Este trabajo estudia la elección de noticias de portada hecha por los editores de los mayores periódicos impresos de Estados Unidos. Se documenta que las portadas están sesgadas hacia ciertas combinaciones de noticias, más allá de estar sesgadas a cierto tipo de noticias. Para identificar las medidas de sesgo, se explota la variación en la importancia de las noticias por temas y días. La importancia de las noticias se mide utilizando las decisiones de cobertura de noticias en los principales espacios de una amplia muestra de medios de masas en Estados Unidos. Como consecuencia, las medidas de sesgo son relativas al sesgo general de los medios. Además, este trabajo aporta un modelo de maximización de lectores para la decisión de portadas que se utiliza para interpretar las medidas de sesgo empíricas como preferencias de la población de lectores objetivo del periódico. De las estimaciones, se recuperan mapas de complementariedades entre pares de noticias para cada uno de los grandes periódicos estadounidenses. Encuentro que las complementariedades entre noticias contribuyen en gran medida a la probabilidad de que una noticia sobre un tema determinado aparezca en la portada.

Palabras clave: sesgo informativo, elección discreta, complementariedades, productos múltiples.

Códigos JEL: D22, C25, C55, L82.

1 Introduction

Mass Media have a fundamental role in the provision of information to the society in a very literal sense as there is evidence that news affect behavior of political agents (Eisensee and Strömberg, 2007; Durante and Zhuravskaya, Forthcoming), collective action (Hendel et al., 2015) and financial decisions (Fang and Peress, 2009; Garcia, 2013) among others. Little is known about the determinants of the specific choice of lead news, while they receive major public attention every day. Days after Lehman Brothers had collapsed, at the day of the US administration first bail-out proposal, economy news were the most relevant news in the media, followed by political and, far apart, by other topics. However, *The New York Times* decided to devote half of the front page to legal news, while *The Wall Street Journal* had full coverage of economy. There are trade-offs among news characteristics that are relevant for the choice of top news by the media. In this paper, I show that the complementarity or substitution between news items is an important factor to the choice of top news.

This paper studies the choice of front page news made by editors of major newspapers in the US. Front page news not only have a special role in determining the public awareness of events, but they are also a clear-cut observable outcome on a daily basis, and therefore amenable to systematic scrutiny. The literature has so far centered on measures of news slant associated to political ideology along the one-dimensional left-right divide. The basic motivation for such focus has been that information about politics affects political attitudes, and eventually voting behavior and political outcomes. However, the notion of bias may be applied more generally to the choice of lead topic or bundles of topics that a media outlet decides to emphasize among the relevant news on a particular date. Understanding such reporting patterns is important because of its potential effects not only on political attitudes, but more broadly on lifestyles, including values, world views, or health.

During the period 2007 to 2012, the two main front page news for major US newspapers were similar as shown by their relative frequencies of choice in Figures 1-5. Full front page coverage in the US was more likely for foreign and political news. While more common front page combinations were on foreign and political, economy and political, economy and foreign and development and foreign news, I recover measures of bias for each newspaper that reveal differences among their choice of front page news.

In this paper I make three contributions. The first one is to provide an empirical framework for measuring news slant across bundles of topics among major US daily newspapers. I wish to measure the inclination of a newspaper towards or against certain combinations of topics in the choice of front page news relative to the importance of such combination of topics in the media as a whole in a given day. My measure of slant or bias of a newspaper concerning two topics is a double difference average quantity. The first difference measures the bias of the newspaper to a topic and the second difference the bias to the combination of topics. Therefore, this magnitude is a measure of leaning towards a combination of topics net of topic bias. To implement this calculation I use a multinomial logit (ML) model defined over the space of all possible pairs of topics and a specification of the importance of news based on data for the media as a whole.

While my measure of slant could be described as a measure of multidimensional specialization, I shall refer to it as a measure of complementarities. Indeed, the second contribution of the paper is to provide an explicit model under which my empirical measure of multidimensional slant captures the preferences of readers for particular sets of news given the importance of those news in a day. In this model I also assume that newspaper editors choose front page news to maximize readership among a target population of readers, which is given in the short run. In principle, newspaper firms could respond to reader preferences, but it could also be the case that they sacrifice profits in exchange of expressing their ideas. In the context of political ideology, Gentzkow and Shapiro (2010) found that slant is mostly explained by firms responding to consumer preferences rather than the opposite. If this is so in the political arena, a model of reader maximization seems natural for interpreting choice data over broader news topics.

The third contribution of the paper is empirical. Although front pages are observable, there is no obvious way to estimate multidimensional measures of slant. My empirical strategy consists of, first, providing a way to organize news data into elements of the newspaper choice set, and second, analyzing newspaper front page choices in response to different news market scenarios. To identify my measures of bias, I exploit variation in media relevance of news across different topics and days. To construct measures of news relevance I use lead news choices of other mass media. I found in the News Coverage Index Data (NCID) from the Pew Research Center for Journalism an interesting dataset for the purpose that the literature has not exploited. It contains lead news choices from a vast number of US mass media for the period 2007 to 2012.

As an output of the estimation I obtain maps of complementarities among pairs of topics for each of the major five US newspapers. These maps show the extent and the directions in which each newspaper deviates from the aggregate flow of news in the US media. Moreover, under the utility-based interpretation, they also speak about the preferences over topics of the target population of readers associated with each newspaper. The existence of such complementarities imply that the probability that some piece of news in a particular topic makes it to the front page of a newspaper will not only depend on the importance of the news in that particular day and the newspaper bias to the news but also on a cross-effect of the satisfaction of appearing alongside other topics.

I find that *The New York Times* (NYT) complementarities to economy news imply that half of the probability that it publishes on economy news in the front page is due to the satisfaction of combining it alongside other topics. At the average day in the sample, 9.6 out of 18.0 percentage points (pp) of the probability that economic news are published in the front page of the New York Times (NYT) are due to cross-effects with topics other than economics. Something similar happens with political news where 8.5 out of 21.5 pp of the probability of making it to the front page are due to positive cross-effects. Alternatively, the probability that legal news makes it to *The Wall Street Journal* (WSJ) front is half of what it would be if it had not such dissatisfaction of publishing it alongside other topics. The probability of legal news is 3.5 percent, -2.3 pp of it are due to negative cross-effects with other topics. The empirical results also say that at days in which a particular topic is dominant the relative importance of complementarity drops, although complementarities still play an important role.

The estimates of slant for each newspaper and combination of news are robust to a different classification of news, alternative measures of news relevance and to the correlation of news relevance across pairs of topics. Moreover, I perform multiple testing procedures to document the statistical significance of complementarities in news reporting for major US newspapers.

Relation to the Literature. To my knowledge this is the first work that empirically studies the decision of which news go to the front. Previous works on front page decisions are found in sociology and communication literature, such as Reisner (1992) and Clayman and Reisner (1998) that provide a sociological study of news selection through conversational analysis at editorial meetings. Other strand of works have studied the determinants of news coverage at all. In the literature of media and communication, Berkowitz (1990)

uses content analysis in television industry to account for how much news characteristics such as timeliness, significance and story typology matter for news selection. He finds that these factors account for 20 percent of the variation in choices, supporting the relevance of social practices related to the gatekeeping process for news selection.

Sociology studies again focus in editors preferences for characteristics of news, such as Peterson (1979) and Chang and Lee (1992). Most of the attention has been given to the study of whether media choose the most relevant news for the society and which deviations can be found. Originally, Galtung and Ruge (1965) and Sande (1971) hypothesized on a number of "news factors" that they though increase the chance of being perceived and reported. Some of those being unambiguity, frequency, threshold, personal elitism or elite nations but also meaningfulness. Later empirical works gave some evidence on the relevance of these factors for media choices of news. Peterson (1981) analyzes the determinants of international news coverage at the Times comparing covered and non-covered news across different "news factors". Shoemaker et al. (1991) find that events significant to the US are more likely covered in major US television channels. This paper contributes to this strand of the literature, first, by providing an empirical approach to the study of media choices, and second, by giving empirical evidence on the importance of topics and their interactions over and above the relevance of these topics for the choice of lead news.

It has been in the literature of Economics of the Media and the Political Economy of Mass Media where major contributions have been done to the measurement of systematic deviations of media content from objective measures. Groseclose and Milyo (2005) was the first to provide a measure for newspaper political slant comparing citations of think tanks in the texts of newspapers to citations of think tanks by politicians. Gentzkow and Shapiro (2010) measure newspaper democratic-republican leaning using text analysis techniques that account for phrase usage compared to politicians' speech. Other works, such as Di Tella and Franceschelli (2011) measure newspaper coverage of corruption scandals and Branton and Dunaway (2009) study immigration slant in the press. However, all existing empirical work on media bias has focused on one-dimensional measures of slant. My paper contributes by providing a measurement strategy for multi-dimensional slant to this literature. These measures of slant for each newspaper can be used to learn how much (less)more likely is a piece of news on a topic to be leading news because of the

(dis)taste in combining it with other topics. This is the first paper to provide evidence of such slant in the mass media.

My modeling approach to multidimensional slant connects in several aspects with existing models in the literature on complementary choices. The interpretation of bias for a combination of topics as the complementarity in the utility of readers for those topics relates to models that study competition in demand between products. In particular, Gentzkow (2007) studies whether on-line and printed newspapers are substitute or complement goods. In that paper, complementary goods are parameterized as the cross-category transfer of utility between the two products. In the current paper, substitution patterns between news are similarly modeled in terms of cross-topic transfers of utility, thereby transporting these ideas to the the study of multidimensional bias in the media. Moreover, I contribute by extending the applicability of such models for the study of multiple competing products in a marketplace.

Content bundling is pervasive in all media sectors. In this paper, I provide a model for bundles of news, which is consistent with the simultaneous decision among multiple candidates to the front page. The excess of publication of certain bundles of news is related to preferences of the target population of readers. The decision of which target population to choose entails problems of demand and supply characteristics, market structure and advertising which have been studied in works such as Bakos and Brynjolfsson (1999), Koschat and Putsis (2002), Crawford and Yurukoglu (2012), and more recently in Jeon and Menicucci (2012) and Zhou (2017).

My paper also contributes to existing models for news markets by introducing a distinction between media short-run and long-run decisions. Mullainathan and Shleifer (2005) consider a model for unidimensional slant where media outlets face a population of readers and competitors. In this model, slant is a long run decision because the outlet decides which part of the distribution of reader preferences serves, while choosing how much to slant information. In my model, preferences for topics and combination of topics of the population of readers that the newspaper targets is exogenous, part of some ex-ante decision process, and the editor of the newspaper chooses front page news to maximize aggregate readership among this subpopulation of readers. It does so, because it needs to react to the daily flow of news.

To the best of my knowledge, this is the first paper that empirically studies the trade-offs between news relevance and preferences for characteristics of the news under the

same decision framework. Earlier work has considered either one or the other factor for newspaper decisions. For example, Gentzkow and Shapiro (2010) studies the importance of demand preferences for political parties in the newspaper choice of political slant. Sen and Yildirim (2015) studies how much popularity of a story matters for the overall online coverage of that story in the newspaper. In this paper I can distinguish between preference for topics and preference for relevance of topics within the newspaper by comparing days with different number of news for each topic in the media and the repeated choice of the newspaper. Evidence on these forces is also a test for the market behavior of the mass media.

The structure of the paper is as follows. I introduce the measure of multidimensional slant in Section 2. In Section 3, I describe the dataset and the construction of measures of news relevance. Section 4 deals with the empirical methodology, including a discussion about the identification of parameters of interest. In Section 5, I define the decision problem for newspaper choice of front page news and discuss a set of assumptions under which the theoretical model matches the empirical specification; thus, I offering a utility-based interpretation for the evidence on complementary news in the data. Section 6 contains the full set of empirical results and Section 7 concludes.

2 Multidimensional Slant

2.1 Slant across topics

In this work I am concerned with a notion of slant in the choice of bundles of lead news. Existing empirical work concentrated on measures of slant representing political ideology on the one-dimensional left-right divide. They have considered the choice of phrases and events reported in a newspaper discretized in a for/against reference category. Instead I focus on media choices of lead news from a set of different topics, such as political, economy or disaster news.

I wish to show that on top of news outlets being biased to certain topics and to the importance of news on a particular topic an additional force for the choice of lead news is a (dis)taste for bundles of topics. This force provokes an excess(scarcity) of publication of topics which I call multidimensional slant.

2.2 Measurement Approach

In this section I detail my measure of slant. I classify front pages between those with two main pieces of news and those with one main piece of news. I proceed this way for simplicity and because most front pages can be accommodated in this way. My measure of slant or bias of newspaper n concerning two topics ,j, s, Γ^n_{js}, can be thought of as the double difference average quantity,

$$\Gamma^n_{js} = \frac{1}{2} E[\sum_{k=j,s} [(lnp^n_{js,t} - ln\pi_{js,t}) - (lnp^n_{kk,t} - ln\pi_{kk,t})]] \quad (1)$$

where $p^n_{js,t}$ is the probability of publishing front news j and s at day t, and $\pi_{js,t}$ rates the importance of news js at day t. The expectation is defined over multiple different days t. Hence, I take the expectation of double differences over multiple days. The difference between $lnp^n_{kk,t}$ and $ln\pi_{kk,t}$ is a measure of newspaper n bias to topic k at day t. If the newspaper had no bias for news k, this difference is zero. Therefore, my double difference magnitude is a measure of leaning towards topics js net of unidimensional bias. One could compute the $p^n_{js,t}$ non-parametrically from the observation of newspaper choices of js at days like t. The news importance of js at day t could also be computed non-parametrically as the relative importance of news js at days like day t. Therefore, the different elements of this formula are totally unrestricted. The one half in the formula states for the simplifying assumption that the two main news in the front page bundle weight the same to net out unidimensional bias to each topic[1]. In Section 4, I choose a logistic specification for $p^n_{js,t}$ for all js. In the robustness checks, I argue on alternative measures for $\pi_{js,t}$ for all js.

The measure of slant gives as an output an index for each newspaper and bundle of news that can take values in the space of real numbers. The index accounts for the excess frequency in the publication of bundles of topics relative to the overall media relevance of those news and the bias of the newspaper to those separate topics. Gentzkow and Shapiro (2010) measures local newspapers political slant by computing the slope of a regression of a newspaper's frequency of politically differentiating phrases on the frequency of usage of

[1] One could think about a different weighting scheme for the elements in the bundle of front page news by incorporating features into the choice, such as the space devoted or the order of news. We do not dispose of the front page space of news and, although we know the order of news, a larger time span would be required to study other weighting schemes.

those phrases for all congressmen of the Republican or Democrat party. With this strategy they can net out the common factor for both parties and newspapers. In the sense that the measure allows for a comparison between newspapers in Republican versus Democrat speech metric, the Gentzkow and Shapiro (2010) measure can also be considered as a relative measure of slant.

In contrast, my measure is relative to the overall media relevance of the news. Groseclose and Milyo (2005) also provide a relative measure of newspaper political slant, which in this case, is based on think tank citations. Di Tella and Franceschelli (2011) relate political corruption coverage to advertising of the political party in the newspaper, producing a measure of bias relative to political advertising in the newspaper.

3 Data

I use data on US mass media choices from the News Coverage Index Data (NCID) from Pew Research Center for Journalism[2] for the period January 2007 to June 2012. This dataset contains hand-coded news for a large number of media outlets in the US, i.e. printed newspapers, on-line news sites, network tv, cable tv and radio. It identifies events covered by country-relevant news outlets. It samples from media sources in proportion to the sector and program/issue rating/readership. The data was coded on a daily basis and, for feasibility reasons, it rotates outlets each day. We observe The New York Times every day, and Washington Post, USA Today, Wall Street Journal and Los Angeles Times every two days. NCID was collected to construct a measure of US media news agenda but I exploit the time series dimension to study major newspaper biases.

Table 1 provides descriptive statistics about the sample coverage of media outlets. On average there are 26 different media outlets in the sample every day. The average share of broadcast news in the sample is 52 percent. The sectoral composition reveals on average 30.1 percent of the news in the sample are from newspapers, 23.8 percent from online newspapers, 15.6 percent from network tv, 25 percent from cable tv and 9.4 percent from radio. This news sources vary from day to day due to the rotation of outlets and reduced sampling at special days, but as shown by standard deviations this variation is small.

[2]Pew Research Center (http://www.pewresearch.org/) is the source of the data and the Center bears no responsibility for the interpretations presented or conclusions reached based on analysis of the data.

From this dataset I extract two types of information. I extract data from main front page news choices for the five major US printed newspapers, which are my dependent variables. To construct measures of news relevance at particular days, which are independent variables in the empirical analysis, I use the lead news choices made by all the media outlets in the sample. Note the reader that these choices are not only for printed newspapers or for the two main news in the front page but also for any media outlet in the sample and any number of lead news published, e.g. for television lead news are the main 30-minute segment of news.

3.1 Top News

We observe the leading news for all the sampled media. For newspapers, it amounts to front page news with information on the position of each piece of news within the front page. For broadcasts, they coded from fifteen minutes to the full emission depending on its relevance. Mostly news programs are in the sample, although for radio and cable tv we observe some shows. The dataset offers three types of story classifications, broad topic, big story and sub-story, that go from less to more transitory characteristics of the news. Due to sample size restrictions I decide to focus on the more aggregate classification that I describe in Table 16. There were 26 broad topics in the original classification which describe general topics in the news and I aggregate them to 8 categories by topic affinity as shown in the table. This classification is sufficiently large to exploit different socio-economic characteristics across events and narrow enough so that the categories do not become time varying. As a robustness check, I provide results for a classification of 21 topics as documented in Section 6.2.In Table 2, I offer the list of final topics jointly with the relative presence as main front page news at newspapers and also at the rest of the media by sectors. We can compare aggregate differences across media sector in publication of content. Cable tv publishes a 31.5% of their news on politics, followed by newspapers who do it on 20.74% and radio with 20.57%.

A first challenge that this paper faces is which is the unit of analysis we want to look at and whether it is feasible for that unit. News have a strong time-varying component, while looking at these features is interesting, this approach will make the number of potential choices considerably large. Moreover, the interpretation of a model with time varying choices will be difficult. Other approach could have been to allow time varying choices in a model of characteristics and specify parameters for a set of characteristics. In this

sense, the choices are "anonymous", i.e. they refer to the set of characteristics, this fact complicates the interpretation of the results for an analysis of complementary news.

3.2 News Relevance

In order to estimate my measure of slant I need to construct measures of news relevance. The relevant frequency for the analysis is the day, since printed newspapers are produced at this frequency. Every day there are news shocks that change the importance of topics in the media which I propose to capture in measures of news relevance.

Using this dataset we can construct several measures exploiting the number of news covering each topic at each date. The most basic measure I implement accounts for the total number of publications about a topic at different days. Summary statistics for this key explanatory variable are presented in Table 3. The table shows there is considerable variation in the measures of news relevance across days. For example, the relevance of economic news is on average 41 news per day, while the standard deviation is 23 news, ranging from 2 to 114 news, the median is 34 news, the interquartile range 30 news and the total number of day-outlet-news observations on economic news during the sample period is 45,151 news while for any topic is 326,007.

The basic measures of news relevance I propose capture actual major events in their respective topics. In Figure 6, I depict the series for economy news and that of the Dow Jones Industrial Average Index. These two indices are negatively correlated reflecting the fact that the media tends to publish more on negative economic events than on positive news, something that has been tested in the literature about the interaction of financial markets and the media (Garcia, 2013, 2014). Important dates such as the Lehman Brothers collapse that maybe the largest economic scandal in this period appears as the major inflection point. In Figure 7, I compare the series of political news relevance aggregated to the week jointly with the Google Trends index for the term "political" in the US in the same period. We observe that these two indices are positively correlated specially at major events. In this case important political events in the US such as the Presidencial Elections (PE) of 2008 or House of Representative Elections (HE) of 2010 are also captured by the measure of political news relevance.

Given the variation in the number of outlets sampled in a day and the number of news released on particular days, I propose two different measures of news relevance other than topic counts. One is the measure of topic share which accounts for the share of news on

a topic over the total number of publications in the sample at a given day. This measure corrects for the possibility that one day there are more news on a topic while not relatively more than other day, which depends on the number of outlets sampled and the number of lead news published per source at a particular day. Let X_{jt} the number of news on topic j at day t and X_t the number of lead news in the sample at day t:

$$topicshare_{jt} = \frac{X_{jt}}{X_t} \tag{2}$$

Summary statistics for this measure are in Table 4. The relevance of economy news is on average 18 percent of the total news per day, while the standard deviation is 0.09 percent, ranging from 1.1 percent to 46.2 percent, the median is 15.8 percent, the inter-quantile range 12 percent and the total number of days where there are observations of economy news is 1,414. The other measure I propose is one I call topic newshole. This measure combines the share of minutes and the share of words a news on a topic is covered using the sample shares of broadcast publications relative to written publications, hence weighting the the two types of coverage by the relevance of broadcast coverage at each day. This measure also adds up to one everyday. I statistically describe this measure in Table 5. Let M_{jt} be amount of minutes that broadcast news on topic j are lead news at day t, W_{jt} the amount of words that written news of topic j are lead news at day t, let B_t the share of broadcast news items at day t:

$$topicnewshole_{jt} = \frac{M_{jt}}{M_t} B_t + \frac{W_{jt}}{W_t}(1 - B_t) \tag{3}$$

The measure of news relevance that I offer may not coincide with potential social planner measures of news relevance. The reader should be aware that the measure of news relevance that I present is one that represents the overall mass media coverage of the news at any given date. We can think of this measure of news relevance as the aggregate expected preference for news of the US mass media. Preferences for news characteristics will be measured with respect to the mass media aggregate preference for news. The measures of bias are going to be relative to this benchmark, hence, we cannot say anything about "unbiased" or "objective" media in a general sense.

4 Empirical Strategy

To identify my measures of slant, I exploit repeated-choice data from media outlets lead news choices across different days. The market of general news is noisy, every day there are shocks to the ordering of news relevance due to the events that take place around the globe according to some exogenous random process. In this paper I profit from the natural variation in the data generating process of general news to measure newspaper slant. Specifically, I rely on the variation of the news relevance measure of topics in the mass media across different days to identify any systematic deviation in the front page choices for each newspaper and bundle of news.

To illustrate the information obtained in this type of data variation let me pose the following example. Let there be news on three potential topics to publish in the front page, e.g. political, economy and legal. Let their values to the the general public be, x_p, x_e, x_l, respectively. If the newspaper chooses political and economy news, by revealed preference, it must be that the utility of that combination, u_{pe} is at least as high as that of any other alternative u_a, for any other alternative a. If we observe sufficiently many different market scenarios (general public valuations), we can effectively exploit the information revealed in the repeated choices.

To implement the estimation of multidimensional slant I estimate a discrete choice model using repeated choice data. By revealed preference I claim I obtain the parameters of interest that account for such measures and I detail in the following section.

4.1 Empirical Model

I specify a discrete choice model for the newspaper decision of front page news. These top news are bundles of up to two different elements. I classify front pages between those with two main pieces of news and those with one main piece of news. I proceed in this way for simplicity and also because observationally most front pages can be comfortably classified in this way. However, my framework could be extended to consideration of front pages with more than two news items, and to consideration of a finer classification of the importance of news in a newspaper. Everyday, the choice set is then $A = \{(a_1, a_2), \forall a_i = \{1, ..., J\}\}$. I specify the log odds probability ratio for elements in the choice set, js, bundles of two different news, and, jj, bundles with one piece of news, as,

$$ln\frac{p^n_{js,t}}{p^n_{bb,t}} = \kappa\delta_j + (1-\kappa)\delta_s + \Gamma_{js} + \kappa\beta_j X_{j,t} + (1-\kappa)\beta_s X_{s,t}$$
$$ln\frac{p^n_{jj,t}}{p^n_{bb,t}} = \delta_j + \beta_j X_{j,t} \quad (4)$$

where $p^n_{jj,t}$ is the probability of newspaper n choosing a bundle only including news j at day t, $p^n_{js,t}$ is the probability of newspaper n choosing a bundle including news j and s at day t. $p^n_{bb,t}$ is the probability of newspaper n choosing the base bundle choice bb at day t. The log odds of choosing bundle js at day t is composed of δ_j, the topic specific bias, weighted by the value of κ. Γ_{js} is the bundle-specific bias parameter net of topic bias, β_j is the marginal value in utility of one extra unit of news relevance of topic j, X_{jt}. If $\kappa = 0.5$ this model exactly estimates the measure of multidimensional slant of equation 1. Due to data limitation, I do not estimate κ, so I assume it takes value 0.5, but in this model one could identify κ with a sufficiently large sample. If I assume that choices are independent one day from another, I can estimate the standard ML.

As long as the researcher disposes of a large time series on choices for the newspaper and regularity conditions are satisfied, all the parameters of model (4) are identified. However, the researcher disposes of a limited amount of data and that makes her take a specific solution to obtain some parameter estimates which are discussed in the Appendix A. As in any discrete choice model, we need to normalize the value of some alternative to zero, which I choose to be foreign news. The empirical strategy relies on the fact that readership does not vary substantially in the period of analysis, F is fixed across dates. The short term framework allows us to avoid the interference of dynamic strategies that aim at readership building.

5 Economic Decision Framework

In this section, I develop a model for the newspaper choice of top news, in the context of my empirical setting, pursuing an economic interpretation for the measures of slant. I make the distinction between media short-run problems and long-run decisions. In the long-run media decide on their product strategy so that this maximizes the interest of the company as a whole, i.e. what is the target readership, the owner, the editors and other production infrastructures. The long-run strategy defines the target readership of the newspaper, which is characterized by being profitable in the market to an owner. A

specific target readership is profitable if it makes advertising contracts, raises subscribers and sells units. Some owners may demand more monetary profitability than others, i.e. some owners of media may be as happier making less profits if in exchange they can express their ideas.

Under the short-run framework of decisions, newspaper editors are going to take target readership as given. Gentzkow and Shapiro (2010) show that US newspapers choice of ideological slant in the news strongly relates to demand ideology as opposed to ownership preferences. Motivated by their results and the fact that I study major commercial newspapers, I assume that newspaper maximize aggregate target readership in order to choose front page news. In the short run framework, the newspaper profit is the amount of readership captured by the publication of the front page, given that marginal costs are known to be close to zero in this industry. Target readership contains targeted subscribers and single-unit readers weighted by their profitability to the newspaper. Parts of the readership distribution with more mass are more valuable to the newspaper.

The literature on Media Economics has recognized that readers enjoy reading news with information bias that is closer to their beliefs. Gentzkow and Shapiro (2010) and Mullainathan and Shleifer (2005) both work under the paradigm of reader preferences for confirmation beliefs. These preferences are explained by motives such as delegation, psychological utility and reputation (Gentzkow et al. (2014)). Preferences for like-minded news could be generalized to the scenario of readers deriving utility from reading news on topics that are closer to their tastes. If this is the case, a newspaper that maximizes quantities will try to satisfy as many readers as possible and it will carry this strategy through the publication of content that suits better the preferences of their target readers.

On the other hand, newspapers confront the changing media landscape facing an evolving trade-off between story relevance and what readers generally like to read. The readers may want to be informed about the true state of nature which is defined by the events that took place in a particular day. Yet, they may be more satisfied with an emphasis on certain events than others because of their proximity to their lives or minds. Not only that, but also they may prefer reading certain combinations of events together more than others. Events can be classified according to multiple characteristics that do not restrict to politics. One can find in the media a variety of news such as crime, terrorism, religion, legal system, sports, business and so on. For sure these news can have certain political tone, but given the political tone, they can be associated to different

socio-economic interests. Some of them may have more correlated characteristics to the tastes of some readers than others.

5.1 Top News Selection Model

I index time with t and omit the newspaper index to avoid redundant notation but the reader should keep in mind that all parameters are newspaper-specific. Let us model the choice of the main part of the front page of a newspaper as a bi-dimensional vector of news, $a = (a_1, a_2)$. Thus, the newspaper choice set is $A = \{(a_1, a_2), \forall a_j = \{1, ..., J\}\}$, where J is the total number of different news available in the media. Let us note a reader r , $r \in R = \{1, ..., R\}$. Let the population of readers that the newspaper targets in the short-run has a distribution F. Readers face top news $a = (a_1, a_2)$ and decide whether they read the newspaper and enjoy utility u^r_{at} or enjoy their outside option with utility u^r_{0t}. The choice set of a reader r is $\Upsilon = \{1, 0\}$, where 1 indicates reading the newspaper and 0 not reading the newspaper.

Without loss of generality, the utility of reader r for bundle a is the weighted sum of standalone utilities of the elements in the bundle plus a potential complementarity between those elements in the bundle, Γ^r_a and ν_{at} is an unobserved homogeneous shock to readers valuation of bundles in a given day.

$$\tilde{u}^r_{at} = \kappa^r u^r_{a_1 t} + (1 - \kappa^r) u^r_{a_2 t} + \Gamma^r_a 1(a_1 \neq a_2) + \nu_{at} \tag{5}$$

this implies that $\Gamma^r_a = u^r_{at} - \kappa^r u^r_{a_1 t} - (1 - \kappa^r) u^r_{a_2} - \nu_{at} - u^r_0$ for all $a = (j, k)$ where $j \neq k$. Thus $\Gamma^r_a = 0$ for all $a = (j, j)$ for $j = \{1, ..., J\}$, implying there is no additional value in reading two top news of the same type other than the sum of standalone utilities in the bundle. We specify the standalone utility, u^r_{jt}, for $a_i = j$,

$$u^r_{jt} = \delta^r_j + \beta^r_j X_{jt} \tag{6}$$

where δ^r_j is a reader-specific value for news of type j and β^r_j is a reader-specific valuation of each unit of news relevance, X_{jt}, for news of type j. Reader r will read top news a if and only if $\tilde{u}^r_{at} \geq u^r_{0t}$.

The newspaper problem consists of choosing the bundle of top news that maximizes aggregate readership. The potential aggregate target readership for each choice is defined over all the readers in the population that the newspaper target whose preferences are

such that they would prefer reading the newspaper. Among the population of readers that are relevant to the newspaper, F, a measure of the aggregate quantity of readers that prefer a to the outside option is given by:

$$P(a|X_{at}) = \int 1(\tilde{u}^r_{at} \geq u^r_{0t}) dF(u^r_{at}; X_{at}) \tag{7}$$

where F is the distribution of preferences of the readers in the target of the newspaper. F is allowed to differ across newspapers, reflecting the fact that in the short-term each newspaper will weight different type of readers differently, e.g. subscribers and single-unit type or readers in Utah and readers in NY for different newspapers. This model does not consider the endogeneity of F, which will matter from the point of view of the newspaper long-run policy. I make a distinction between the short-run market share optimization and the long-run, which is taken as given. Nevertheless, F could be allowed to change exogenously in the empirical analysis if the time-span is sufficiently large. The newspaper optimization problem consists of choosing top stories $a \in A$ that produce the largest aggregate readership among the potential readership figures:

$$\max_{a \in A} \{P(a|X_a)\} \tag{8}$$

The estimates that come from the discrete choice model of front page news (4) reveal features of the distribution of preferences of the population of readers that each newspaper targets in the short run, that is, features of F. However, these are going to be features of the aggregated target readership. In the following subsections I make use of aggregation theory to establish the assumptions on F that allow us to interpret the estimates of (4) in terms of a explicit characterization of target readers for each newspaper.

5.1.1 Distributional Assumptions

In the following lines, I state assumptions for the target readership that if satisfied lead to model (4). Let define the random variable $W^r_{jk,t}$ for bundle $a = (j, k)$, $\forall t \in \{1, ..., T\}$, and for a reader r as

$$W^r_{jk,t} = \kappa^r(\delta^r_j + \beta^r_j X_{jt}) + (1 - \kappa^r)(\delta^r_k + \beta^r_k X_{kt}) + \Gamma^r_{jk} + \nu_{jkt} - u^r_{0t} \tag{9}$$

and the vector of parameters $\zeta^r = (\kappa^r, (\beta^r_j)_{j=1,...,J}, (\delta^r_j)_{j=1,...,J}, (\Gamma^r_a)_{a=1,...,A}, (u^r_{0t})_{t=1,...,T})$. I allowed the outside option to take different values for different days, e.g. there are days where it is relatively more valuable to read the newspaper than others. The multivariate

distribution of ζ^r defines the target readership distribution. For any distribution of ζ^r of a newspaper, F, we can compute aggregate readership for any choice $a \in A$ as

$$P(a|X_{at}) = Pr(W^r_{at}(\zeta^r, X_{at}) > 0) \tag{10}$$

If $\zeta^r \sim F(\bar{\zeta}, \Omega_\zeta)$ is distributed Normal, the newspaper objective function can be rewritten as a linear utility index for each choice, $Z_{at} = \frac{\bar{W_{at}}}{\sigma_{W_{at}}}$ where

$$\begin{aligned} \bar{W_{jk,t}} = \bar{\kappa}(\bar{\delta}_j + \bar{\beta}_j X_{jt}) + (1-\bar{\kappa})(\bar{\delta}_k + \bar{\beta}_k X_{kt}) + \bar{\Gamma}_{jk} + \nu_{jk,t} - \bar{u}_{0t} \\ \sigma^2_{W_{jk,t}}(X_t) = Var(W^r_{jk,t}) \end{aligned} \tag{11}$$

The variance of each alternative's utility index, $\sigma^2_{W_{js,t}}(X_t)$, is homoskedastic if we impose that $(\beta^r_j)_{j=1,...,J} = (\beta_j)_{j=1,...,J}$ and $\kappa^r = \kappa \ \forall r$,

$$\begin{aligned} \bar{W_{jk,t}} = \kappa(\bar{\delta}_j + \beta_j X_{jt}) + (1-\kappa)(\bar{\delta}_k + \beta_k X_{kt}) + \bar{\Gamma}_{jk} + \nu_{jk,t} - \bar{u}_{0t} \\ \sigma^2_{W_{jk,t}} = Var(\kappa\delta^r_j + (1-\kappa)\delta^r_k + \Gamma^r_{jk} - u^r_{0t}) \end{aligned} \tag{12}$$

A final assumption that we have to make for the interpretation of these estimates is that $\sigma^2_{W_{at}} = \sigma^2$, $\forall a \in A$, $\forall t \in \{1, ..., T\}$, that is, the variance of the utility index is constant across alternatives. This assumption does not limit Ω_ζ to be diagonal but it bounds the variance of each alternative utility index $\sigma^2_{W_{at}}$ to a constant that must be the same for all bundles in the choice set of the newspaper. This also restricts the variance of u^r_{0t} to be constant across time.

$$\begin{aligned} \bar{W_{jk,t}} = \kappa(\bar{\delta}_j + \beta_j X_{jt}) + (1-\kappa)(\bar{\delta}_k + \beta_k X_{kt}) + \bar{\Gamma}_{jk} + \nu_{jk,t} - \bar{u}_{0t} \\ \sigma_{W_{jk,t}} = \sigma \end{aligned} \tag{13}$$

Let define $\nu^*_{jk,t} = \nu^n_{jk,t} - \bar{u}^n_{0t}$, if $\nu^{n*}_{jk,t}$ is i.i.d Type I Error over alternatives and days, and uncorrelated to X_t. Under this assumptions and if regular conditions are satisfied ML estimates of the uni and multidimensional slant of model (13) are interpretable in terms of preferences of the target readership.

5.2 Interpretation of Model Parameters and Discussion

In a short-term framework of decisions, the newspaper caters the representative reader. Under model (9) and preferences for readers of section 5.1.1 the parameters in Θ^n can be interpreted as those of the average target reader for newspaper n. The normality assumption for the distribution of target readership preferences jointly with the assumption

on constant variance across alternatives implies that the newspaper likes the mean of the alternative specific readership.

The literature has been questioning that some media content decisions are pretty affected by the political ideology of the owners of the media outlet. We can obviate this interpretation for bias estimates because, under the short-run framework, target readership is set ex-ante. If target readership was chosen to suit an owner's ideology the representative reader and the owner preferences would coincide.

6 Empirical Results

6.1 Main Results

Table 8 and Figures 8-12 present the results from the ML estimation of model (4) as described in Section 4 for the five major newspapers in the sample. For this estimations I normalize the scale of the model to $\sigma = 1$ and set the base category is full coverage of foreign news. I estimate independent specifications for each newspaper. The value $\kappa = \frac{1}{2}$ implies that the newspaper objective function gives equal weight to topic in any of the two positions. To implement estimation I filter the data by days with at least fifteen news in the sample and those topics that are chosen more than three times by each major newspaper across the sample period.

In Figures 8-12, I present new empirical evidence on multidimensional slant. Each figure contains point estimates for the measures of multidimensional slant for each newspaper that I call complementarities between news motivated by the model in Section 5. Each of the figures represent the position of one newspaper in the map of potentially complementary news. Each cell presents the point estimate of the complementarity for each combination of topics, where the x-axis and y-axis label one of each of the elements of the combination. White cells are non-identified coefficients, however those in the right-lower triangular part are not identified because the model the order of the two main news is not modelled in this specification of front page choices. The scale of colors indicate the magnitude of estimates: hotter colors, such as strong red, state for larger estimates (or complementary pairs of news); colder colors, such as the dark blue, state for smaller coefficients (or substitute pair of news). Only those parameters that are statistically significant at 5% confidence level using point wise standard errors are presented. The Appendix covers the corresponding results using a 10% confidence level.

There are some patterns in the results that are relevant to the choice of lead news. The New York Times is the newspaper that has more statistically significant complementarities between news than any other newspaper. There are patent differences between newspapers. Strong complementarity between economy and political for New York Times and Washington Post while not the case for the rest of newspapers. Strong complementarity between disasters and political in the New York Times while not the case for the rest of newspapers. At the Wall Street Journal we also find strong complementarity between foreign and economy news but a substitution between legal and political news. At the 10% confidence level more complementaries are statistically significant for the Wall Street Journal, the USA Today and Los Angeles Times, however there is no gain for the Washington Post or New York Times, results are presented in Figures C.16-C.20 of the Appendix.

Not only these coefficients are statistically significant but they are also economically meaningful in terms of choice frequency, in Figures 13 to 17. The most significant combination of news for the NYT, political and disaster, makes the bundle frequency more than 2 times different. One of the less significant coefficients for the NYT, foreign and political, makes it approximately 1 time different.

In addition to the evidence on multidimensional slant that I recover, the empirical results provide insights on various factors about the newspaper problem of choosing lead news. In Table 8, I present the estimates for the unidimensional bias parameters and the sensitivity to topic relevance. Unidimensional biases are associated to the model parameters that represent the average reader tastes for particular topics in the newspaper. Given the base category is foreign news, the New York Times choices reveal a preference for foreign, political and developement news. The Washington Post has preference for foreign news followed closely by political news and then development news, being Legal the less preferred topic. The Wall Street Journal has a strong bias to economy, foreign and legal news. The USA Today has a slight preference for development news. The Los Angeles Times is also biased to foreign and political news. The sensitivities to news relevance, β_j, are statistically significant for the majority of topics which, on the one side provides evidence on the newspapers attention to overall media relevance for the choice of front page news and, second, provides empirical support to the measures of news relevance that this paper exploits as an important determinant for major newspapers choice of front page news.

6.2 Robustness Checks

6.2.1 Alternative Controls

In this section, I present results from an alternative specification of model (4). In Table 9 and Figures C.1-C.5 I show the estimates of the following extended model,

$$\begin{aligned} u_{at} = \kappa u_{a_1 t} + (1-\kappa) u_{a_2 t} + \Gamma_a 1(a_1 \neq a_2) + \gamma \langle X_{at} \rangle + \epsilon_{at} \\ u_{jt} = \delta_j + \beta_j X_{jt} \end{aligned} \tag{14}$$

where ϵ_{at} is also distributed i.i.d Type I Error $\forall a, t$. This specification considers a potential additional factor for front page choices, the interaction of market relevance of news within the bundle, $\langle X_{at} \rangle$. This results are robust to the concern about complementarities capturing comovements in the market relevance of topics across different dates, e.g. there are more news about politics and legal because they are related through some unobservable event, like a corruption scandal. Results are robust to the topics relevance comovements but we can observe a reduction in the number of statistically significant complementarities for the WSJ and USA Today.

6.2.2 Alternative Measures of News Relevance

Proper measurement of news relevance is key to correctly interpret the results. The aim of this section is to test the robustness of baseline results to alternative measures of new relevance that we can derive from this data. Baseline results are computed using the daily topic distribution but I provide alternative measures of news relevance in Section 3.2.

I present evidence for the baseline model using the two alternative measures, topic newshole and topic publications. The results for complementarities are presented in Figures C.6 to C.10 of the Appendix for topic newshole and Figures C.11 to C.15 for total publications and those for unidimensional bias and sensitivity to news relevance in Table B.1 and B.2 of the Appendix. Under this two alternative measures we find more statistically significant complementarities for the USA Today and the WSJ.

Since one concern is that results may be sensitive to the fact that newspapers are more affected by past than contemporaneous news relevance, I perform a robustness

check consiting of using lagged news relevance[3]. I find that the results do not change qualitatively if we use lagged media relevance as shown in Figures C.21-C.25.

6.2.3 Alternative Classification of News

Do complementary news exist or the results depend on the specific news classification we exploit? I tried a finer classification of news in the same dataset, one with 21 topics instead of 8, and many complementarities are still significant[4]. Finer classifications require more a larger time series because one has to compute measures of news relevance for more topics at different dates and estimate more topic specific parameters and complementarity parameters. At Table 16 the reader can see the relation of each 8-category topic to the 21-category list. One can observe that complementarities that are significant with 21 topics are still significant with the 8 topic classification, e.g. the Wall Street Journal substitution of legal and economy news is mapped to the substitution of crime and business using the finner classification. However, we gain some other significant complementarity in the 8-topic classification that with 21 it was not statistically significant. Results are presented in Figures C.26 to C.30.

A second line of concern is whether the Pew Research classification of topics is accurate or there were substantial mistakes in classifying news items. The data producers performed reliability checks for several variables that are published in their methodology document, in particular, for the broad topic classification, which is the classification I am using, the level of agreement in the classification was above 80%. The construction of a sample of individual lead news articles in the mass media using text analysis was preliminary explored. Undertaking a project of this magnitude would be of interest but is outside the scope of this paper.

Other answer to these concerns would be to go to original news articles and create an original topic classification using automatic text analysis and statistical learning algorithms to classify text such as k-means or LDA. I have explored that solution but there are

[3] A minor limitation to test this is that the sample does not contain Sundays and some particular days had no coding so we used the last available measure of media relevance which may not necessarily be the former day

[4] The original list of topics consisted of 26 categories but due to few observations for certain of them I aggregate to 21 topics, the ones I merged to obtain this classification were for one group, transport and development, and for other group sports, lifestyle, additional domestic affairs, media and miscellaneous.

two limitations: identification of news that appear exactly at each part of the front page is not trivial from digital data sources, one would have to use printed sources; collecting a time series of mass media lead news coverage available to textual analysis is a project of an enormous scale. Undertaking the latter project, for the US or any other markets too, is of interest not only as a test to the results found in this paper but also as a future venue of research projects on media economics and the application of media data for economics.

6.3 Joint Significance of Complementarities

I test the robustness of the results to potential proliferation of parameters given their combinatorial production. In Table 10 I present the results of a joint significance test for the bulk of the complementarity parameters of model 4 for each of the five US national newspapers. The test statistic is the likelihood ratio test of the unconstrained model, which is the baseline model where all complementarity parameters are present, and the constrained model, which is one with no complementarity between news. Thus we test the null hypothesis that all complementarity parameters are zero against the alternative that at least one of them is statistically different from zero.

The results show that New York Times, Washington Post and the Wall Street Journal have significant complementarities between news, while for the USA Today and the Los Angeles Times I do not find that the model with complementary news is significantly different from that of no complementarity news. These results are robust to the consideration of market relevance interaction as shown in Table 11 and alternative measures of news relevance such as topic newshole or total number of publications per topic as presented in Appendix Tables B.3 and B.4. I also provide evidence that results are robust to the use of past media relevance instead of contemporaneous media relevance in Appendix Table B.5.

This is evidence of complementarities between news being significant for decisions on top news of several massive circulation newspapers in the US. Complementarities are not statistically significantly different from zero for USA Today, which tends not to offer more than one top news in the main front page, and Los Angeles Times.

6.4 Multiple Testing of Complementarity Parameters

This paper is concerned with multidimensional bias in the media. The complementarity parameters for bundles of news capture multiple biases. So far the results have provided pointwise estimates of each complementarity parameter. The standard practise is to test individual hypotheses at a usual nominal rate, 1% or 5%. However, the probability of rejecting a null hypothesis increases with the number of hypotheses to test at the same time. To account for the multiplicity of tests one has to control for the appropriate error rate, see Romano et al. (2010) for further details.

There are a number of multiple testing procedures (MTPs) which try to adjust different error rate formulas and implement different methodologies. In this paper I present the results of the single individual MTP of Bonferroni (1936), the stepwise individual test of Holm (1979) (Bonferroni-Holm) and the dependence-control stepwise test of Romano and Wolf (2005) applied to the estimates of multidimensional bias in US frontpages. The first two are directly implementable using the estimated p-values and adjusting the significance level by the number of hypotheses to test in a single or a stepwise method respectively. Romano and Wolf (2005) asymptotically controls for the familiwise error rate, it is more powerful than single-step methods and it often will reject more false hypotheses; in exchange it requires bootstrap samples of the estimates of the parameters and the construction of a critical value to generate a rejection region that adjusts for the multiplicity of tests.

I present all the results in Figures 18 to 20. The figures are set as those showing the baseline estimates of complementarity parameters but they show which complementarities are statistically significant by each MTP. For the NYT and WP many complementarities are rejected to be zero using any of the three MTP. For WSJ only one complementarity coefficient survives. The results confirm that complementarities are statistically significant to the choice of front page news for major US newspapers. However, for the USA Today and the Los Angeles Times there is no complementarity that survives to any MTP using the baseline results. In the Appendix I also offer the results using the other two alternative measures of news relevance and results are in line, except for USA Today where one complementarity survives in the case where the control is topic newshole.

6.5 Economic Implications

The empirical results on multidimensional bias for major US newspapers quantify the (dis)taste for combining some lead news topic with other in the coverage of such topic in the front page. The (dis)taste for combinations make some topics be (less)more covered than what one-dimensional bias and relevance forces would predict. This fact makes the major newspapers stress certain dimensions of the US life more than others.

Tables 13-14 present evidence on the explanatory power of complementary news for the choice of topics in the front page at different news market scenarios. Column (2) contains the contribution of complementary news to each topic predicted marginal probabilities for particular days in the sample using baseline estimation of model (4). Each two-column pair in the table gives results to one of the news relevance scenarios that I describe in Table 12. I show that on the average news day, the predicted probability of foreign news at NYT is 37.7% and 14.5 pp of this amount are explained by complementarities of foreign news to other topics. The probability of political news in the front page is 21.5%, contributing complementarities to other topics by 8.5 pp. For WSJ, the probability of politics in the front is 14.7% but only 2.8 pp due to the complementarities with other topics. On the other hand, the probability of legal news is 3.5% and -2.3 pp is the contribution of negative complementarities to other topics impliying that half of the probability of covering legal in the front page is reduced by the distaste of combining it with other news.

On a day following a notable disaster, such as the climate disaster of the Super Outbreak of April 2011, the predicted probability that disaster makes it to the front page was 16.1% in the NYT and 6.9% in the WSJ, however the contribution of cross-effects to these probabilities is 8.6% for the NYT and -4.8% for the WSJ. This shows that for the NYT the complementarities of disaster news to other news push up the decision to cover disaster in the front page by 53% while for the WSJ the substitution of other news to disaster pushes down the decision to covered it in the front by 40%. In contrast for the USA Today, the probability of disaster news was 30.4% and only -5.7 pp accounted for complementarities.

The 15th September 2008, the day after the Lehman Brother's collapse, the probability of economy news in the front pages was over 40% for any newspaper, however, whereas for the NYT it was 51.2% and 18.3 pp due to complementarities, for the WSJ it was 65.5% and only 3.1 pp due to complementarities. On top of this, the probability of disaster or

legal news in the WSJ is reduced by 1.9 pp and 1.4 pp over 3 and 1.8 respectively in such scenario.

The results also give evidence that some complemetarities are relevant for the probability of certain topics even at days that when the media does not pay attention to them. For example, following the Super Outbreak, the probability that NYT publishes foreign news in the front page is 43%, and 15.1 pp of these are due to complementarities of foreign to other topics.

Finally, I perform an out of sample validation of the model with complementarities and the model without complementarities to compare the prediction power of each of the models based on the mean square error of predicting lead news choices. The model is trained in a random 80% of the original sample and tested in the other 20%. The results are presented in Table 15. The mean square errors of the model of complementarities are smaller than those of the model without complementarities. This evidence supports that the model with complementarities predicts at least as good as the model without complementarities in a random 20% test sample. I did the same using other measures of media relevance and obtain similar conclusions, although we note mean square errors are slightly larger.

7 Conclusion

I provide a new empirical framework that I use to document the presence of multidimensional media bias in the choice of lead news by major US newspapers. Conditional on the market relevance of news and unidimensional bias of newspapers into topics, there is an excess of publication of certain bundles of news. As an implication, I find that the choice probability of news on a topic in a given day is biased due to the (dis)satisfaction of appearing alongside other topics. Moreover, I obtain maps of complementarities among pairs of topics for each newspaper. These maps show that each newspapers is located at a different set of complementary news, providing additional support for the market behavior of these newspapers. After giving an explicit utility-based interpretation of these biases through a model of top news selection I refer to these biases as complementarities between news. The model stresses the short-run framework for these daily decisions of front pages as opposed to long-run decisions which affect other structures such as what is the target readership, which news departments to invest and so on.

To account for these biases I estimate a discrete choice model of front page news using daily choices of the five most read newspapers in the US. I exploit a dataset of US mass media choices of top news to extract two types of information, one is the choices for the five major US printed newspapers, and the second to construct measures of news relevance. The measures of news relevance account for the importance of different news at each date. I perform several robustness checks that my results on multidimensional slant resist. They are robust to alternative measures of news relevance, alternative classification of news and to the correlation of topic relevances. I test for the joint significance of all complementarity parameters and I cannot reject their presence for top news selection. I also implement several multiple testing procedures to provide more evidence of the statistical significance of the results. Using this model we can predict the front page coverage of topics for a given media outlet under specific news market scenarios better than a model without complementarities would do. Furthermore, I find complementarities between news items affect meaningfully to the probability that a topic appears in the front page at a given day.

This is the first paper to document the existence of multidimensional media slant and to test it for newspapers front pages. Furthermore, this paper opens several venues for research. An open question is how to model the market importance of news, I use the equilibrium importance but this may be endogenous to various dimensions of the newspaper decision being interesting from a policy perspective of media content. To empirically answer to this question lower frequency data over a longer time horizon seems necessary in this setup. The potential effects of media competition on the choice of lead news are not studied in this paper; the presence of such effects would affect both importance of a piece of news and the value of different alternatives in a given day for the outlet. It is a challenging question to answer but data on media returns to choices would be required, although it is also an interesting path for future work.

Tables and Figures

Table 1: NCID Sample Characteristics

	Mean	Std.	Median	MIN	Max	Weeks
Media Outlets	26.63	10.03	29	2	36	1695
Broad Topics	19.34	4.77	21	3	26	1695
Share of Broadcast News	0.52	0.23	0.61	0.00	0.83	1695
Sources by news share						
Newspapers	30.15	10.02	28.41	12.65	50	1695
Online Newspapers	23.81	5.10	22.94	11.23	37.35	1413
Network Tv	15.58	2.07	15.77	8.97	28.47	1414
Cable Tv	25.03	2.75	24.76	9.38	32.40	1414
Radio	9.39	1.52	9.55	3.36	19.49	1414

Source: Pew Research Center for Journalism. Note: Each row describes one the following variables: Media Outlets is the number of different outlets that are present each day of the sample; Broad Topics is the number of different broad topics that are coded each day of the sample; Share of Broadcast is the total number of publications that are broadcast as opposed to printed; Sources by news share is the multivariate variable of share of publications per media sector and day, it adds up to one each day of the sample.

Table 2: Topic Frequencies in the US Mass Media Leading News

	Newspaper	Online	Network TV	Cable TV	Radio	Total
Political	20.74	18.35	17.05	31.5	20.57	22.74
Foreign	21.82	32.66	22.91	18.02	18.59	21.62
Economic	17.17	13.64	12.71	10.42	19.35	13.86
Development	16.85	8.06	10.94	7.53	10.42	10.05
Other	8.16	8.81	17.29	9.95	13.69	12.36
Disaster	4.73	6.75	9.3	8.1	7	7.71
Legal	4.33	8.12	6.41	9.43	6.12	7.28
Race	6.21	3.61	3.39	5.04	4.25	4.39

Source: Pew Research Center for Journalism. Note: Relative frequencies of each topic for each media sector in the period 2007 to 2012.

Table 3: News Publications per Topic and Day

Topic	Mean	Std.	Median	Min	Max	Iqr	Obs.
Political	65.98	34.48	58	5	208	44	74138
Foreign	56.96	25.17	53	3	142	35	70414
Economic	40.83	23.32	34	2	114	30	45151
Disaster	35.08	25.97	29	0	153	32	25284
Other	34.84	17.68	32	4	133	21	40482
Development	31.00	19.74	25	1	99	25	32416
Legal	28.91	22.57	22	0	141	23	23917
Race	15.86	11.48	12	0	56	14	14205
Total	45.90	29.73	39	0	208	37	326007

Source: Pew Research Center for Journalism. Note: It describes the measure of total publications per day and topic across the sample period 2007 to 2012. Std. stands for the standard deviation, Min for the minimum, Max for the maximum, Iqr for the interquartile range and Obs. for total observations. Observation unit is the news and the day.

Table 4: Topic Share per Day

Topic	Mean	Std.	Median	Min	Max	Iqr	Obs
Political	0.292	0.138	0.272	0.022	0.844	0.182	1,414
Foreign	0.256	0.113	0.238	0.012	0.765	0.138	1,414
Economic	0.180	0.093	0.158	0.011	0.462	0.120	1,414
Disaster	0.162	0.120	0.132	0.000	0.622	0.146	1,414
Other	0.153	0.071	0.138	0.024	0.554	0.081	1,414
Development	0.136	0.083	0.111	0.004	0.494	0.104	1,389
Legal	0.128	0.102	0.099	0.000	0.762	0.102	1,405
Race	0.073	0.052	0.057	0.000	0.274	0.063	1,407
Total	0.204	0.127	0.179	0.000	0.844	0.163	11,271

Source: Pew Research Center for Journalism. Note: It describes the multivariate measure of number of news on a topic over total number of news in a given day across the sample period 2007 to 2012. Std. stands for the standard deviation, Min for the minimum, Max for the maximum, Iqr for the interquartile range and Obs. for total observations. Observation unit is the day.

Table 5: Topic Newshole share per Day

Topic	Mean	Std.	Median	Min	Max	Iqr	Obs
Political	0.270	0.145	0.245	0.017	0.914	0.197	1414
Foreign	0.214	0.114	0.192	0.023	0.840	0.130	1414
Economic	0.125	0.090	0.103	0.003	0.543	0.100	1414
Development	0.108	0.078	0.088	0.000	0.611	0.066	1414
Other	0.099	0.064	0.085	0.008	0.667	0.066	1414
Disaster	0.068	0.087	0.037	0.000	0.663	0.068	1389
Legal	0.066	0.074	0.044	0.000	0.759	0.057	1405
Race	0.052	0.045	0.040	0.000	0.352	0.044	1407
Total	0.125	0.117	0.089	0.000	0.914	0.126	11271

Source: Pew Research Center for Journalism. Note: It describes the multivariate measure of topic newshole across the sample period 2007 to 2012. Std. stands for the standard deviation, Min for the minimum, Max for the maximum, Iqr for the interquartile range and Obs. for total observations. Observation unit is the day.

Table 6: Words per Topic and Day

Topic	Mean	Std.	Median	Min	Max	Iqr	Obs
Foreign	14023.99	5569.23	13447	720	38427	7389	19258
Political	11559.92	6567.48	10210	201	48237	7182	13205
Economic	9391.17	5192.37	8462	181	30540	6304	10545
Development	7577.40	4541.58	6611	86	26487	5097.5	8100
Disaster	5950.01	5212.73	4415	62	30733	5478	4058
Other	5305.96	3723.94	4379	72	30471	3781	5842
Legal	5073.39	4424.00	3841	81	28359	4228	4515
Race	3915.89	2714.92	3276	84	14057	3109	3237
Total	9800.09	6311.38	8715	62	48237	8572	68760

Source: Pew Research Center for Journalism. Note: It describes the total number of words per topic in a given day across the sample period 2007 to 2012. Std. stands for the standard deviation, Min for the minimum, Max for the maximum, Iqr for the interquartile range and Obs. for total observations. Observation unit is a news in a day.

Table 7: Minutes per Topic and Day

Topic	Mean	Std.	Median	Min	Max	Iqr	Obs
Political	9170.20	5056.49	8126	18	24549	6983	60933
Foreign	5164.87	3524.96	4115	88	18827	4452	51156
Disaster	3673.49	3661.65	2445	4	20812	3792	21226
Economic	3548.84	2971.40	2579	8	14621	3400	34606
Development	3440.65	2797.90	2437	16	13813	3197	24316
Legal	3250.80	3714.63	1956	9	24606	2963	19402
Other	3238.64	2555.42	2503	128	17578	2484	34640
Race	2088.75	1821.65	1527	4	10250	2016	10968
Total	5075.27	4405.62	3620	4	24606	5218	257247

Source: Pew Research Center for Journalism. Note: It describes the total number of minutes per topic in a given day across the sample period 2007 to 2012. Std. stands for the standard deviation, Min for the minimum, Max for the maximum, Iqr for the interquartile range and Obs. for total observations. Observation unit is a news in a day.

Table 8: Measures of Unidimensional Bias and Market Parameters

	New York Times	Washington Post	Wall Street Pos	USA Today	LA Times
Unidimensional Bias					
Economics	-2.58***	-3.54***	1.75***	-0.29	-1.27**
	(0.47)	(0.63)	(0.46)	(0.62)	(0.55)
Other				-0.56	-1.54**
				(0.71)	(0.65)
Development	-2.54***	-2.27***	-2.86***	1.01*	-0.59
	(0.52)	(0.58)	(0.85)	(0.56)	(0.54)
Disasters	-4.26***	-3.41***	-2.61***	-0.68	-2.88***
	(0.81)	(0.72)	(0.90)	(0.65)	(0.75)
Race		-2.01***		0.44	-1.01*
		(0.62)		(0.65)	(0.61)
Politics	-1.14***	-0.69	-1.27**	0.66	0.04
	(0.43)	(0.48)	(0.58)	(0.58)	(0.50)
Legal	-4.42***	-4.22***	-1.17*	-1.02	-1.77***
	(0.81)	(0.91)	(0.62)	(0.73)	(0.68)
Betas					
Economics	8.00***	8.92***	3.35***	6.79***	6.59***
	(0.78)	(1.05)	(0.79)	(0.94)	(0.90)
Other	5.02**	3.61	4.64	4.34***	5.61***
	(2.31)	(2.58)	(3.11)	(1.35)	(1.30)
Development	6.64***	6.31***	6.77***	4.78***	3.41***
	(0.97)	(1.10)	(1.64)	(1.07)	(1.05)
Disasters	6.55***	6.99***	5.33***	6.88***	7.46***
	(1.01)	(1.09)	(1.39)	(0.98)	(1.10)
Race	11.77***	5.71***	2.68	-0.83	4.55**
	(2.40)	(2.21)	(7.02)	(2.97)	(2.24)
Politics	3.74***	3.71***	4.12***	3.36***	3.10***
	(0.46)	(0.54)	(0.61)	(0.55)	(0.53)
Legal	9.22***	9.26***	4.37***	7.08***	5.72***
	(1.36)	(1.63)	(1.34)	(1.46)	(1.30)
Foreign	3.91***	2.69***	3.45***	3.30***	3.99***
	(0.55)	(0.62)	(0.67)	(0.74)	(0.64)
Observations	1086	697	671	690	697
Bundles	28	29	20	35	33
Goodness Fit	0.25	0.22	0.26	0.12	0.15

Note: Each column contain multinomial logit estimates of model (4) for each newspaper. Empty cells appear because some parameters are not identified in isolation. The baseline topic for all newspapers choices is Foreign news. Standard errors are in parenthesis. The news relevance measures used are the daily topic shares. ***,** and * indicate coefficients are significant at the 1%, 5% and 10% level respectively.

Table 9: Measures of Unidimensional Bias and Market Parameters controlling for News Relevance Interactions

	New York Times	Washington Post	Wall Street Pos	USA Today	LA Times
Unidimensional Bias					
Economics	-2.57***	-3.54***	1.76***	-0.29	-1.33**
	(0.48)	(0.64)	(0.46)	(0.62)	(0.55)
Other				-0.57	-1.58**
				(0.71)	(0.65)
Development	-2.53***	-2.26***	-2.85***	1.00*	-0.66
	(0.52)	(0.59)	(0.86)	(0.56)	(0.55)
Disasters	-4.25***	-3.41***	-2.59***	-0.70	-2.96***
	(0.81)	(0.72)	(0.90)	(0.65)	(0.76)
Race		-2.01***		0.41	-1.09*
		(0.62)		(0.66)	(0.62)
Politics	-1.14***	-0.68	-1.27**	0.70	0.01
	(0.43)	(0.48)	(0.58)	(0.59)	(0.50)
Legal	-4.41***	-4.22***	-1.16*	-1.04	-1.84***
	(0.81)	(0.91)	(0.62)	(0.73)	(0.68)
Betas					
Economics	8.04***	8.94***	3.45***	6.49***	6.28***
	(0.85)	(1.12)	(0.89)	(1.02)	(0.96)
Other	5.05**	3.63	4.71	4.08***	5.27***
	(2.32)	(2.61)	(3.12)	(1.39)	(1.36)
Development	6.67***	6.32***	6.86***	4.52***	3.14***
	(1.01)	(1.14)	(1.68)	(1.12)	(1.09)
Disasters	6.59***	7.01***	5.43***	6.60***	7.13***
	(1.06)	(1.15)	(1.45)	(1.05)	(1.16)
Race	11.81***	5.72**	2.78	-1.08	4.30*
	(2.42)	(2.24)	(7.04)	(3.00)	(2.26)
Politics	3.78***	3.74***	4.26***	2.93***	2.69***
	(0.61)	(0.73)	(0.85)	(0.77)	(0.70)
Legal	9.26***	9.28***	4.49***	6.80***	5.41***
	(1.39)	(1.67)	(1.42)	(1.51)	(1.34)
Foreign	3.96***	2.72***	3.57***	2.96***	3.53***
	(0.70)	(0.82)	(0.85)	(0.86)	(0.83)
Observations	1086	697	671	690	697
Bundles	28	29	20	35	33
Goodness Fit	0.25	0.22	0.26	0.12	0.15

Note: Each column contain multinomial logit estimates of model (4) for each newspaper. This model includes a control for the interaction effect of news relevance on top of complementarities. Empty cells appear because some parameters are not identified in isolation. The baseline topic for all newspapers choices is Foreign news. Standard errors are in parenthesis. The news relevance measures used are the daily topic shares. ***,** and * indicate coefficients are significant at the 1%, 5% and 10% level respectively.

Table 10: Joint Significance of Newspaper Complementarities using Topic Share

Newspaper	T	k	LR	Chi2	Reject	G.Fit
NYT	1086	20	124.67	31.41	1	0.25
WP	697	21	39.10	32.67	1	0.22
WSJ	671	12	35.74	21.03	1	0.26
USATODAY	690	27	33.91	40.11	0	0.12
LATIMES	697	25	31.91	37.65	0	0.15

Note: T states for number of day observations, k for degrees of freedom, LR is the likelihood ratio test statistic that tests a model with complementarity parameters against the model without them, Chi2 is the upper-tail critical value of the Chi-square distribution for a 5% level significance test and k degrees of freedom. Reject takes value 1 when we can reject the null that all complementarity parameters are equal to zero. G.Fit is the goodness of fit statistic.

Table 11: Joint Significance of Newspaper News Complementarity controlling for the Interaction of Market Relevance

Newspaper	T	k	LR	Chi2	Reject	G.Fit
NYT	1086	20	124.68	31.41	1	0.25
WP	697	21	39.11	32.67	1	0.22
WSJ	671	12	35.80	21.03	1	0.26
USATODAY	690	27	34.53	40.11	0	0.12
LATIMES	697	25	32.68	37.65	0	0.15

Note: T states for number of day observations, k for degrees of freedom, LR is the likelihood ratio test statistic that tests a model with complementarity parameters and interaction of market relevance against the model without both type of parameters, Chi2 is the upper-tail critical value of the Chi-square distribution for a 5% level significance test and k degrees of freedom. Reject takes value 1 when we can reject the null that all complementarity parameters and market relevance iteraction are equal to zero. G.Fit is the goodness of fit statistic.

Table 12: News Relevance Measure at Particular Dates

	Average	Lehman Brothers Collapse	Super Outbreak	Osama bin laden death
Economic	13.9%	37.6%	5.9%	5.3%
Other	12.7%	4.0%	8.5%	6.6%
Development	9.7%	4.0%	4.6%	2.0%
Disaster	7.8%	20.8%	30.7%	24.3%
Race	4.2%	0.4%	1.3%	2.6%
Political	23.0%	22.1%	20.9%	7.9%
Legal	7.7%	0.0%	0.0%	1.3%
Foreign	21.0%	11.1%	28.1%	50.0%

Note: Each of the columns presents a proposed share of topic relevance under different scenarios. The Average day is a scenario where all topic shares are on their average. Each of the following days in columns 3-5 are actual news relevance scenarios that happened: the 15th September 2008 for the Lehman Brothers Collapse; the 27th April 2011 for the Super Outbreak and the 2nd of May of 2011 for Osama bin Laden's death. All columns add up to one.

Table 13: Relative Effect of Complementarities between News (I)

	Average Day		Lehman Brothers		Super Outbreak		Osama bin laden	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
New York Times								
Economic	18.0%	9.6%	51.2%	18.3%	9.3%	5.2%	7.5%	4.5%
Other	1.7%	0.0%	0.9%	0.0%	1.2%	0.0%	1.1%	0.0%
Development	9.9%	4.9%	4.7%	2.6%	6.9%	3.5%	4.8%	2.6%
Disaster	3.9%	2.2%	5.4%	2.8%	16.1%	8.6%	8.6%	4.7%
Race	3.6%	0.0%	1.3%	0.0%	2.2%	0.0%	2.4%	0.0%
Political	21.5%	8.5%	14.6%	7.0%	19.9%	8.8%	9.7%	4.4%
Legal	3.7%	1.4%	1.3%	0.7%	1.5%	0.6%	1.5%	0.7%
Foreign	37.7%	14.5%	20.7%	11.5%	43.0%	15.1%	64.4%	13.8%
		41.1%		42.9%		41.9%		30.6%
Washington Post								
Economic	12.8%	4.6%	48.8%	12.2%	6.0%	2.4%	5.3%	2.2%
Other	3.3%	0.0%	1.9%	0.0%	3.2%	0.0%	2.7%	0.0%
Development	8.4%	0.7%	3.1%	0.1%	5.4%	0.3%	4.5%	0.5%
Disaster	4.8%	0.8%	8.1%	2.5%	21.4%	3.2%	13.4%	2.8%
Race	6.1%	0.6%	2.3%	-0.1%	4.1%	0.6%	5.1%	1.2%
Political	28.6%	6.1%	18.4%	6.5%	23.4%	4.6%	14.6%	4.0%
Legal	2.9%	0.2%	1.0%	0.2%	1.1%	0.0%	1.3%	0.1%
Foreign	33.1%	8.9%	16.3%	6.3%	35.4%	8.9%	53.1%	9.8%
		22.0%		27.8%		20.1%		20.6%
Wall Street Journal								
Economic	44.8%	4.8%	65.5%	3.1%	36.9%	3.6%	29.7%	5.4%
Other	2.1%	0.0%	1.3%	0.0%	1.7%	0.0%	1.3%	0.0%
Development	5.6%	1.6%	3.1%	1.1%	3.7%	0.9%	2.4%	0.4%
Disaster	1.6%	-1.3%	3.0%	-1.9%	6.9%	-4.8%	3.5%	-3.5%
Race	1.2%	0.0%	0.7%	0.0%	1.4%	0.0%	1.7%	0.0%
Political	14.7%	2.8%	10.3%	1.8%	13.7%	2.9%	7.0%	1.8%
Legal	3.5%	-2.3%	1.8%	-1.4%	2.4%	-1.7%	2.3%	-1.6%
Foreign	26.5%	4.5%	14.3%	3.3%	33.4%	2.9%	52.1%	3.7%
		10.2%		6.0%		3.8%		6.1%

Note: Column (1) contains predicted topic probability and column (2) contains the relative effect of complementarities in terms of topic probability. Predictions and relative effects of complementarities are made using baseline Model (4) estimated using daily topic shares as measures of news relevance for the different news scenarios that are presented in Table 12.

Table 14: Relative Effect of Complementarities between News (II)

	Average Day		**Lehman Brothers**		**Super Outbreak**		**Osama bin laden**	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
USA Today								
Economic	15.0%	-0.3%	48.9%	-1.7%	7.3%	-0.8%	6.8%	-1.2%
Other	13.7%	4.7%	6.9%	2.8%	10.0%	3.0%	9.3%	2.8%
Development	18.0%	-0.2%	8.9%	-0.2%	11.1%	-2.0%	10.5%	-1.2%
Disaster	5.9%	-1.9%	10.0%	-3.2%	30.4%	-5.7%	20.3%	-3.2%
Race	7.3%	-0.9%	4.1%	-1.5%	6.2%	-1.6%	6.6%	-1.3%
Political	22.6%	1.9%	14.4%	1.0%	17.4%	-0.4%	11.8%	0.1%
Legal	4.5%	-2.0%	1.4%	-1.0%	1.7%	-0.8%	2.0%	-1.3%
Foreign	13.0%	-0.7%	5.4%	-1.1%	15.8%	-0.2%	32.7%	-0.7%
		0.5%		-5.0%		-8.6%		-5.9%
Los Angeles Times								
Economic	13.5%	1.4%	44.7%	2.6%	6.6%	0.9%	5.6%	1.3%
Other	8.3%	0.2%	3.8%	0.0%	5.8%	0.6%	4.2%	0.5%
Development	12.4%	2.9%	7.2%	1.5%	9.5%	3.1%	7.1%	2.5%
Disaster	5.2%	1.9%	9.0%	2.1%	23.5%	8.1%	13.7%	6.0%
Race	6.8%	0.4%	3.8%	0.0%	4.0%	0.3%	4.1%	0.6%
Political	20.5%	1.4%	14.2%	0.7%	16.2%	2.0%	8.9%	1.4%
Legal	4.7%	-0.4%	1.9%	-0.5%	2.2%	0.0%	2.7%	0.6%
Foreign	28.6%	7.4%	15.4%	5.0%	32.3%	9.9%	53.6%	11.2%
		15.3%		11.4%		24.9%		24.0%

Note:Column (1) contains predicted topic probability and column (2) contains the relative effect of complementarities in terms of topic probability. Predictions and relative effects of complementarities are made using baseline Model (4) estimated using daily topic shares as measures of news relevance for the different news scenarios that are presented in Table 12.

Table 15: Mean Square Errors in the OOS validation

	Complementarities	No complementarities
NYT	0.859	0.875
WP	0.901	0.904
WSJ	0.849	0.846
USATODAY	0.934	0.933
LATIMES	0.939	0.940

Note: Column (1) the mean square errors calculated using real outcomes and model fit using the model with complementary news and column (2) the mean square error calculated the model without complementarities. The model uses daily topic share as measures of news relevance.

Table 16: Topic Aggregation from Original NCID Broad Topics

Label	**NCID Broad Topics**
Political	Government - Campaign - Defense
Foreign	US - Non Us Foreign News
Economic	Business-Economics
Disaster	Disaster-Domestic Terrorism
Other	Celebrities-Sports-Lifestyle-Media-Miscellaneous
Development	Development-Environment-Transportation-Education-Religion-Health/Medicine-Science/Technology
Legal	Legal - Crime
Race	Race/Gender/Gay Issues - Immigration - Other Domestic Affairs

Figure 1: Relative Frequency of Bundles in the Front Page for New York Times.

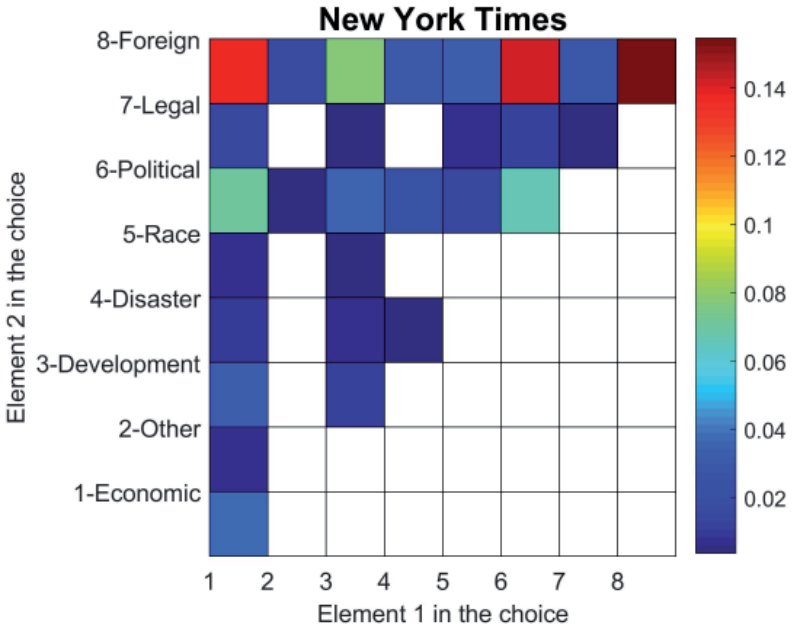

Figure 2: Relative Frequency of Bundles in the Front Page for Washington Post

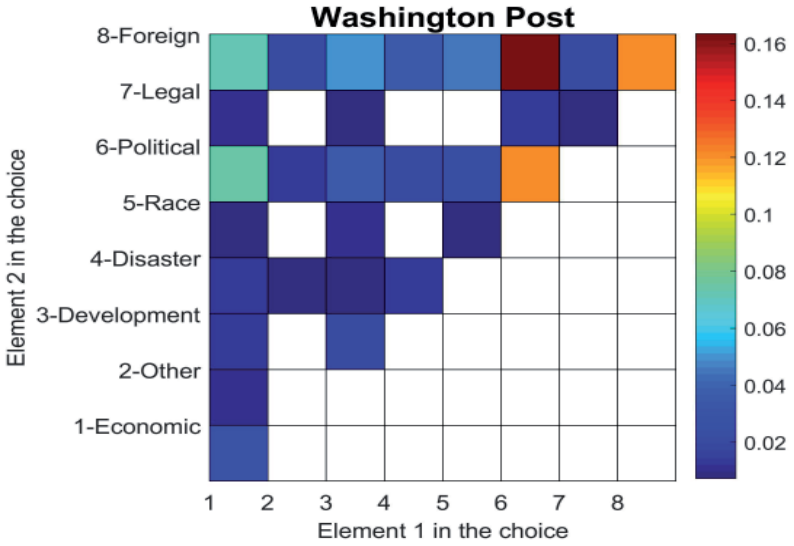

Figure 3: Relative Frequency of Bundles in the Front Page for Wall Street Journal

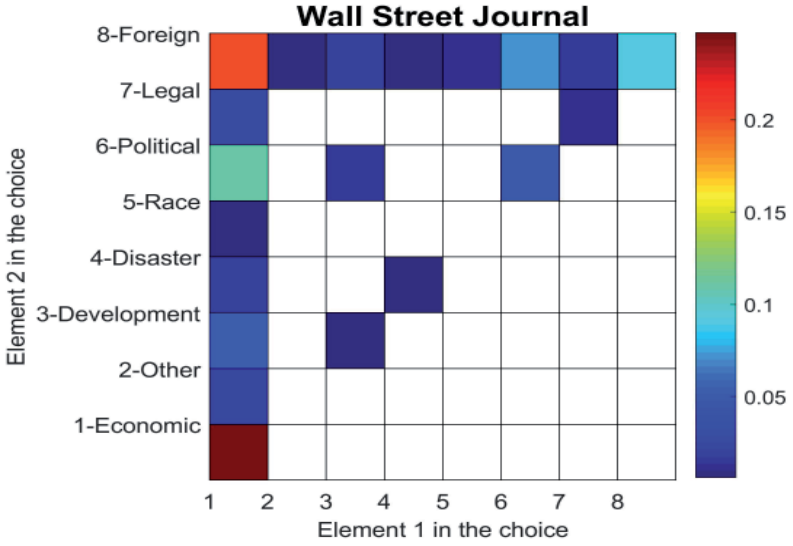

Figure 4: Relative Frequency of Bundles in the Front Page for USA Today

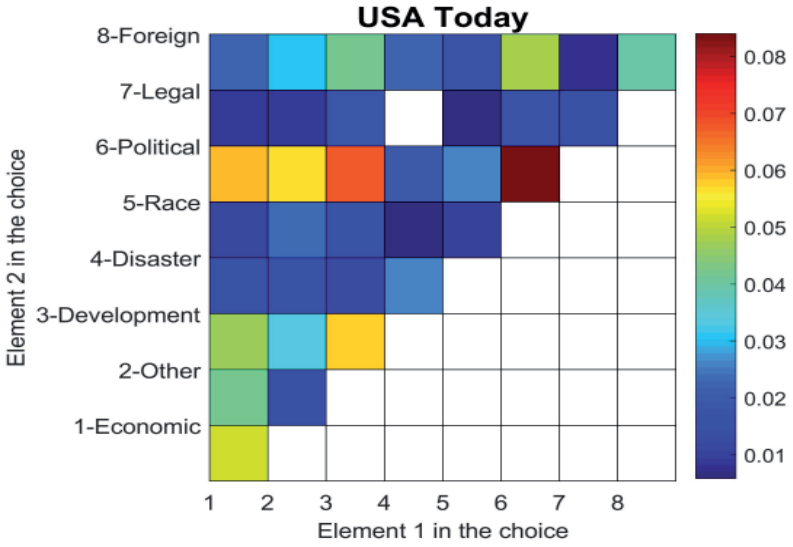

Figure 5: Relative Frequency of Bundles in the Front Page for Los Angeles Times

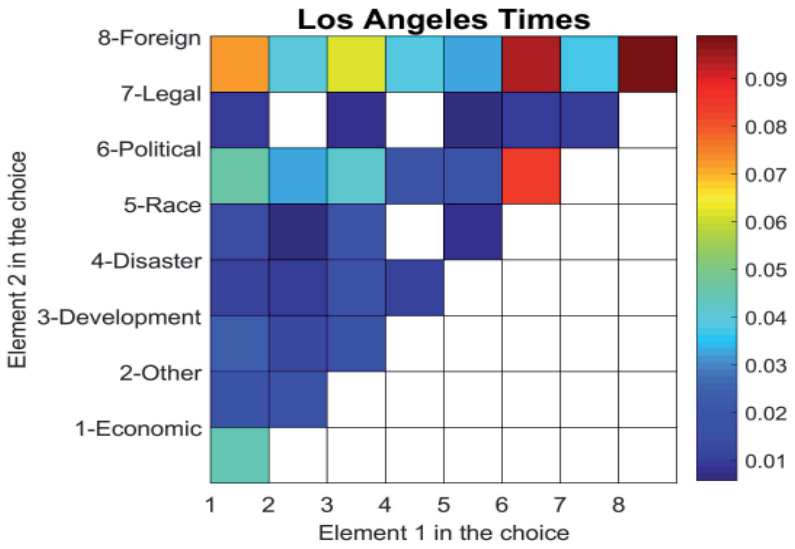

Figure 6: Economic News Relevance to Dow Jones Industrial Average Index

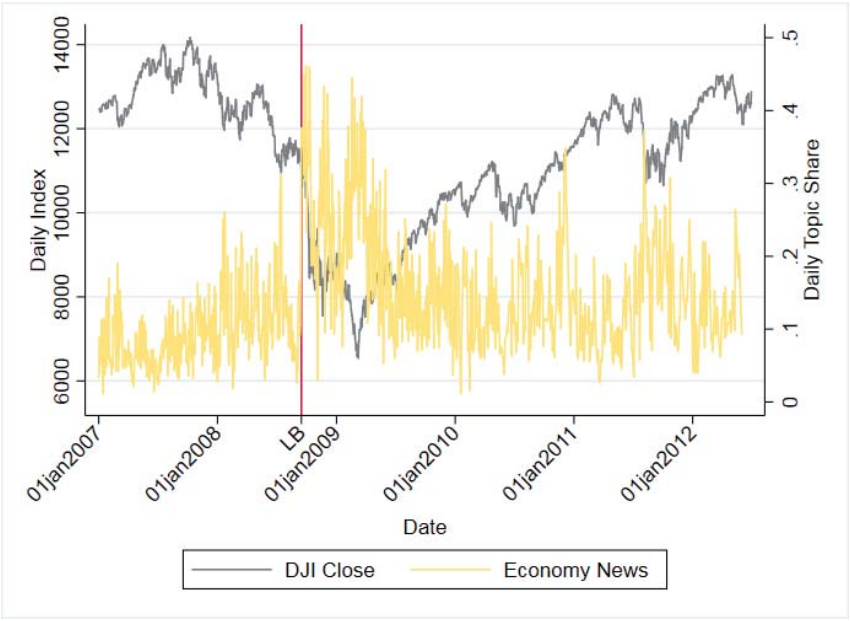

Figure 7: Political News Relevance to Google Trends Index for "political"

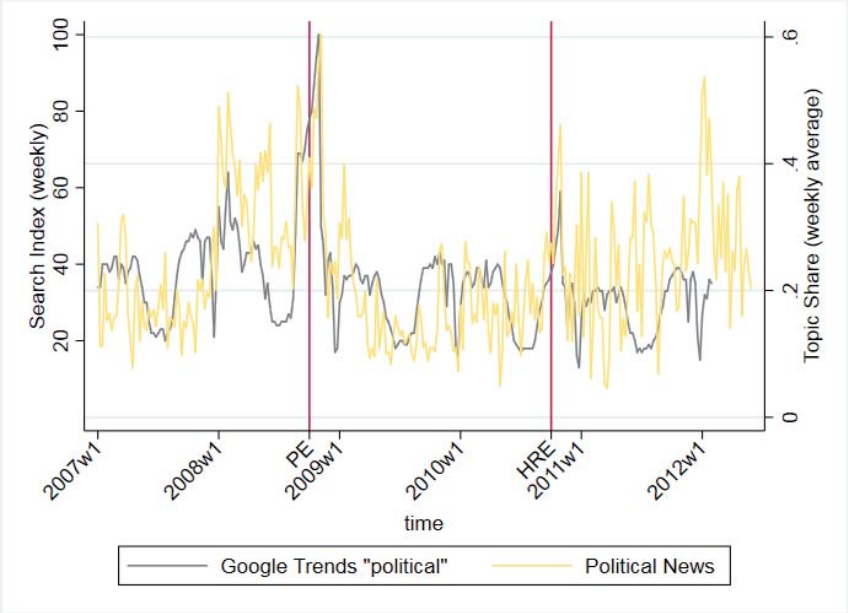

Figure 8: New York Times's Complementarity News using Topic Shares

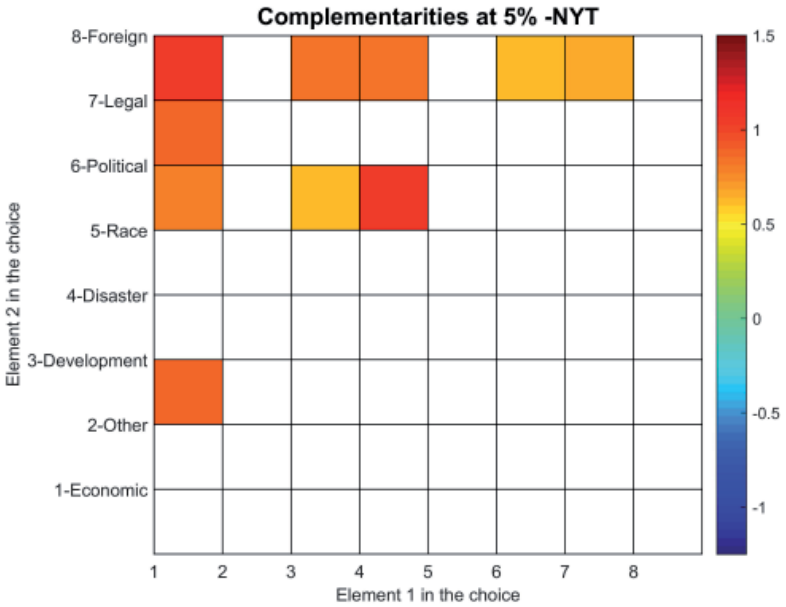

Figure 9: The Washington Post's Complementary News using Topic Shares

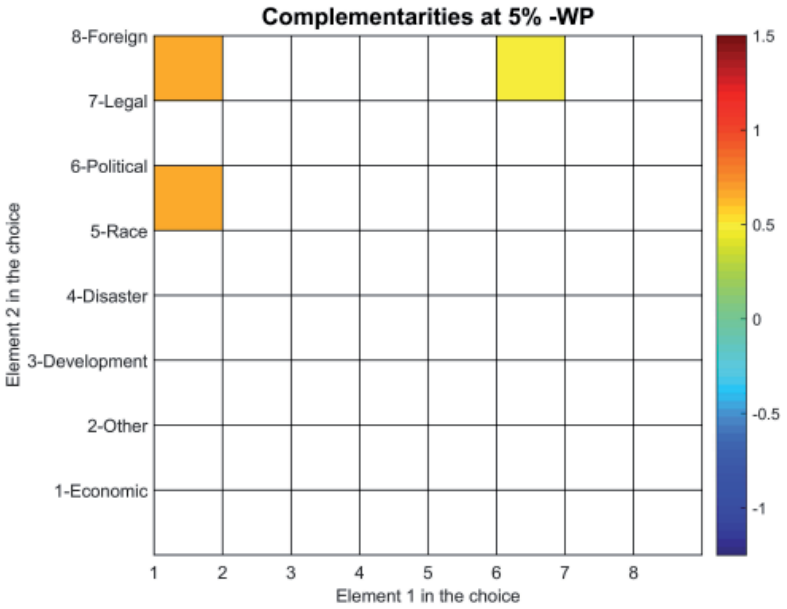

Figure 10: The Wall Street Journal's Complementary News using Topic Shares

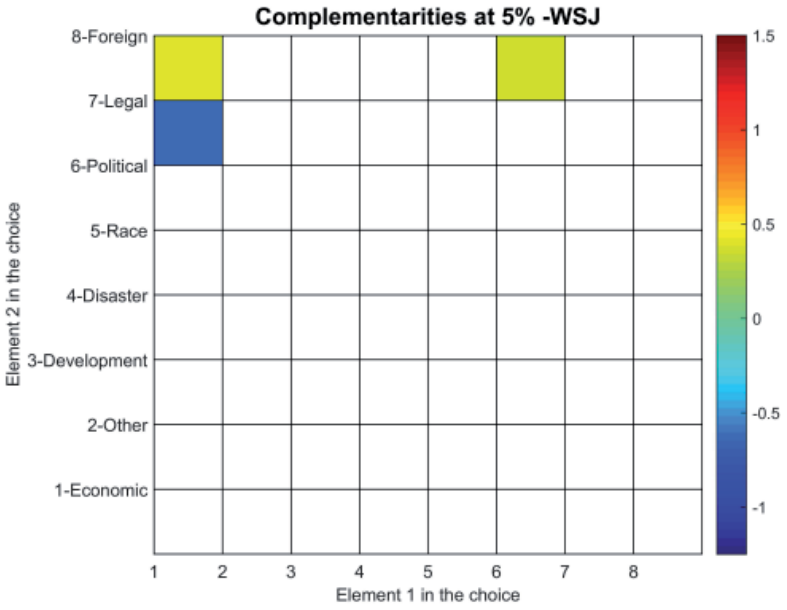

Figure 11: The USA Today's Complementary News using Topic Shares

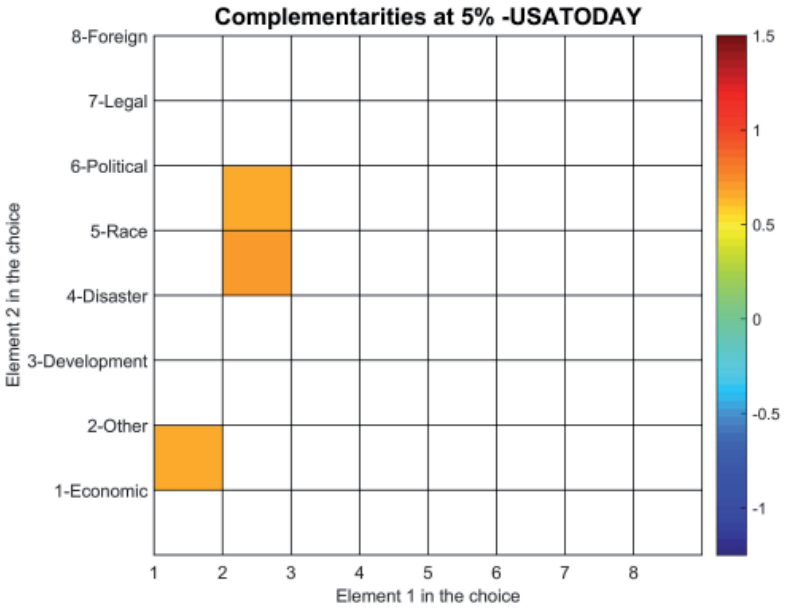

Figure 12: Los Angeles Times's Complementary News using Topic Shares

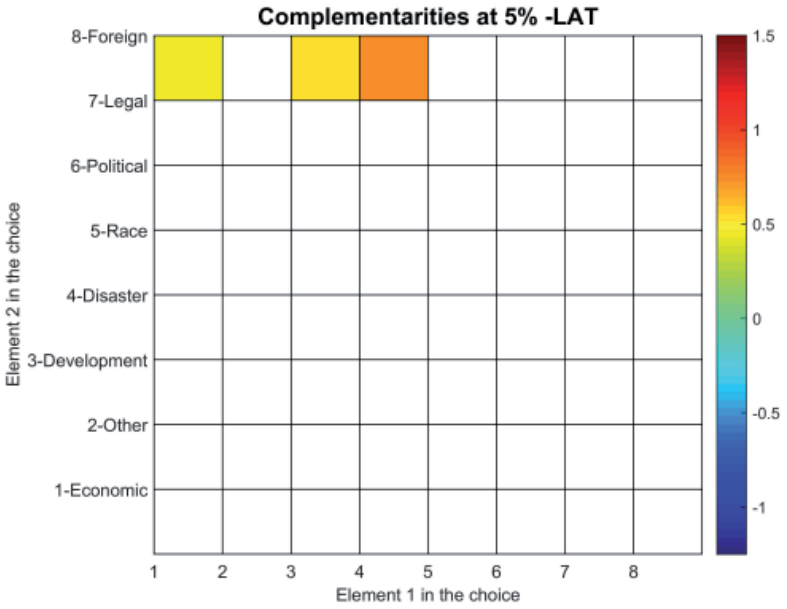

Figure 13: The New York Times's Magnitude of Complementarities

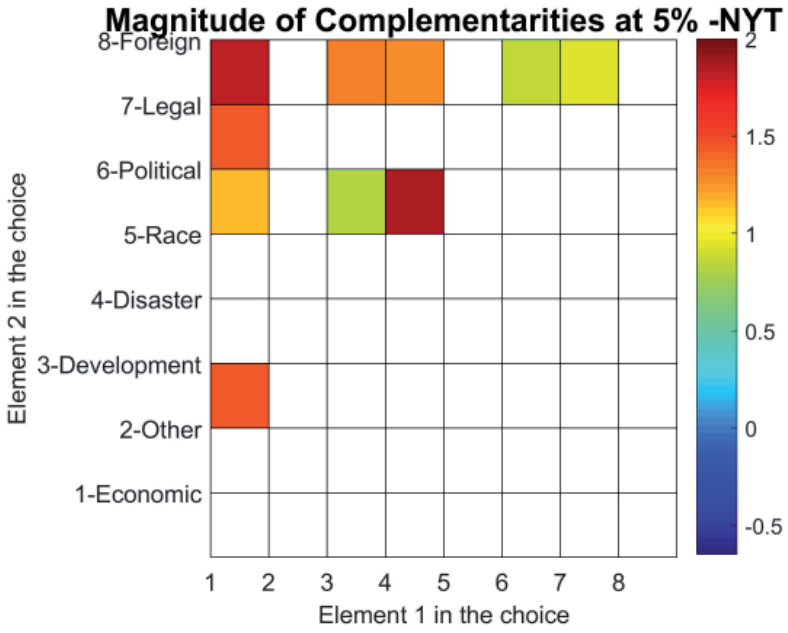

Figure 14: The Washington Post's Magnitude of Complementarities

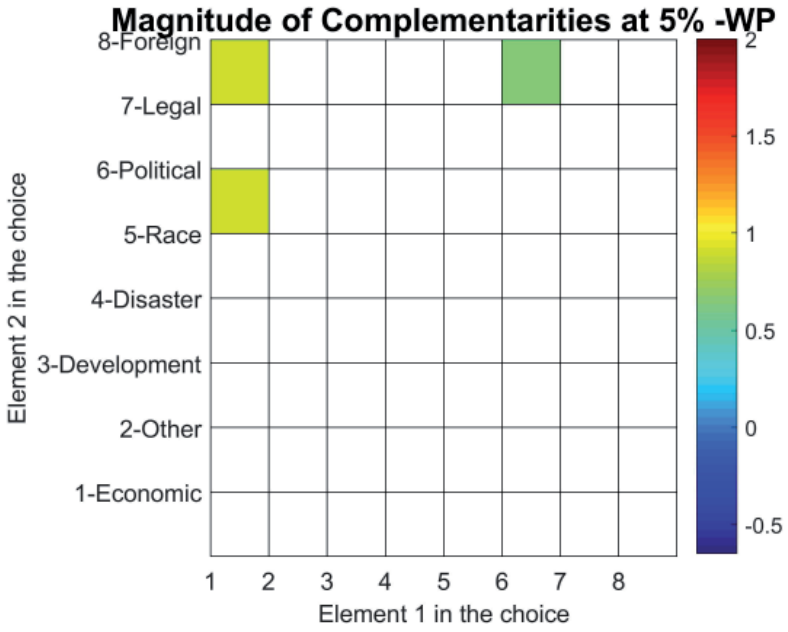

Figure 15: The Wall Street Journal's Magnitude of Complementarities

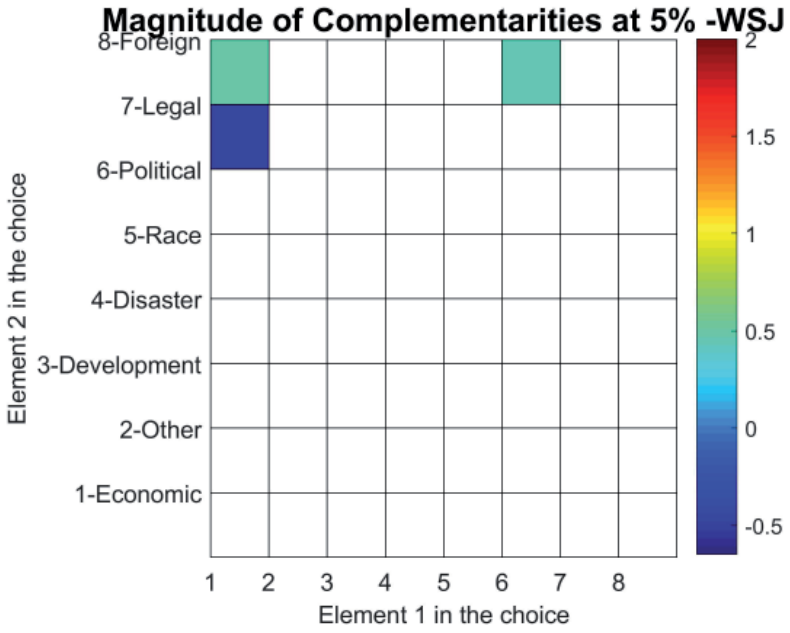

Figure 16: USA Today's Magnitude of Complementarities

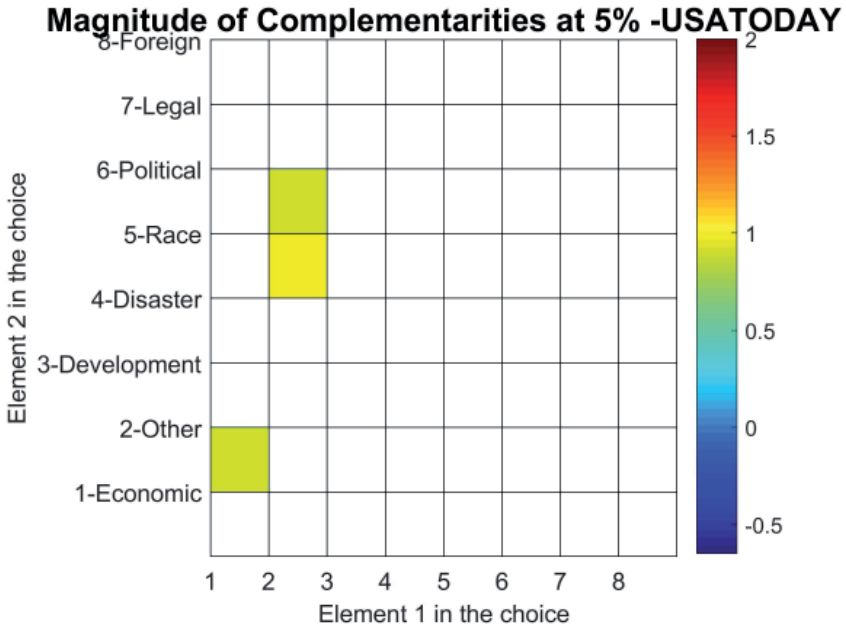

Figure 17: The Los Angeles Times's Magnitude of Complementarities

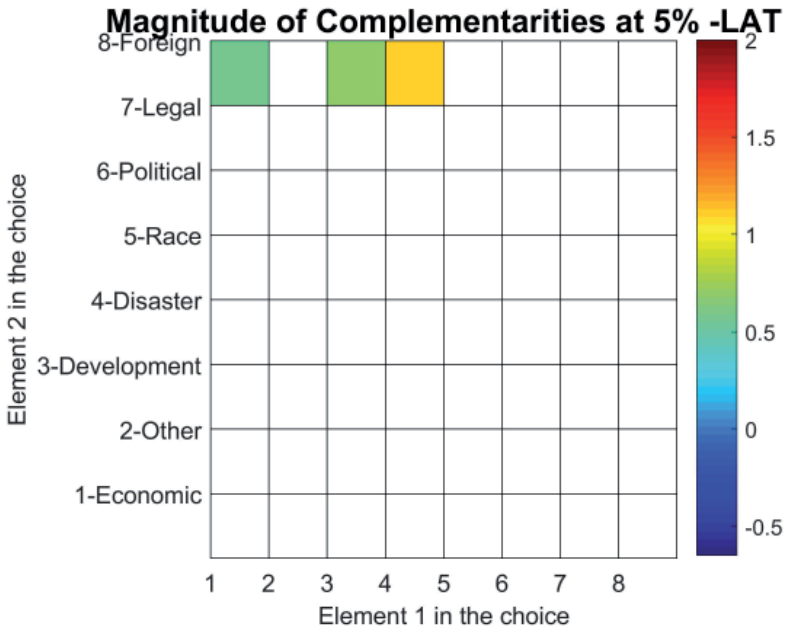

Figure 18: MTP Surviving Complementarities for NYT using Topic Share

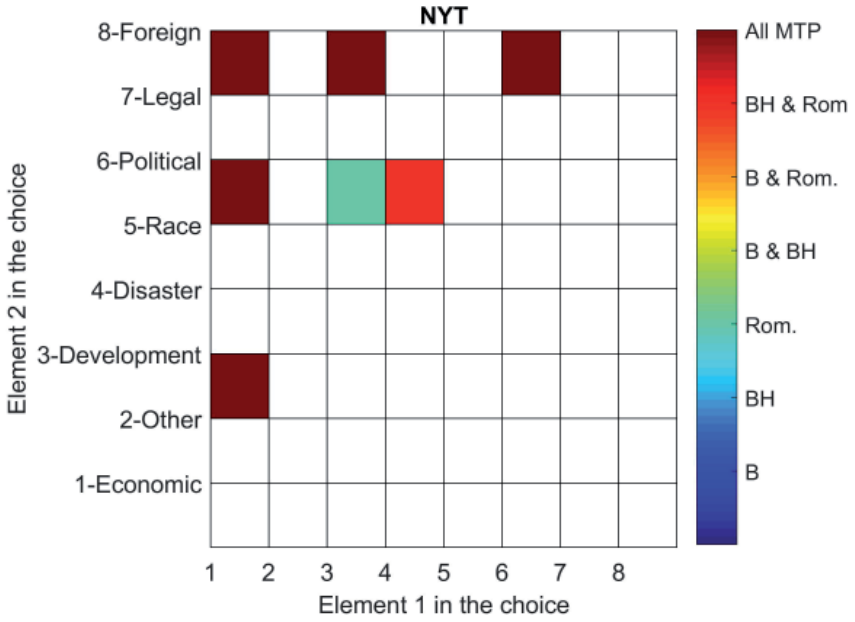

Note:Cells are colored at complementarities for the New York Times which end up being significant using a particular Multiple Testing Procedure, either Bonferroni (1936) (B), the stepwise individual test of Holm (1979) (Bonferroni-Holm, BH) or the dependence-control stepwise test of Romano and Wolf (2005) (Rom.). Tests performed over baseline results that use topic shares as news relevance measures.

Figure 19: MTP Surviving Complementarities for WP using Topic Share

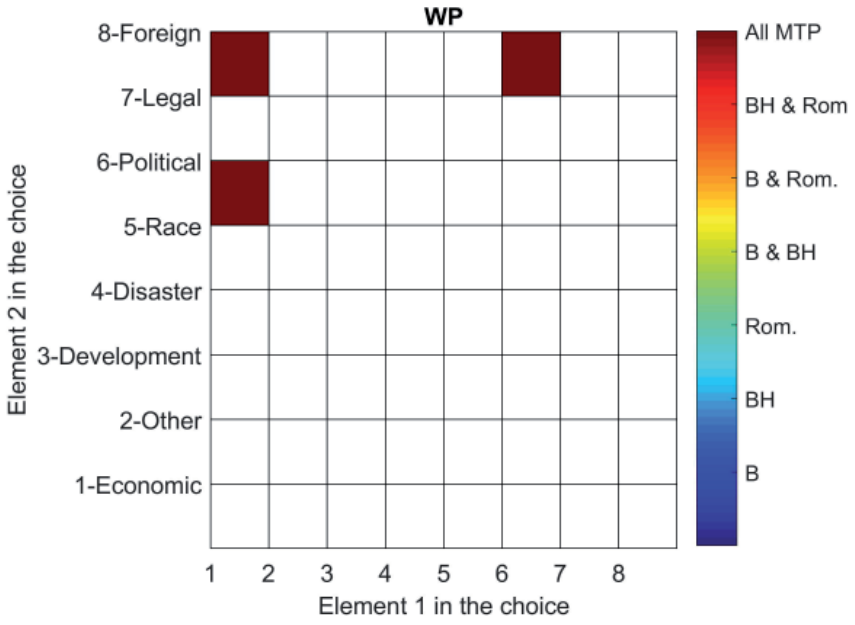

Note:Cells are colored at complementarities for the Washington Post which end up being significant using a particular Multiple Testing Procedure, either Bonferroni (1936) (B), the stepwise individual test of Holm (1979) (Bonferroni-Holm, BH) or the dependence-control stepwise test of Romano and Wolf (2005) (Rom.). Tests performed over baseline results that use topic shares as news relevance measures.

Figure 20: MTP Surviving Complementarities for WSJ using Topic Share

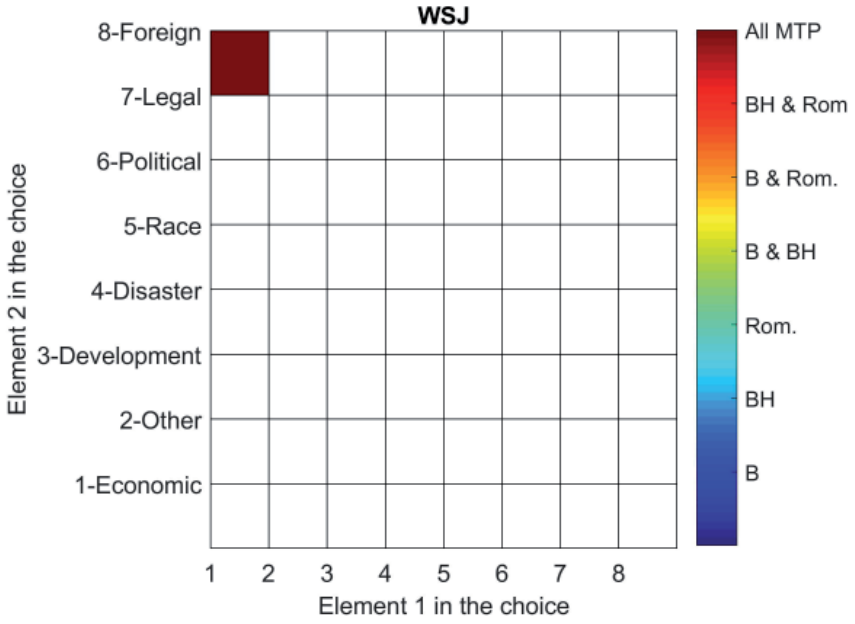

Note:Cells are colored at complementarities for the Wall Street Journal which end up being significant using a particular Multiple Testing Procedure, either Bonferroni (1936) (B), the stepwise individual test of Holm (1979) (Bonferroni-Holm, BH) or the dependence-control stepwise test of Romano and Wolf (2005) (Rom.). Tests performed over baseline results that use topic shares as news relevance measures.

A Roots to Model Parameters

Given that the number of topics can vary depending on the interest of the researcher, the number of potential elements in the choice set can become very large. The researcher would like to know the values of all the parameters associated to any potential choices that the newspaper could take. As the researcher becomes more demanding on the potential choices a newspaper can make, the number of different choices actually made by the newspaper may be fewer than those in the potential set in a limited time period. Let us assume that we know the value of κ. If the newspaper never chose government and crime, we cannot separately identify the value of their complementarity with respect to the preference for the topic government or crime. If the newspaper did not choose full front page on crime, we are not able to separately identify the taste for crime. Neither we separately identify the complementarity between government and crime news even if we are able to identify the taste for government news. The complementarity parameter, Γ_{jk} is separately identified from the choices of the bundle once we know δ_j and δ_k.

The lack of some choices blocks separate identification of parameters associated to other choices and it makes the empirical specification more complex. To overcome this complication, I take a two-stage approach for estimation. I first estimate the parameters in a model with a single bundle specific constant,

$$ln\frac{p_{at}^n}{p_{At}^n} = \theta_a + \beta_a X_{at} \tag{15}$$

where p_{at}^n is the newspaper probability of choice a at day t and p_{At}^n is the newspaper probability of baseline choice A at day t. θ_a is the bundle-specific bias parameter. $\beta_a = (\beta_{j|a_1=j}, \beta_{g|a_2=g})'$ and $X_{at} = (X_{j|a_1=j}, X_{g|a_2=g})$. I obtain estimates $\hat{\theta}_a$ and $\hat{\beta}_a$, $\forall a$. In a second step, I use the structure of θ_a, specified in equation (4), to compute the values of the identified parameters of interest, Γ_a and δ_j $\forall a \in \hat{A}, j \in \hat{J}$, where $\hat{A}$ refers to the set of bundles for which Γs are identified and $\hat{J}$ refers to the set of topics for which δs are identified using the available data. Analytically,

$$\hat{\theta} = H\theta \tag{16}$$

where θ is a vector $Nx1$ that contains the set of potentially identified parameters, this is, $\theta = (\delta_1, ..., \delta_{J-1}, \Gamma_1, ..., \Gamma_A)$. H is a matrix of size MxN that represents the structure for θ which is the one of model (4). $\hat{\theta}$ is a vector of size $Mx1$ that contains the estimated parameters from ML on (15). In general, $M \leq N$, so H is not a squared matrix and it

will not be invertible. We can decompose H into a triangular, one of them being a square matrix, such that,

$$\tilde{\theta} = LU\theta \tag{17}$$

where L is an MxM lower triangular matrix and U is an MxN upper triangular matrix.

$$L^{-1}\tilde{\theta} = U\theta$$

where $U\theta$ is the vector of parameters that can be identified out of the model and the data. $L^{-1}\tilde{\theta}$ is the value of those parameters.

Identification works is the following way. If the newspaper never chose topic j or k, but chose the bundle jk at some point, we identify $\theta_a = \delta_j + \delta_k + \Gamma_{jk}$. In the case that the newspaper chose j but never k, but also chose the bundle jk, we identify δ_j and $\delta_k + \Gamma_{jk}$. If it also chose other bundle with k, such as gk, we additional identify $\Gamma_{gk} - \Gamma_{jk}$. The last case is where the newspaper ever chose j, k and jk, in this case, we identify all δ_j, δ_k and Γ_{jk}.

An approach to the problem of identification of all the δs in the model due to sample size is to use other sources of variation in the data. One possibility is to add an intensity equation to the model, the intensity will be related to how much each topic occupies in the front page, e.g inches, number of words, size of the characters, position. The intensity adds a second layer to the choice problem where the newspaper first chooses the topic and then how much to publish in the front page. It may be that in the second layer one identifies the parameters for the taste of each topic in the intensity model, but yet these are parameters with different meanings to those of the first layer of the decision. The idea is to think of a set of assumptions that helps identify δ_j even when j is never chosen in isolation.

B Appendix Tables

Table B.1: Multinomial Logit Measures of Unidimensional Bias and Market Parameters using Topic Newshole

	New York Times	Washington Post	Wall Street Pos	USA Today	LA Times
Unidimensional Bias					
Economics	-2.20***	-2.76***	1.69***	-0.20	-0.86
	(0.46)	(0.61)	(0.42)	(0.58)	(0.53)
Other				-0.75	-1.17*
				(0.67)	(0.61)
Development	-2.66***	-2.16***	-3.07***	1.09**	-0.52
	(0.54)	(0.60)	(0.86)	(0.53)	(0.55)
Disasters	-5.00***	-3.22***	-2.52***	-0.66	-2.84***
	(0.90)	(0.75)	(0.89)	(0.64)	(0.78)
Race		-2.35***		-0.04	-1.37**
		(0.69)		(0.65)	(0.66)
Politics	-1.08**	-0.57	-1.36**	0.66	0.11
	(0.44)	(0.50)	(0.57)	(0.57)	(0.52)
Legal	-4.79***	-4.19***	-1.20**	-0.79	-1.59**
	(0.86)	(0.97)	(0.61)	(0.70)	(0.69)
Market Relevance					
Economics	7.54***	7.84***	3.91***	6.65***	6.24***
	(0.72)	(0.93)	(0.72)	(0.84)	(0.80)
Other	4.76***	3.66	4.93**	5.90***	5.99***
	(1.78)	(2.26)	(2.33)	(1.20)	(1.20)
Development	6.37***	5.86***	6.22***	4.32***	3.46***
	(0.81)	(0.91)	(1.31)	(0.87)	(0.83)
Disasters	7.93***	7.33***	5.12***	7.17***	7.83***
	(1.03)	(1.08)	(1.22)	(0.98)	(1.09)
Race	13.00***	8.17***	3.14	4.19**	7.15***
	(2.02)	(1.77)	(5.60)	(2.07)	(1.73)
Politics	3.35***	3.47***	3.68***	3.00***	2.82***
	(0.40)	(0.48)	(0.53)	(0.47)	(0.46)
Legal	9.33***	8.96***	4.36***	6.06***	5.63***
	(1.20)	(1.46)	(1.09)	(1.14)	(1.10)
Foreign	4.38***	3.41***	3.51***	3.40***	4.37***
	(0.53)	(0.60)	(0.59)	(0.64)	(0.61)
Observations	1086	697	671	690	697
Bundles	28	29	20	35	33
Goodness Fit	0.28	0.25	0.27	0.14	0.18

Note: Each column contain multinomial logit estimates of model (4) for each newspaper using topic newshole as measure for news relevance. Empty cells appear because some parameters are not identied in isolation. The baseline topic for all newspapers choices is Foreign news. Standard errors are in parenthesis. The news relevance measures used are the daily topic shares. ***,** and * indicate coecients are signicant at the 1%, 5% and 10% level respectively.

Table B.2: Multinomial Logit Measures of Unidimensional Bias and Market Parameters using Total Publications

	New York Times	Washington Post	Wall Street Pos	USA Today	LA Times
Unidimensional Bias					
Economics	-2.72***	-3.74***	0.56	0.17	-1.14**
	(0.47)	(0.63)	(0.46)	(0.62)	(0.55)
Other				-0.54	-1.67***
				(0.71)	(0.65)
Development	-2.62***	-2.47***	-3.81***	1.07**	-1.17**
	(0.52)	(0.58)	(0.85)	(0.56)	(0.54)
Disasters	-4.22***	-3.46***	-3.13***	-0.46	-3.64***
	(0.81)	(0.72)	(0.90)	(0.65)	(0.75)
Race		-2.23***		0.62	-1.48**
		(0.62)		(0.65)	(0.61)
Politics	-1.25***	-1.03**	-1.71***	1.16**	0.08
	(0.43)	(0.48)	(0.58)	(0.58)	(0.50)
Legal	-4.54***	-4.23***	-2.12***	-0.97	-1.95***
	(0.81)	(0.91)	(0.62)	(0.73)	(0.68)
Betas					
Economics	8.00***	8.92***	3.35***	6.79***	6.59***
	(0.78)	(1.05)	(0.79)	(0.94)	(0.90)
Other	5.02**	3.61	4.64	4.34***	5.61***
	(2.31)	(2.58)	(3.11)	(1.35)	(1.30)
Development	6.64***	6.31***	6.77***	4.78***	3.41***
	(0.97)	(1.10)	(1.64)	(1.07)	(1.05)
Disasters	6.55***	6.99***	5.33***	6.88***	7.46***
	(1.01)	(1.09)	(1.39)	(0.98)	(1.10)
Race	11.77***	5.71***	2.68	-0.83	4.55**
	(2.40)	(2.21)	(7.02)	(2.97)	(2.24)
Politics	3.74***	3.71***	4.12***	3.36***	3.10***
	(0.46)	(0.54)	(0.61)	(0.55)	(0.53)
Legal	9.22***	9.26***	4.37***	7.08***	5.72***
	(1.36)	(1.63)	(1.34)	(1.46)	(1.30)
Foreign	3.91***	2.69***	3.45***	3.30***	3.99***
	(0.55)	(0.62)	(0.67)	(0.74)	(0.64)
Observations	1086	697	671	690	697
Bundles	28	29	20	35	33
Goodness Fit	0.25	0.22	0.26	0.12	0.15

Note: Each column contain multinomial logit estimates of model (4) for each newspaper using number of news per topic in a day as market relevance. Empty cells appear because some parameters are not identified in isolation. The baseline topic for all newspapers choices is Foreign news. Standard errors are in parenthesis. The news relevance measures used are the daily topic shares. ***,** and * indicate coefficients are significant at the 1%, 5% and 10% level respectively.

Table B.3: Joint Significance of Newspaper News Complementarity using Topic Newshole

Newspaper	T	k	LR	Chi2	Reject	G.Fit
NYT	1086	20	124.67	31.41	1	0.25

Note: T states for number of day observations, k for degrees of freedom, LR is the likelihood ratio test statistic that tests a model with complementarity parameters that uses topic share as market relevance measure against the model without both type of parameters, Chi2 is the upper-tail critical value of the Chi-square distribution for a 5% level significance test and k degrees of freedom. Reject takes value 1 when we can reject the null that all complementarity parameters are equal to zero. G.Fit is the goodness of fit statistic.

Table B.4: Joint Significance of Newspaper News Complementarity using Total Publications

Newspaper	T	k	LR	Chi2	Reject	G.Fit
NYT	1086	18	122.80	28.87	1	0.25
WP	697	19	40.39	30.14	1	0.22
WSJ	671	10	35.87	18.31	1	0.25
USATODAY	690	25	34.36	37.65	0	0.12
LAT	697	23	32.27	35.17	0	0.15

Note: T states for number of day observations, k for degrees of freedom, LR is the likelihood ratio test statistic that tests a model with complementarity parameters that uses the total publications per topic in a day as market relevance measure against the model without both type of parameters, Chi2 is the upper-tail critical value of the Chi-square distribution for a 5% level significance test and k degrees of freedom. Reject takes value 1 when we can reject the null that all complementarity parameters are equal to zero. G.Fit is the goodness of fit statistic.

Table B.5: Joint Significance of Newspaper News Complementarity using Lagged Topic Share Relevance

Newspaper	T	k	LR	Chi2	Reject	G.Fit
NYT	1085	20	125.83	31.41	1	0.25
WP	697	21	44.31	32.67	1	0.23
WSJ	671	12	37.26	21.03	1	0.26
USATODAY	690	27	35.78	40.11	0	0.12
LATIMES	697	25	32.01	37.65	0	0.16

Note: T states for number of day observations, k for degrees of freedom, LR is the likelihood ratio test statistic that tests a model with complementarity parameters that uses the topic shares at the previous day available in the sample as market relevance measure against the model without both type of parameters, Chi2 is the upper-tail critical value of the Chi-square distribution for a 5% level significance test and k degrees of freedom. Reject takes value 1 when we can reject the null that all complementarity parameters are equal to zero. G.Fit is the goodness of fit statistic.

C Appendix Figures

Figure C.1: New York Times' Complementarity News controlling for Market Interaction

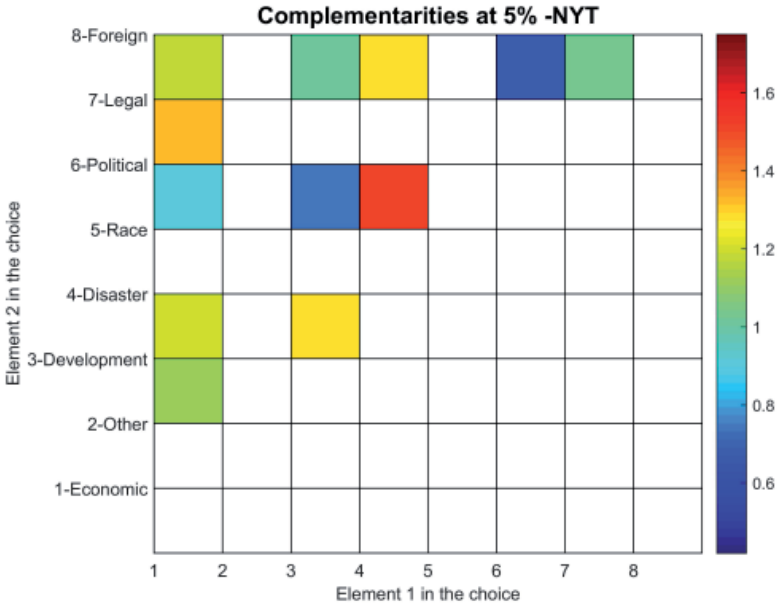

Figure C.2: Washington Post's Complementarity News controlling for Market Interaction

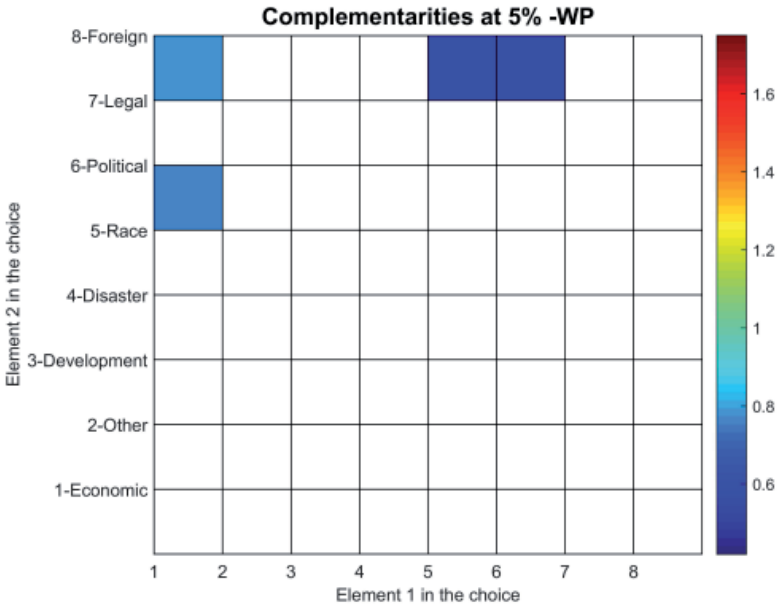

Figure C.3: The Wall Street Journal's Complementarity News controlling for Market Interaction

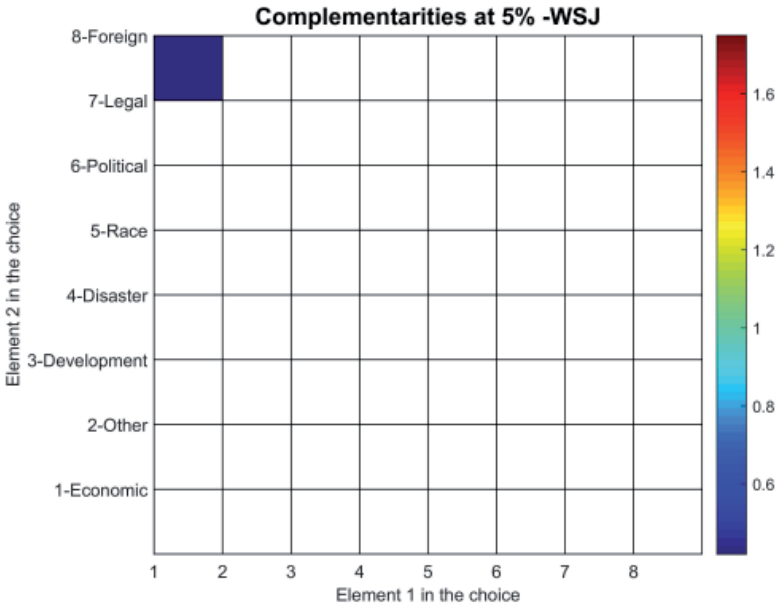

Figure C.4: The USA Today's Complementarity News controlling for Market Interaction

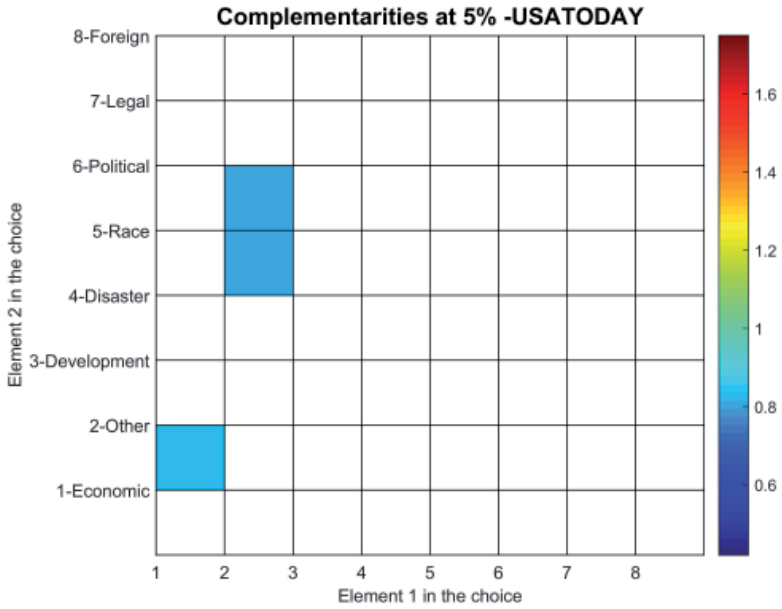

Figure C.5: The Los Angeles Times' Complementarity News controlling for Market Interaction

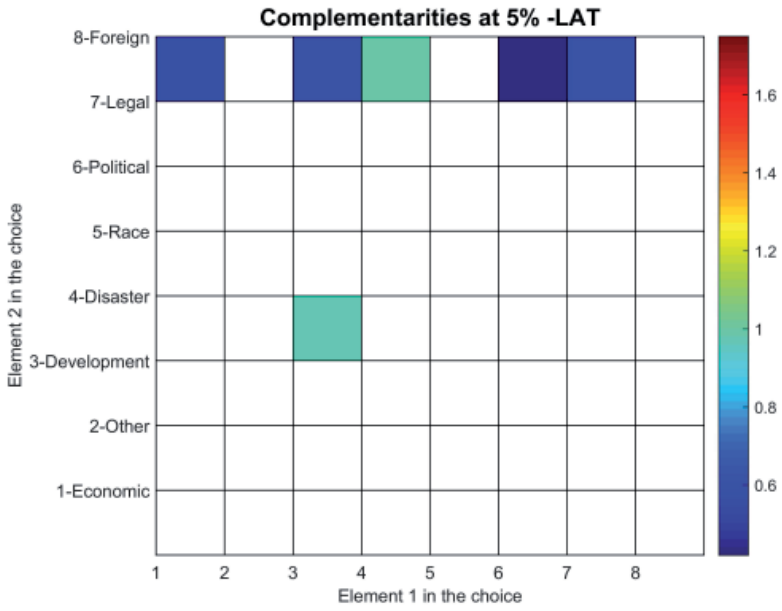

Figure C.6: New York Times's Complementarity News using Topic Newshole Relevance

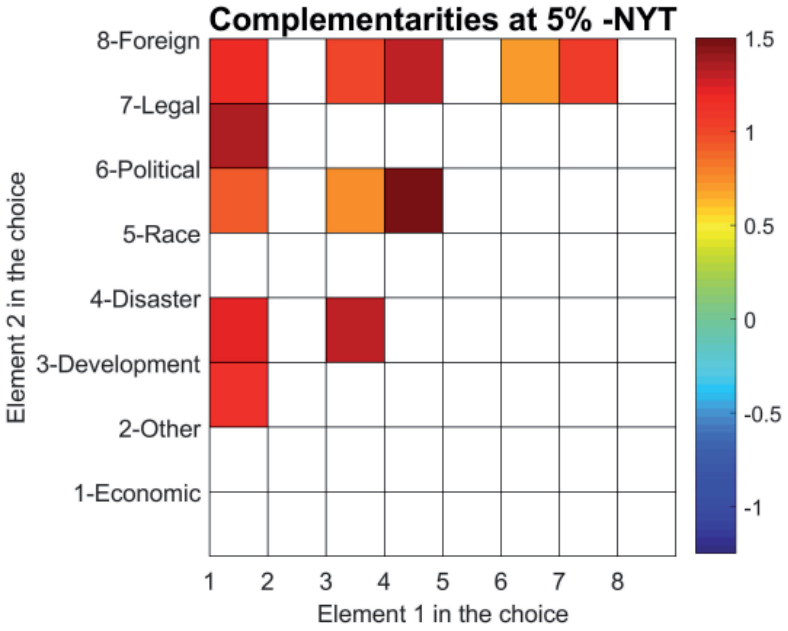

Figure C.7: The Washington Post's Complementary News using Topic Newshole Relevance

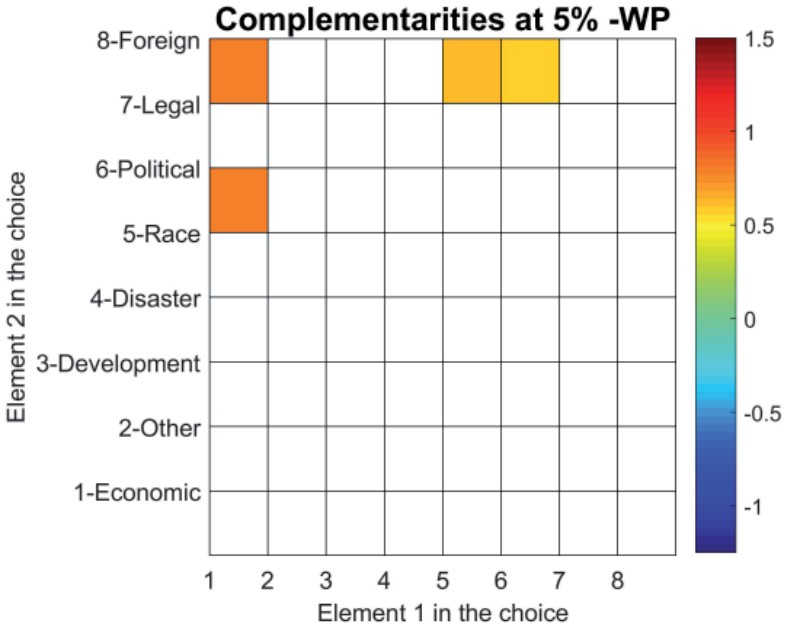

Figure C.8: The Wall Street Journal's Complementary News using Topic Newshole Relevance

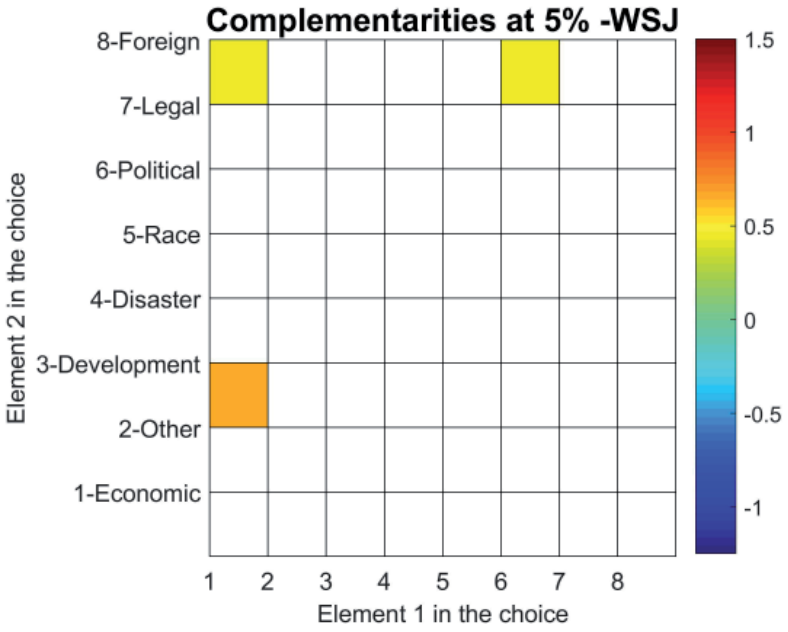

Figure C.9: The USA Today's Complementary News using Topic Newshole Relevance

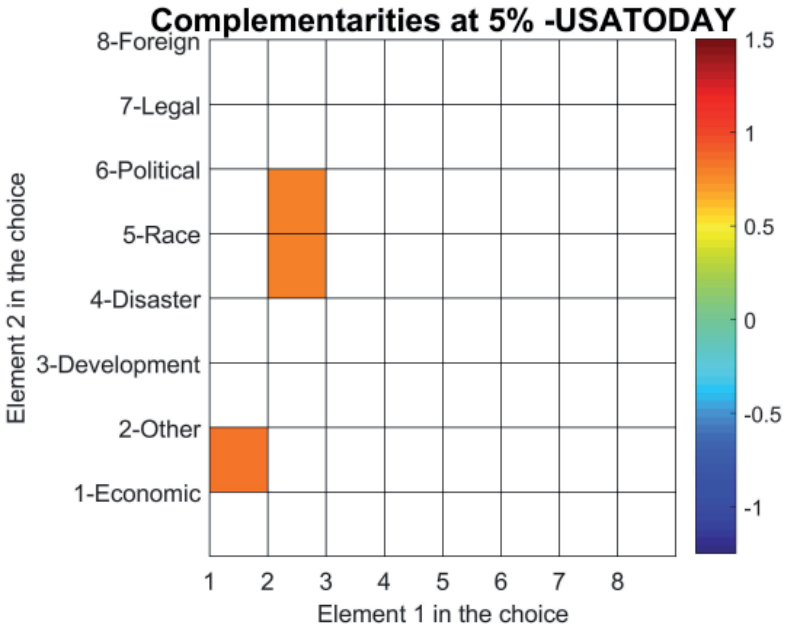

Figure C.10: Los Angeles Times's Complementary News using Topic Newshole Relevance

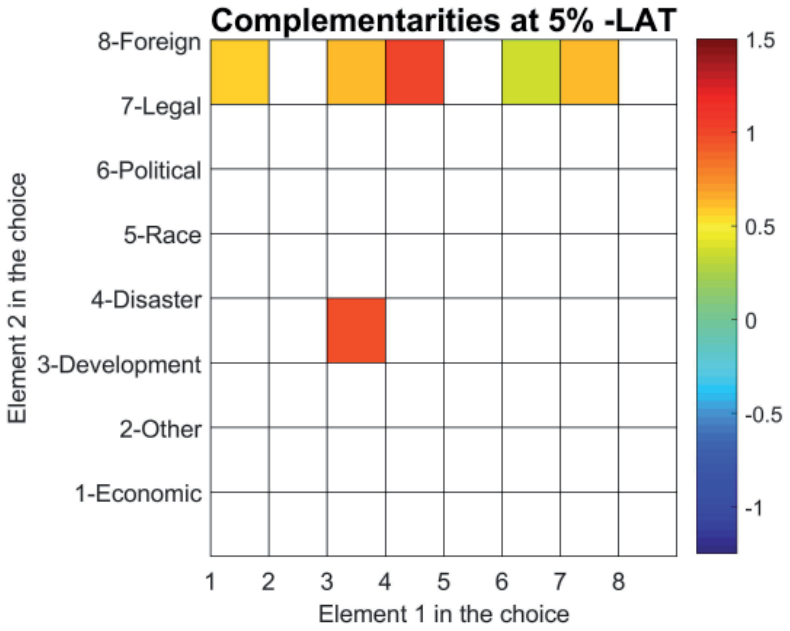

Figure C.11: New York Times' Complementarity News using Total Publications

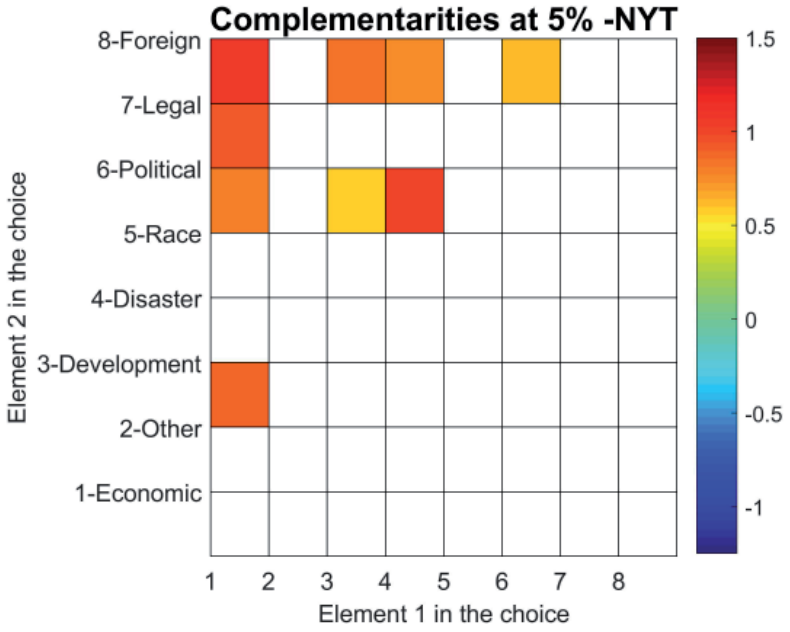

Figure C.12: Washington Post's Complementarity News using Total Publications

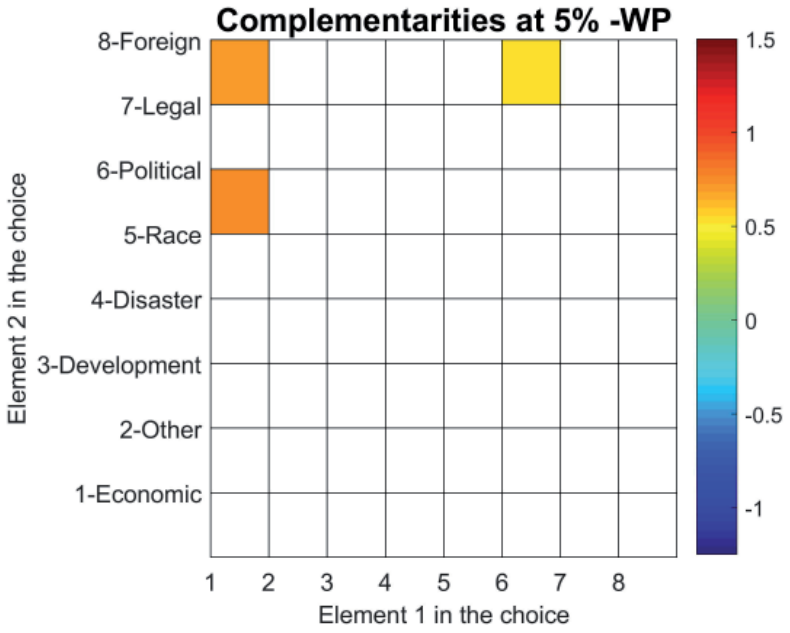

Figure C.13: The Wall Street Journal's Complementarity News using Total Publications

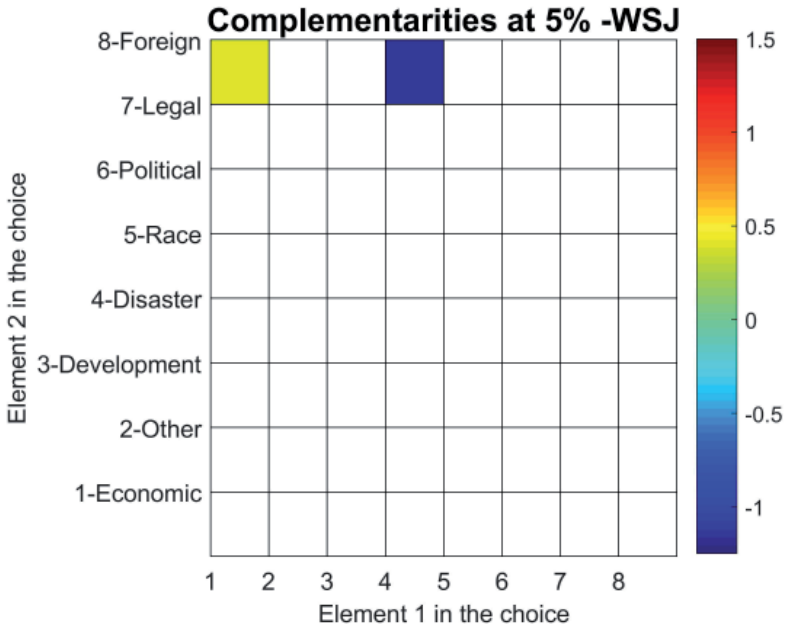

Figure C.14: The USA Today's Complementarity News controlling using Total Publications

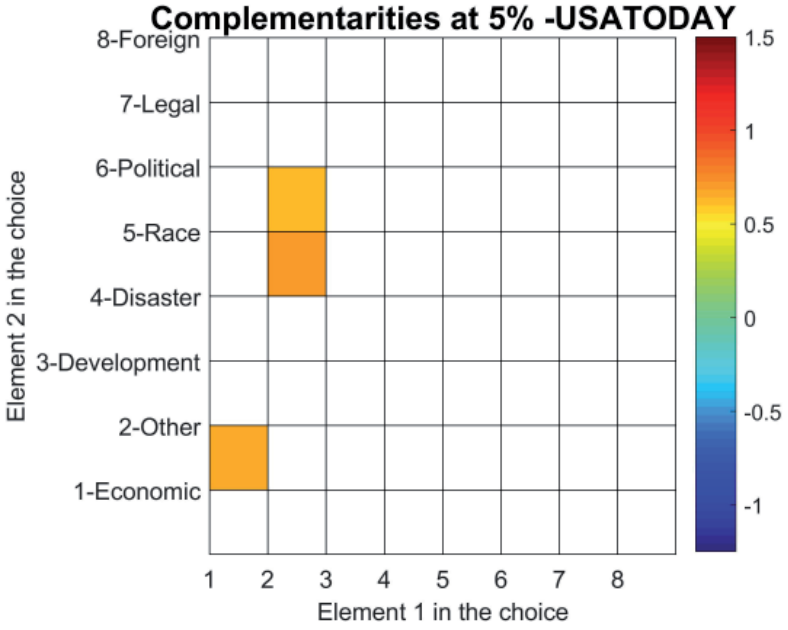

Figure C.15: The Los Angeles Times' Complementarity News using Total Publications

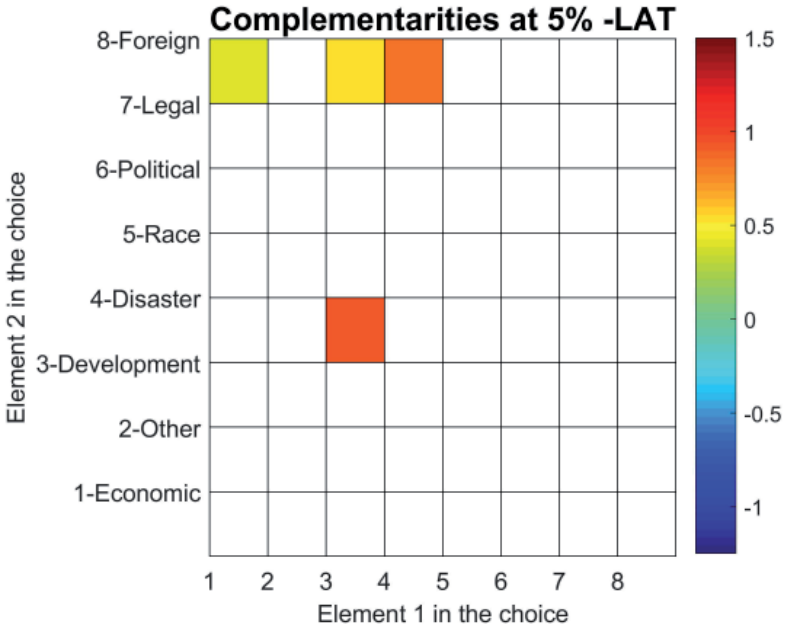

Figure C.16: New York Times' Complementarity News using Topic Share Relevance

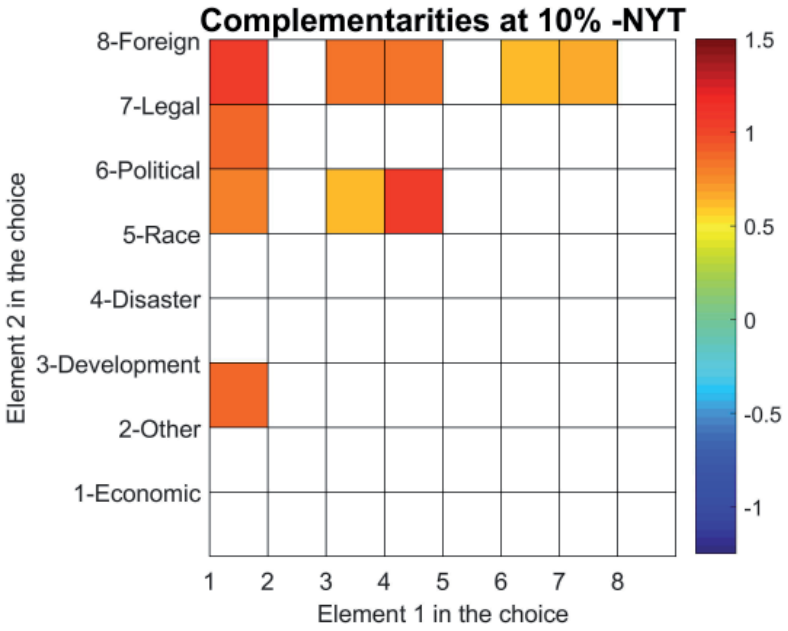

Figure C.17: Washington Post's Complementarity News using Topic Share Relevance

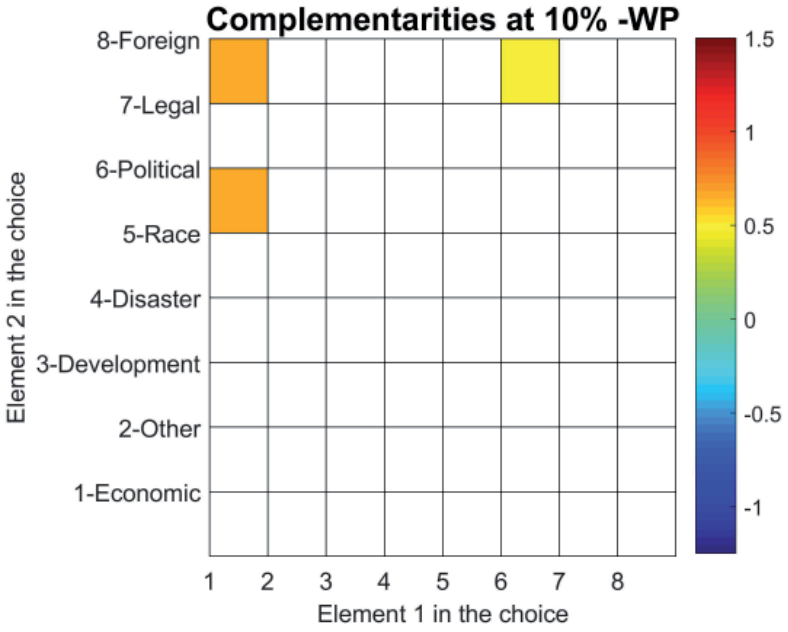

Figure C.18: The Wall Street Journal's Complementarity News using Topic Share Relevance

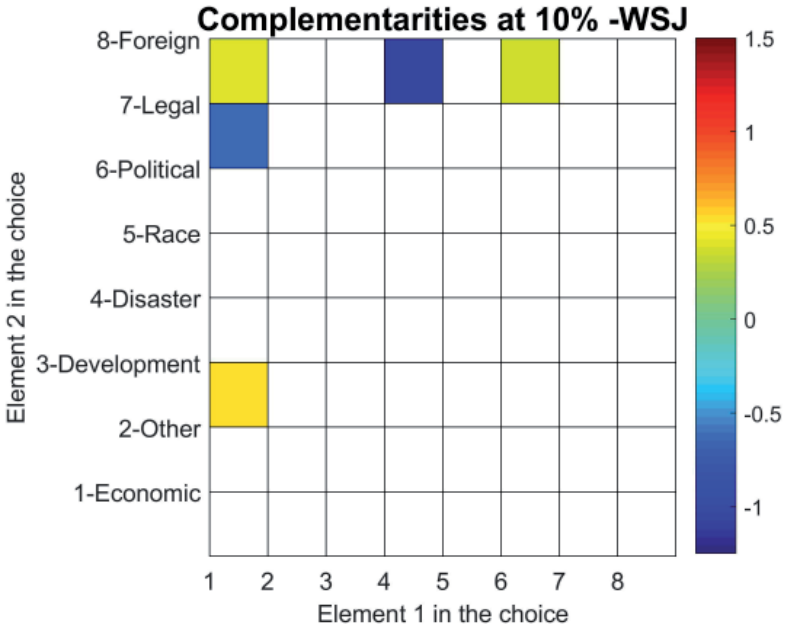

Figure C.19: The USA Today's Complementarity News controlling using Topic Share Relevance

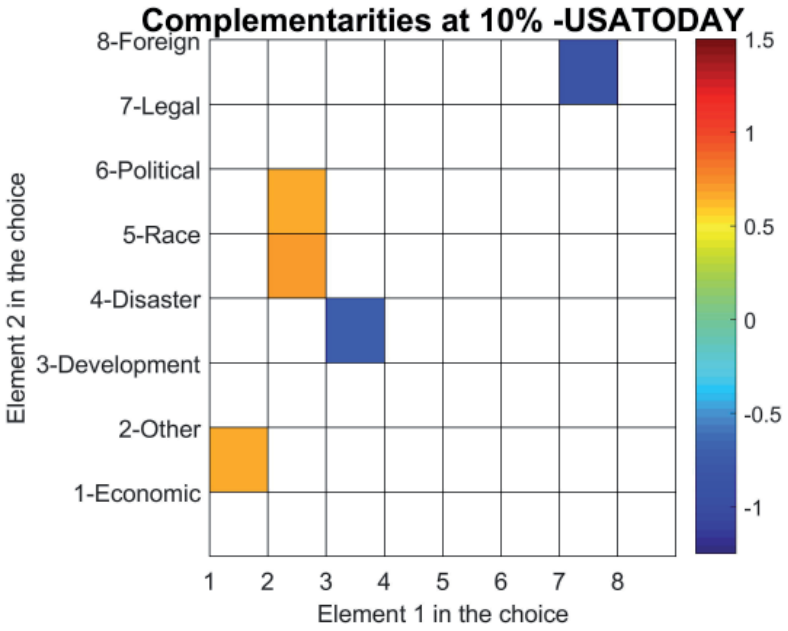

Figure C.20: The Los Angeles Times' Complementarity News using Topic Share Relevance

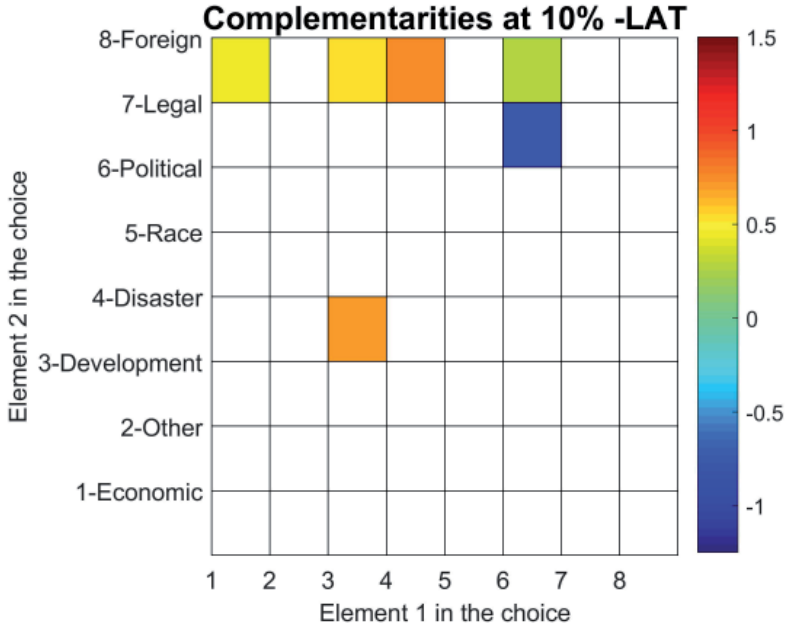

Figure C.21: New York Times' Complementarity News using Last Topic Share Relevance

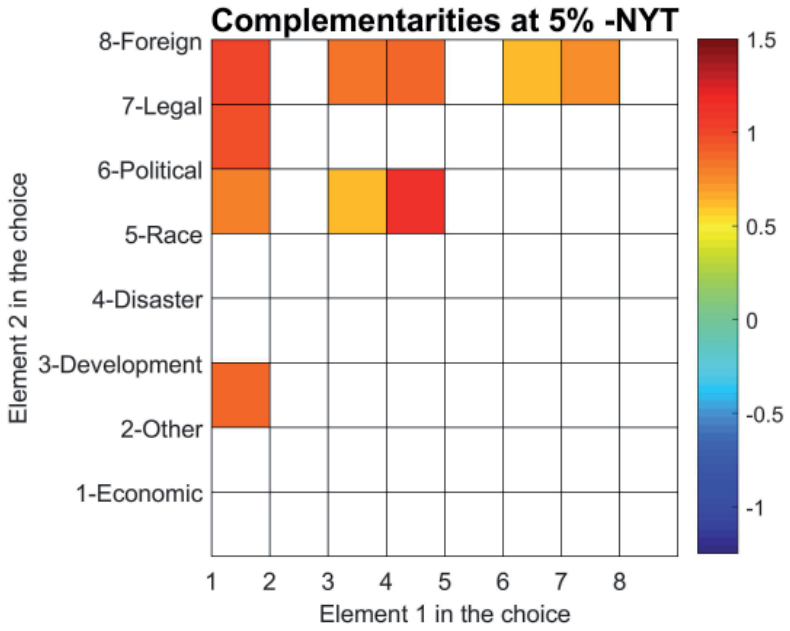

Figure C.22: Washington Post's Complementarity News using Last Topic Share Relevance

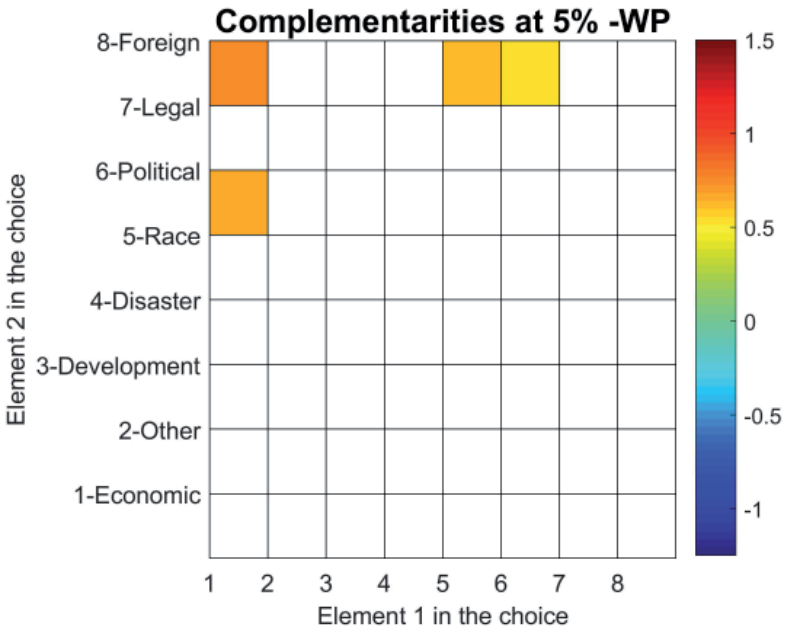

Figure C.23: The Wall Street Journal's Complementarity News using Last Topic Share Relevance

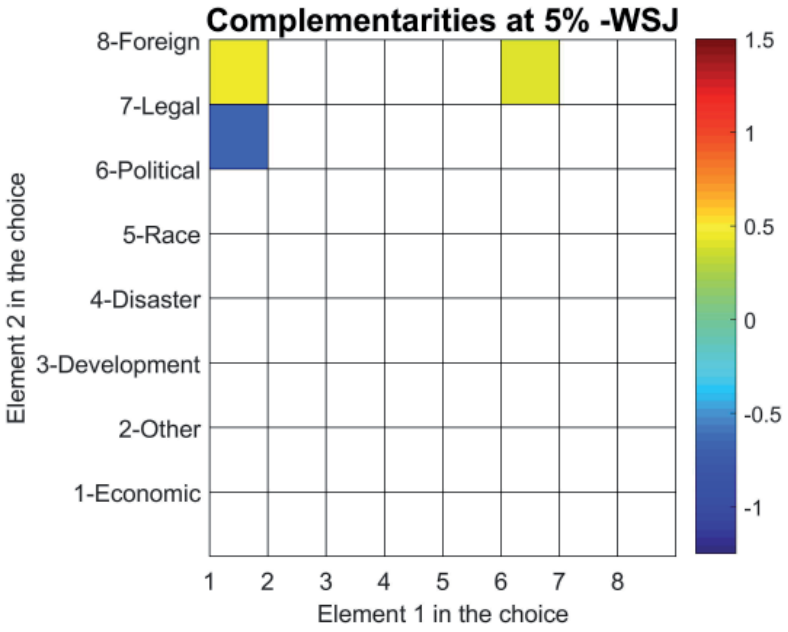

Figure C.24: The USA Today's Complementarity News controlling using Last Topic Share Relevance

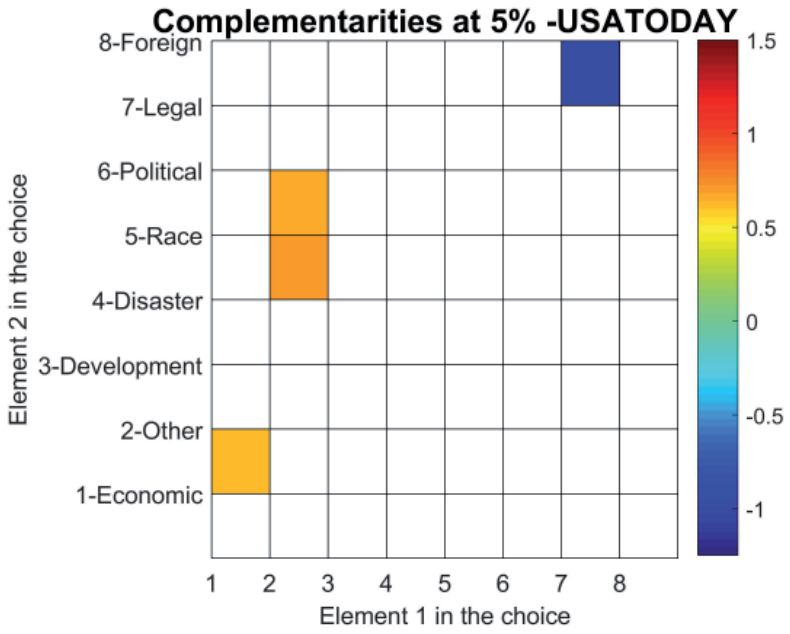

Figure C.25: The Los Angeles Times' Complementarity News using Last Topic Share Relevance

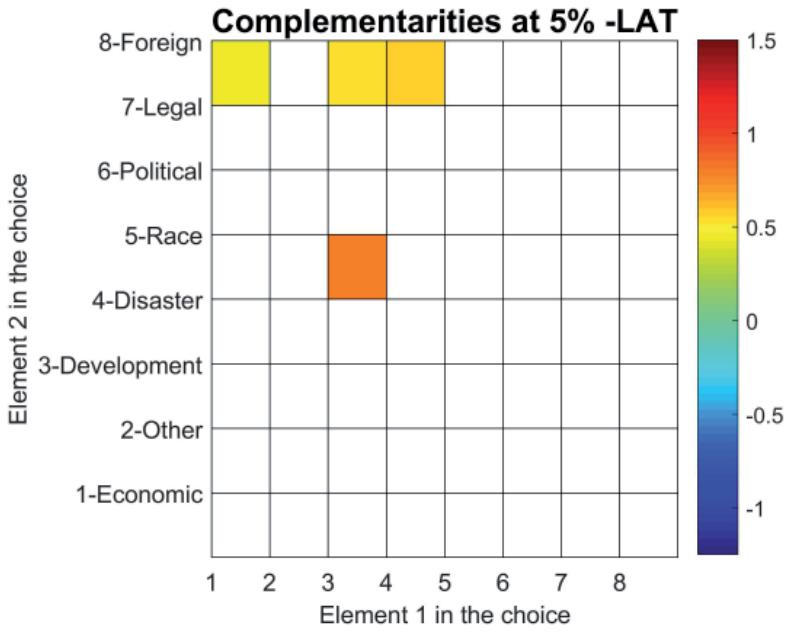

Figure C.26: New York Times' Complementarity News using Last Topic Share Relevance

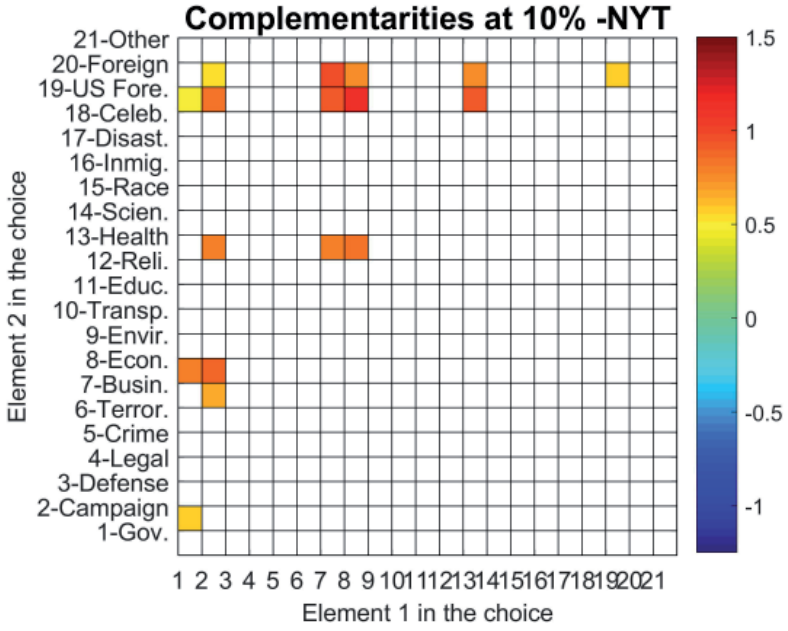

Figure C.27: Washington Post's Complementarity News using Last Topic Share Relevance

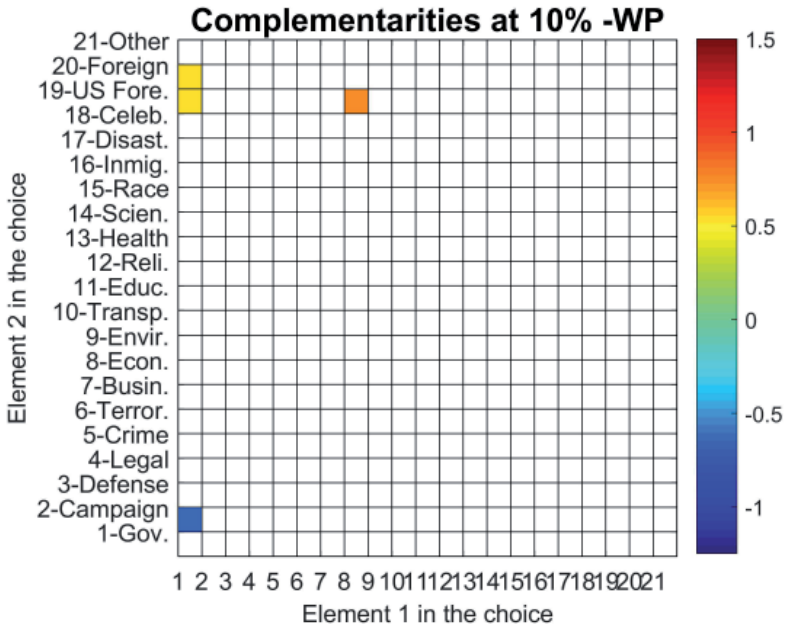

Figure C.28: The Wall Street Journal's Complementarity News using Last Topic Share Relevance

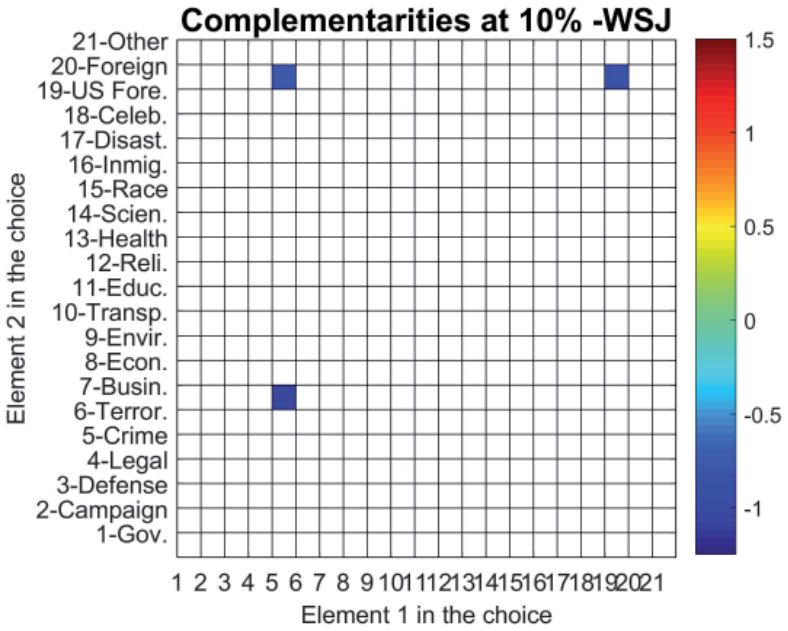

Figure C.29: The USA Today's Complementarity News controlling using Last Topic Share Relevance

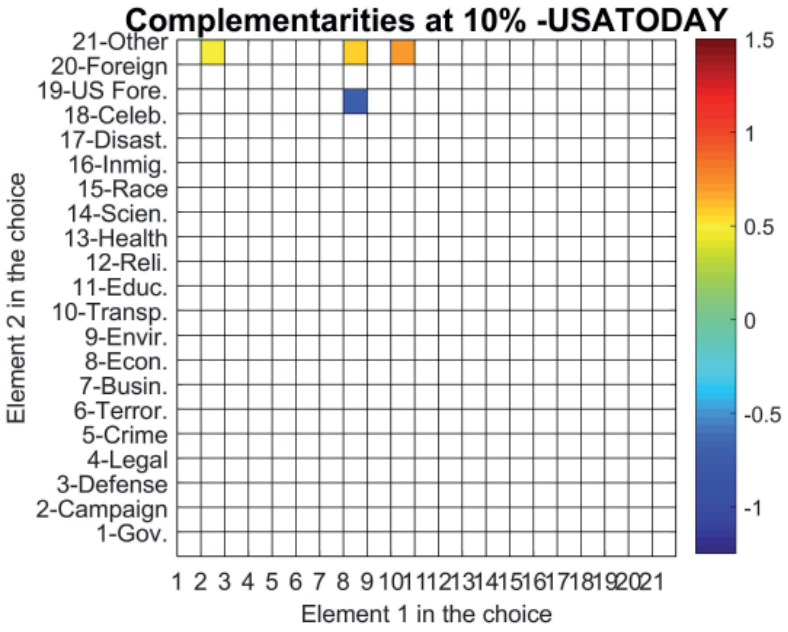

Figure C.30: The Los Angeles Times' Complementarity News using Last Topic Share Relevance

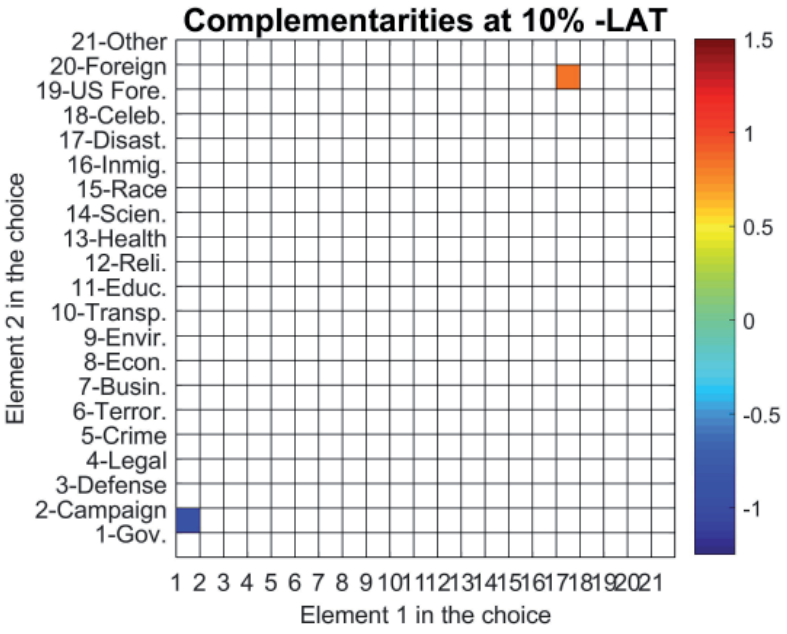

References

Yannis Bakos and Erik Brynjolfsson. Bundling information goods: Pricing, profits, and efficiency. *Management Science*, 45(12):1613–1630, 1999.

Dan Berkowitz. Refining the gatekeeping metaphor for local television news. 1990.

Carlo E Bonferroni. *Teoria statistica delle classi e calcolo delle probabilita*. Libreria internazionale Seeber, 1936.

Regina P. Branton and Johanna Dunaway. Slanted newspaper coverage of immigration: The importance of economics and geography. *Policy Studies Journal*, 37(2):257–273, 2009. ISSN 1541-0072.

Julia Cagé. Media competition, information provision and political participation. *Job Market Paper*, 2013.

Tsan-Kuo Chang and Jae-Won Lee. Factors affecting gatekeepers' selection of foreign news: A national survey of newspaper editors. *Journalism & Mass Communication Quarterly*, 69(3):554–561, 1992.

Steven E Clayman and Ann Reisner. Gatekeeping in action: Editorial conferences and assessments of newsworthiness. *American Sociological Review*, pages 178–199, 1998.

Gregory S. Crawford and Ali Yurukoglu. The welfare effects of bundling in multichannel television markets. *American Economic Review*, 102(2):643–85, April 2012.

Rafael Di Tella and Ignacio Franceschelli. Government advertising and media coverage of corruption scandals. *American Economic Journal: Applied Economics*, 3(4):119–51, October 2011.

Ruben Durante and Ekaterina Zhuravskaya. Attack when the world is not watching? international media and the israeli-palestinian conflict. *Journal of Political Economy*, Forthcoming.

Thomas Eisensee and David Strömberg. News droughts, news floods, and u. s. disaster relief. *The Quarterly Journal of Economics*, 122(2):693–728, 2007.

Lilly Fang and Joel Peress. Media coverage and the cross-section of stock returns. *The Journal of Finance*, 64(5):2023–2052, 2009. ISSN 1540-6261.

Guido Friebel and Matthias Heinz. Media slant against foreign owners: Downsizing. *Journal of Public Economics*, 120:97 – 106, 2014. ISSN 0047-2727.

Jean J. Gabszewicz, Didier Laussel, and Nathalie Sonnac. Press advertising and the political differentiation of newspapers. *Journal of Public Economic Theory*, 4(3):317–334, 2002. ISSN 1467-9779.

Johan Galtung and Mari Holmboe Ruge. The structure of foreign news the presentation of the congo, cuba and cyprus crises in four norwegian newspapers. *Journal of peace research*, 2(1):64–90, 1965.

Diego Garcia. Sentiment during recessions. *The Journal of Finance*, 68(3):1267–1300, 2013.

Diego Garcia. The kinks of financial journalism. *Available at SSRN 2517591*, 2014.

Matthew Gentzkow. Valuing new goods in a model with complementarity: Online newspapers. *American Economic Review*, 97(3):713–744, 2007.

Matthew Gentzkow and Jesse M Shapiro. What drives media slant? evidence from us daily newspapers. *Econometrica*, 78(1):35–71, 2010.

Matthew Gentzkow, Jesse Shapiro, and Daniel Stone. Media bias in the marketplace: Theory. NBER Working Papers 19880, National Bureau of Economic Research, Inc, 2014.

Fabrizio Germano and Martin Meier. Concentration and self-censorship in commercial media. *Journal of Public Economics*, 97:117 – 130, 2013. ISSN 0047-2727.

Tim Groseclose and Jeffrey Milyo. A measure of media bias. *The Quarterly Journal of Economics*, 120(4):1191–1237, 2005.

Umit.G. Gurun and Alexander.W. Butler. Don't believe the hype: Local media slant, local advertising, and firm value. *The Journal of Finance*, 67(2):561–598, 2012. ISSN 1540-6261.

James Hamilton. All the news that's fit to sell: How the market transforms information into news, regulation through revelation. *Princeton University Press*, 2003.

Matthias Heinz and Johan Swinnen. Media slant in economic news: A factor 20. *Economics Letters*, 132:18 – 20, 2015. ISSN 0165-1765.

Igal. E. Hendel, Saul Lach, and Yossi Spiegel. Consumers' activism: The cottage cheese boycott. CEPR Discussion Papers 10460, C.E.P.R. Discussion Papers, 2015.

Sture Holm. A simple sequentially rejective multiple test procedure. *Scandinavian journal of statistics*, pages 65–70, 1979.

Doh-Shin Jeon and Domenico Menicucci. Bundling and competition for slots. *American Economic Review*, 102(5):1957–85, May 2012.

Martin A. Koschat and William P. Putsis. Audience characteristics and bundling: A hedonic analysis of magazine advertising rates. *Journal of Marketing Research*, 39(2): 262–273, 2002.

Sendhil Mullainathan and Andrei Shleifer. The market for news. *American Economic Review*, pages 1031–1053, 2005.

Sophia Peterson. Foreign news gatekeepers and criteria of newsworthiness. *Journalism and Mass Communication Quarterly*, 56(1):116, 1979.

Sophia Peterson. International news selection by the elite press: A case study. *Public Opinion Quarterly*, 45(2):143–163, 1981.

Ann E Reisner. The news conference: How daily newspaper editors construct the front page. *Journalism & Mass Communication Quarterly*, 69(4):971–986, 1992.

Joseph P Romano and Michael Wolf. Stepwise multiple testing as formalized data snooping. *Econometrica*, 73(4):1237–1282, 2005.

Joseph P Romano, Azeem M Shaikh, and Michael Wolf. Hypothesis testing in econometrics. *Annu. Rev. Econ.*, 2(1):75–104, 2010.

Øystein Sande. The perception of foreign news. *Journal of Peace Research*, pages 221–237, 1971.

Ananya Sen and Pinar Yildirim. Clicks and editorial decisions: How does popularity shape online news. 2015.

Pamela J Shoemaker, Lucig H Danielian, and Nancy Brendlinger. Deviant acts, risky business and us interests: The newsworthiness of world events. *Journalism & Mass Communication Quarterly*, 68(4):781–795, 1991.

Jidong Zhou. Competitive bundling. *Econometrica*, 85(1):145–172, 2017.

BANCO DE ESPAÑA PUBLICATIONS

WORKING PAPERS

1720 LUIS J. ÁLVAREZ and ANA GÓMEZ-LOSCOS: A menu on output gap estimation methods.
1721 PAULA GIL, FRANCISCO MARTÍ, JAVIER J. PÉREZ, ROBERTO RAMOS and RICHARD MORRIS: The output effects of tax changes: narrative evidence from Spain.
1722 RICARDO GIMENO and ALFREDO IBÁÑEZ: The eurozone (expected) inflation: an option's eyes view.
1723 MIGUEL ANTÓN, SERGIO MAYORDOMO and MARÍA RODRÍGUEZ-MORENO: Dealing with dealers: sovereign CDS comovements.
1724 JOSÉ MANUEL MONTERO: Pricing decisions under financial frictions: evidence from the WDN survey.
1725 MARIO ALLOZA: The impact of taxes on income mobility.
1726 DANILO LEIVA-LEON: Measuring business cycles intra-synchronization in US: a regime-switching interdependence framework.
1727 PIERRE GUÉRIN and DANILO LEIVA-LEON: Model averaging in Markov-Switching models: predicting national recessions with regional data.
1728 MÁXIMO CAMACHO and DANILO LEIVA-LEON: The propagation of industrial business cycles.
1729 JAMES COSTAIN: Costly decisions and sequential bargaining.
1730 MARIO ALLOZA: Is fiscal policy more effective in uncertain times or during recessions?
1731 PIERRE GUÉRIN and DANILO LEIVA-LEON: Monetary policy, stock market and sectoral comovement.
1732 HENRIK JENSEN, IVAN PETRELLA, SØREN HOVE RAVN and EMILIANO SANTORO: Leverage and deepening business cycle skewness.
1733 CÉSAR MARTÍN MACHUCA: External stress early warning indicators.
1734 RODOLFO G. CAMPOS: International migration pressures in the long run.
1735 ANDREA ARIU, ELENA BIEWEN, SVEN BLANK, GUILLAUME GAULIER, MARÍA JESÚS GONZÁLEZ, PHILIPP MEINEN, DANIEL MIRZA, CÉSAR MARTÍN MACHUCA and PATRY TELLO: Firm heterogeneity and aggregate business services exports: micro evidence from Belgium, France, Germany and Spain.
1736 LEONARDO GAMBACORTA, STEFANO SCHIAFFI and ADRIAN VAN RIXTEL: Changing business models in international bank funding.
1737 ENRIQUE MORAL-BENITO and FRANCESCA VIANI: An anatomy of the Spanish current account adjustment: the role of permanent and transitory factors.
1738 MARÍA J. NIETO and LARRY D. WALL: Cross-border banking on the two sides of the Atlantic: does it have an impact on bank crisis management?
1739 JACOPO TIMINI: Currency unions and heterogeneous trade effects: the case of the Latin Monetary Union.
1740 PAULINO FONT, MARIO IZQUIERDO and SERGIO PUENTE: Subsidising mature age employment or throwing coins into a wishing well: a quasi-experimental analysis.
1741 THOMAS FUJIWARA and CARLOS SANZ: Norms in bargaining: evidence from government formation in Spain.
1742 ENRIQUE ALBEROLA, ÁNGEL ESTRADA and FRANCESCA VIANI: Global imbalances from a stock perspective.
1743 ÓSCAR ARCE, RICARDO GIMENO and SERGIO MAYORDOMO: Making room for the needy: the credit-reallocation effects of the ECB's corporate QE.
1744 M. D. GADEA-RIVAS, ANA GÓMEZ-LOSCOS and EDUARDO BANDRÉS: Clustering regional business cycles.
1745 NEZIH GUNER, YULIYA KULIKOVA and JOAN LLULL: Marriage and health: selection, protection, and assortative mating.
1746 SERGIO MAYORDOMO and MARÍA RODRÍGUEZ-MORENO: Did the bank capital relief induced by the supporting factor enhance SME lending?
1747 KATALIN BODNÁR, LUDMILA FADEJEVA, MARCO HOEBERICHTS, MARIO IZQUIERDO PEINADO, CHRISTOPHE JADEAU and ELIANA VIVIANO: Credit shocks and the European Labour market.
1748 PIERRE GUÉRIN, DANILO LEIVA-LEON and MASSIMILIANO MARCELLINO: Markov-switching three-pass regression filter.
1749 ISABEL ARGIMÓN: Decentralized multinational banks and risk taking: the Spanish experience in the crisis.
1750 BING XU: Permissible collateral and access to finance: evidence from a quasi-natural experiment.
1751 GERGELY AKOS GANICS: Optimal density forecast combinations.
1801 OLYMPIA BOVER, LAURA HOSPIDO and ERNESTO VILLANUEVA: The impact of high school financial education on financial knowledge and choices: evidence from a randomized trial in Spain.
1802 IGNACIO HERNANDO, IRENE PABLOS, DANIEL SANTABÁRBARA and JAVIER VALLÉS: Private Saving. New Cross-Country Evidence Based on Bayesian Techniques.

1803 PABLO AGUILAR and JESÚS VÁZQUEZ: Term structure and real-time learning.
1804 MORITZ A. ROTH: International co-movements in recessions.
1805 ANGELA ABBATE and DOMINIK THALER: Monetary policy and the asset risk-taking channel.
1806 PABLO MARTÍN-ACEÑA: Money in Spain. New historical statistics. 1830-1998.
1807 GUILHERME BANDEIRA: Fiscal transfers in a monetary union with sovereign risk.
1808 MIGUEL GARCÍA-POSADA GÓMEZ: Credit constraints, firm investment and growth: evidence from survey data.
1809 LAURA ALFARO, MANUEL GARCÍA-SANTANA and ENRIQUE MORAL-BENITO: On the direct and indirect real effects of credit supply shocks.
1810 ROBERTO RAMOS and CARLOS SANZ: Backing the incumbent in difficult times: the electoral impact of wildfires.
1811 GABRIEL JIMÉNEZ, ENRIQUE MORAL-BENITO and RAQUEL VEGAS: Bank lending standards over the cycle: the role of firms' productivity and credit risk.
1812 JUAN S. MORA-SANGUINETTI and ROK SPRUK: Industry vs services: do enforcement institutions matter for specialization patterns? Disaggregated evidence from Spain.
1813 JAMES CLOYNE, CLODOMIRO FERREIRA and PAOLO SURICO: Monetary policy when households have debt: new evidence on the transmission mechanism.
1814 DMITRI KIRPICHEV and ENRIQUE MORAL-BENITO: The costs of trade protectionism: evidence from Spanish firms and non-tariff measures.
1815 ISABEL ARGIMÓN, CLEMENS BONNER, RICARDO CORREA, PATTY DUIJM, JON FROST, JAKOB DE HAAN, LEO DE HAAN and VIKTORS STEBUNOVS: Financial institutions' business models and the global transmission of monetary policy.
1816 JOSE ASTURIAS, MANUEL GARCÍA-SANTANA and ROBERTO RAMOS: Competition and the welfare gains from transportation infrastructure: evidence from the Golden Quadrilateral of India.
1817 SANDRA GARCÍA-URIBE: Multidimensional media slant: complementarities in news reporting by US newspapers.

BANCO DE **ESPAÑA**
Eurosistema

Unidad de Servicios Auxiliares
Alcalá, 48 - 28014 Madrid
E-mail: publicaciones@bde.es
www.bde.es