Quantitative Methods in Economic Research on Media and Communication

Marcel Garz*

Jönköping University

November 2018

1. Introduction

This chapter reviews recent trends in empirical research on media and communication, especially in reference to (a) data collection, (b) measurement, and (c) causal inference. Researchers have been making increasing efforts to *collect data* by harvesting publicly available information from online sources; for instance, by retrieving data via websites' application programming interfaces (APIs) and by using methods to crawl, scrape, and parse web data. At the same time, there has been a trend to compile original datasets by digitizing information from historical sources. When it comes to *measurement*, the literature has benefited from automated methods that allow to analyze text as data, including text mining and natural language processing. Due to their increasing sophistication and the rising processing power of computers, such tools are frequently chosen when human content analysis would be too slow or too costly. Finally, a growing number of studies has been applying techniques that allow for *causal inference* with observational data. Methods such as instrumental variable regression, differences-in-differences, and regression discontinuity have become popular alternatives when controlled laboratory experiments are not feasible or desirable. This chapter does not attempt to systematically survey the literature but aims to provide selected examples of studies that illustrate the applicability of tools to gather media-related data, create text-based measures, and distinguish cause and effect. The review has a stronger focus on studies in economics than management, due to the widespread use of the above-mentioned methods in the former discipline.

---

* In preparation for the volume Management and Economics of Communication – Handbooks in Communication Science, edited by M. Bjørn von Rimscha and Stephanie Kienzler.
Email: marcel.garz@ju.se

2. Data collection

This section focuses on data sources and methods of data collection that have been increasingly used recently, including the digitization of information from historical sources, the use of digital newspaper archives and broadcasting transcripts, and the automated retrieval of information from online sources. Data from these sources are often used in studies that address topics related to news markets, media effects, media bias, and the relationship between mass communication and politics.

2.1 Historical data sources

Early mass media – such as newspapers, radio, and television – have played a pivotal role in human history. A recent strand of literature investigates research questions related to the introduction of these media. At the intersection of media economics and economic history, this literature studies the role of economic and regulatory conditions for the dissemination of new media technologies, as well as the effect of these technologies on economic, political, and social outcomes. The research usually involves great efforts to digitize analog data sources, including historical year books, reference books, and directories. While it is possible to use optical character recognition (OCR) in some cases, it is often necessary to manually locate and process the desired information. Thus the costs of retrieving data from such sources are high. Since the resulting datasets have not been investigated before, analyzing them usually allows for original contributions though.

For example, Gentzkow, Shapiro, and Sinkinson (2011) and (2014) have collected comprehensive data on several thousand US newspapers between 1869 and 2004. Using sources such as the American Newspaper Annual, the American Newspaper Directory, and the Editor and Publisher Yearbook, the authors compile a dataset that includes information on newspaper entry and exit, newspapers' location, circulation, subscription and copy prices, as well as their political affiliation. They use these data to investigate the effect of newspapers on electoral politics, among other things. For a study on the impact of media competition on the quantity and quality of news, Cagé (2017) collects similar data on French newspapers as of 1944, including revenues, expenditures, and demographic information of journalists. Cagé and Rueda (2016) use the World Atlas of Christian Missions to obtain information on printing activities by mission stations in sub-Saharan Africa in the early 20th century. The authors use these data to evaluate the effects of the printing press on newspaper readership, education, and political participation.

In a study on the empowerment of the Nazis, Adena et al. (2015) work with historical documents to collect information on the location and power of radio transmitters in Germany in the early 1930s. They also add region- and city-level data on radio subscription rates, the number of newspapers and cinemas, and the number of speeches given by Hitler. With these data, the authors are able to estimate effects on Nazi party membership, election outcomes, and anti-Semitic acts. Using the availability of radio in US households in the same period of time, Strömberg (2004) studies the impact of the technology on public spending. Among other things, counties with higher degrees of radio penetration received more funds from New Deal relief programs, due to the positive impact of radio on voter information.

Gentzkow (2006) evaluates implications of the introduction of television in the US between 1940 and 1970. The author gathers county-level data on the availability of TV stations from various issues of the Television Factbook. The findings indicate a negative effect on voter turnout, likely because television consumption partially replaced the use of other media with more political coverage. Using similar data, Baker and George (2010) find that exposure to television advertising increased household debt, whereas the results of Campante and Hojman (2013) indicate a decrease in ideological polarization of the US Congress. Bursztyn and Cantoni (2016) collect data on the location and transmission power of West German antennas in 1989 to study effects of reception of Western television in the German Democratic Republic. According to their results, East Germans living in areas with access differed in their consumption behavior after reunification from those living in areas without access.

2.2 Digital newspaper and TV archives

Research on news markets often requires measures of news output, which is why many recent studies have greatly relied on full-text archives of newspaper articles and, to a lesser extent, transcripts of television news. Data extracted from such sources have been used to create variables that count news items, or that capture forms of media bias or media sentiment.

The NewsLibrary and ProQuest databases mostly offer access to US newspapers. Among other things, researchers have used these sources to investigate ideological media bias (Gentzkow and Shapiro, 2010; Puglisi and Snyder, 2015), voter information (Snyder and Strömberg, 2010), and effects on financial markets (Engelberg and Parsons, 2011). The Factiva and Nexis databases include newspapers from all over the world, although their focus is on Western and English-speaking titles. Data extracted from these archives have been used in studies on media bias against foreign companies in Germany (Friebel and Heinz, 2014), advertising bias in US newspapers (Gurun and Butler, 2012), and coverage about government scandals in the UK (Latham, 2015).

All databases have browser interfaces that support advanced keyword searches, including Boolean operators. Searching the NewsLibrary database is free but users are charged for full-text access. ProQuest, Factiva, and Nexis require a subscription, which many universities and libraries have. However, in the case of academic subscribers, there are usually limits on the number of search results and especially the number of downloads of full-text articles.

A commonly used source of data on US newscasts is the Vanderbilt Television News Archive (e.g., Eisensee and Strömberg, 2007; Durante and Zhuravskaya, 2018). This archive offers access to ABC, CBS, and NBC as of 1968, and to CNN and Fox News for more recent periods. Their browser interface allows for keyword searches on story-level abstracts. The retrieval of results is free but access to video clips requires a subscription. Other studies rely on non-public or proprietary data (Durante and Knight, 2012; Martin and McCrain, 2018) or use broadcasting transcripts from Nexis (Martin and Yurukoglu, 2017).

2.3 Web data

As many other disciplines, the field has greatly benefited from the possibility to gather all sorts of data from online sources. The Internet has vastly increased the amount of information accessible to researchers. Most of these data are pre-structured and can thus be harvested at low cost. Researchers have been increasing their efforts to collect data by crawling, scraping, and parsing publicly available information from websites. For simple applications, it is possible to apply ready-to-use packages (e.g., in Python or R) to implement the retrieval and organization of information. In most cases, it is necessary to program customized routines though, which can be a complex task if a website structures the desired information in a non-trivial way. In addition to technical aspects, further limitations may arise due to the terms and conditions of websites and privacy issues. In the case of media economics, the attractiveness of web data mostly arises from the prospect of creating measures of news output. For instance, Berger and Milkman (2012), Budak, Goel, and Rao (2016), and Garz et al. (2018) collect data from US online news outlets to construct measures of content, placement, and popularity, whereas Szeidl and Szucs (2016) and Simonov and Rao (2018) retrieve information from Hungarian and Russian outlets, respectively. Going one step further, Jo (2017) collects data on South Korean users via a custom-made and self-distributed mobile news application. The app not only monitors users' news consumption, but also facilitates the implementation of field experiments. The latter feature allows the author to study the causal link between selective exposure and polarization.

Websites and online platforms sometimes offer access to their data via APIs. Information retrieval through this channel is usually fast and reliable, and the data can often be downloaded in pre-formatted blocks. Lee, Hosanagar, and Nair (2017) use data obtained via the Facebook API to analyze effects of social media advertising on user engagement (i.e., the likelihood to like, share, or comment). Garz, Sörensen, and Stone (2018) are interested in the same outcome variable but investigate the role of politically like-minded versus counter-attitudinal information. Müller and Schwarz (2017) use the Facebook API to construct measures of anti-refugee sentiment, which they relate to xenophobic violence. Data downloaded via the Twitter API have been analyzed as well. Halberstam and Knight (2016) investigate connections between like-minded users and exposure to like-minded information, whereas Allcott, Gentzkow, and Yu (2018) study user interactions with misinformation.

Other researchers take advantage of click data provided by specialized market research companies. For example, comScore has shared their proprietary data for academic purposes. Studying ideologically motivated news consumption, Gentzkow and Shapiro (2011) and Lelkes, Sood, and Iyengar (2017) use samples of individuals who have agreed that the company tracks their browsing activities. ComScore maintains a panel of over one million users who installed software that monitors their online behavior via desktop, tablet, or mobile Internet, in exchange for cash and other rewards. The data allow researchers to describe and analyze people's browsing activities in great detail, with the caveat that panelists might behave differently than users who are not tracked. The same applies to similar data collected by Microsoft via a voluntary extension to its browser. These data have been used by Flaxman, Goel, and Rao (2016) to investigate the role of social media for partisan selective exposure.

3. Text-based measures

Researchers in media economics have a natural interest to capture characteristics of media reports, social media posts, or other communication on websites. These characteristics can be coded manually; for example, by conducting human content analysis. For instance, the Swiss-based company Media Tenor International analyzes media reports published by news outlets from all over the world. Their data – which have been frequently used for academic purposes (e.g., Tausch and Zumbuehl, 2017; Beckmann, Dewenter, and Thomas, 2017; Ulbricht, Kholodinin, and Thomas, 2017) – include information about the topic, the referenced time, region, and actors, as well as the tone of individual news items. Recently, researchers have also used crowd-sourced content analysis, which involves click workers recruited on Amazon Mechanical Turk (Budak, Goel, and Rao, 2016).

However, an increasing amount of studies have been applying automated methods from areas such as information retrieval, text mining, and natural language processing. As a first step, these methods require text to be cleaned. The cleaning usually comprises the removal of stop words (e.g., "the", "is", and "on"), punctuation, numbers, and extra whitespace. Words can then be stemmed to avoid complications related to grammatical issues – such as conjugation or comparison – and converted to lower case. The clean text is usually stored as a term-frequency matrix. In this form, simple measures of news output can be created by counting indicative terms. For example, Groseclose and Milyo (2005) propose to capture media bias by counting citations of think tanks in US media reports. Qin, Strömberg, and Wu (2018) count occurrences of various items in Chinese media, including mentions of political leaders, citations of the state-run press agency Xinhua News, and references to politically relevant events. The authors use principal component analysis to condense these counts into a single measure of political bias.

Term-frequency matrices can also be used to extract the sentiment from text. Researchers often use algorithms that retrieve weights from sentiment dictionaries, based on which sentiment scores related to the term frequencies observed in the text segments under consideration are computed. Basic sentiment dictionaries – such as SentiWords (Gatti, Guerini, and Turchi, 2016) – rank words on a one-dimensional scale from negative to positive. More complex dictionaries feature the ranking of words on further dimensions. For example, the Harvard IV-4 dictionary comprises of 77 categories, including emotions, roles, and motivations. Sentiment dictionaries have also been developed for different languages and specific contexts. For instance, Loughran and McDonald (2011) compile a dictionary that captures sentiment in financial applications. Other dictionaries support the detection of word classes, negations, and amplifiers, which allows to take syntax into account (see, for instance, the R package "Syuzhet" by Matthew Jockers, which currently offers this kind of functionality for at least 23 languages).

The idea of comparing texts with dictionaries has been applied by the media bias literature in a more general way. Instead of using dictionaries, researchers have compared the language used in news reports with the language used in reference texts that have known ideological positions. Gentzkow and Shapiro (2010) and Martin and Yurukoglu (2017) create a measure of political slant by using speech protocols of Democratic and Republican politicians in the US Congress. Garz, Sörensen, and Stone (2018) capture the slant of posts by the Facebook pages of German news outlets by comparing their language with that used in the election programs of the main political parties. Beattie (2017) measures the tone of news reports related to climate change by using reports of the Intergovernmental Panel on Climate Change (IPCC) and the more skeptical Nongovernmental International Panel on Climate Change (NIPCC) as reference texts. For their investigation of censorship in China, Qin, Strömberg, and Wu (2017) analyze several million posts published on the popular social media platform Sina Weibo. To identify unofficial government accounts, the authors use a support vector machine and compare the language of potential government accounts with the language used by official government accounts.

The popularity of topic modeling has been increasing as well in the field. Many research questions require to distinguish between topics addressed by media reports. Topic models use algorithms that cluster pieces of text by the frequency of the terms that these pieces include. Using data from an Indian newspaper, Sen and Yildirim (2016) apply topic modeling to find clusters of successive news stories on the same underlying common topic. Simonov and Rao (2018) use a similar approach to identify government-sensitive news in Russia. For their investigation on the decline of news coverage on local politics in the US, Martin and McCrain (2018) run topic models on broadcast transcripts to distinguish segments with local and national relevance.

4. Causal inference

This section discuses four approaches that support causal inference when analyzing observational data. The review starts with the instrumental variable technique, a method that has been increasingly used since the 1990s to disentangle cause and effect, followed by the differences-in-differences and regression discontinuity approaches. The section ends with a discussion of controlled field experiments.

4.1 Instrumental variables

The instrumental variable approach was first proposed by Wright (1928). For a long time, it was mostly used to address problems related to measurement error and omitted variable bias, until Angrist (1990) and Angrist and Krueger (1991) emphasized the role of instrumental variables in identifying causal relationships. The question of causality is relevant in most empirical applications, but it is of particular importance when investigating media markets. The reason is that prices, quantities, product quality, and external effects of media consumption are observed simultaneously. For example, in the market for news, the selection, presentation, and evaluation of news items could be driven by the supply side, if ideologically-motivated media owners, editors, or journalists seek to influence public opinion. The news output could also be driven by the demand side though, because profit-maximizing media companies can increase their revenue by catering to the preferences of their recipients. Thus regressing measures of news output on consumer preferences does not necessarily reveal demand-side effects, and a correlation between news output and outlet ideology is not sufficient to identify supply-side effects.

In their study on political bias in US newspapers, Gentzkow and Shapiro (2010) address the simultaneity by using religiosity as an instrumental variable for consumer preferences. They argue that religiosity is a strong predictor of demand for conservative reporting, and that religiosity is unlikely to be affected by newspaper coverage in the short run. According to their findings, the part of consumer preferences that is driven by religiosity explains some of the differences in the political bias across newspapers, which plausibly identifies demand-side effects.

To illustrate the approach formally, assume that the researcher is interested in estimating the effect of an explanatory variable x_i on an outcome variable y_i. Ignoring control variables, the standard approach would be an ordinary least squares (OLS) regression of the following form:

$$y_i = a_1 + a_2 x_i + e_i \tag{1}$$

where i indices observations, a_1 is a constant, and a_2 is the parameter of interest. However, with OLS, a_2 does not identify the effect of x_i on y_i, if in reality y_i also affects x_i (reverse causality), or if there is a third variable that influences both y_i and x_i (omitted variable bias). The identification problem can be tackled if the researcher has a variable z_i that (a) is exogenous to y_i and x_i, (b) correlates with the endogenous regressor x_i, and (c) affects the outcome variable y_i exclusively through x_i. If such a valid instrument z_i exists, the causal effect of x_i on y_i can be estimated by two-stage least squares:

$$x_i = b_1 + b_2 z_i + \varepsilon_i \tag{2}$$

$$y_i = c_1 + c_2 \hat{x}_i + \epsilon_i \tag{3}$$

Equation (2) denotes the first stage and is used to estimate the effect of the instrument z_i on the x_i. In equation (3), y_i is then regressed on the predicted values of the endogenous regressor $\hat{x}_i$. Thus the parameter c_2 captures the causal effect, which is usually interpreted as a local average treatment effect.

In practice, valid instrumental variables are rare. The instruments most often used in media-related research are either based on random technological conditions or random newsworthy events that compete with the news coverage in question for attention. Examples of studies falling in the first category include Falck, Gold, and Heblich (2015) and Miner (2015), both of which exploit random differences in the rollout and penetration of high-speed Internet in Germany and Malaysia, respectively. More precisely, they use the distance of households to the main distribution frame, as municipalities closely located to the next hub were more likely to receive a high-speed connection early on. Distance to the main distribution frame exogenously predicts Internet penetration, which in turn explains differences in Internet usage and resulting effects on voting behavior. The empirical strategy addresses the possibility of reverse causality (i.e., citizens could vote for politicians that lobby for a rollout in their region) and omitted variable bias due to demographic differences across regions. Enikolopov, Petrova, and Zhuravskaya (2011) study voting behavior too, but they exploit differences in people's exposure to the only independent national TV channel in Russia. Their instrument is based on the strength of the broadcasting signal, which varied for geographical and other random reasons. Martin and Yurukoglu (2017) construct an instrument based on the channel positions of Fox News, CNN, and MSNBC in cable television lineups. Channel positions vary randomly across US counties. However, a channel is watched more if it has a low position, which allows the authors to identify the effect of exposure to the different news channels on voting. Sen and Yildirim (2016) investigate variation in the popularity of online articles published by an Indian newspaper. The authors use power shortages and rainfall as instruments to identify effects of article views on total coverage of a topic.

The second group of often-used instrumental variables exploit random variation in news coverage caused by unrelated but competing events. On a given day, news outlets have finite capacities to cover different stories, while recipients are limited in their attention to these stories. Thus only a subset of all news is covered, and particularly newsworthy events sometimes cause a congestion of the news agenda. Eisensee and Strömberg (2007) take advantage of this mechanism in their study on US disaster relief. They construct a measure of news pressure based on the length of the top three news segments of ABC, CBS, and NBC. Their findings suggest that other newsworthy events – such as the Olympic Games – crowd out news coverage of otherwise similar disasters, which affects the propensity of the US government to support the country or region in question. Webbink, van Erp, and van Gastel (2015) investigate if showing suspects of crime in a Dutch TV show affects the likelihood that these suspects are apprehended. To identify the causal effect, the authors use exogenous variation in the number of viewers of the crime TV show caused by Champions League games that are occasionally broadcasted at the same time on different TV channels. Garz and Sörensen (2017) create an instrumental variable that measures the length of the cover story of the German daily Frankfurter Allgemeine Zeitung. They use this instrument to show that competing events reduce news coverage about criminal investigations into German politicians, which in turn decreases the likelihood that these politicians resign. Jetter (2017) uses the occurrence of natural disasters to study the effects of terrorism coverage by the New York Times on subsequent terrorist attacks. Using a similar instrument, Garz and Pagels (2018) show that media attention to court trials involving celebrity tax offenders increases the likelihood that other tax payers voluntarily declare taxes they evaded.

As these examples show, instrumental variables are often used in observational studies. Because a valid instrument is as good as randomly assigned, the setting is sometimes referred to as a natural experiment. This kind of experiment can be an option if a controlled experiment is not feasible, or if researchers want to take advantage of real-world data instead of some artificial laboratory environment. The rarity of valid instruments limits the researcher's ability to manipulate the desired treatment though. In addition, the empirical applications are usually quite context-specific, which can make it difficult to generalize findings. Another disadvantage it that the validity of an instrumental variable cannot be entirely tested. Specifically, it is not possible to verify the so-called exclusion restriction (i.e., the instrument must affect the outcome variable only through the endogenous regressor). Thus researchers have to rely on theoretical arguments that this condition is met in their setting.

4.2 Differences in differences

Another approach that can help to identify causal effects is the differences-in-differences method. Sometimes referred to as a pretest-posttest controlled design, the approach compares the differential effect of an intervention on some outcome in a treatment and a control group. The 2×2 setting can be implemented by using repeated cross-sections of observational data or, better yet, panel data (i.e., repeated observations on the same individual). These data need to include pre- and post-treatment observations on both the treatment and control group, as illustrated in Table 1.

Table 1: The differences-in-differences setting

	Pre-treatment	Post-treatment
Treatment group	untreated	treated
Control group	untreated	untreated

In this setting, the average change in the control group after the treatment serves as a counterfactual for the treatment group. The underlying assumption is that in the absence of the treatment, both groups would have developed in similar ways. This requirement is also referred to as the common trends assumption, which researchers usually spend great efforts to verify. When plotting the outcome variable over time, both groups should ideally exhibit similar trends before the treatment, and different ones afterwards, as illustrated in Figure 1.

Figure 1: Hypothetical pre- and post-treatment trends

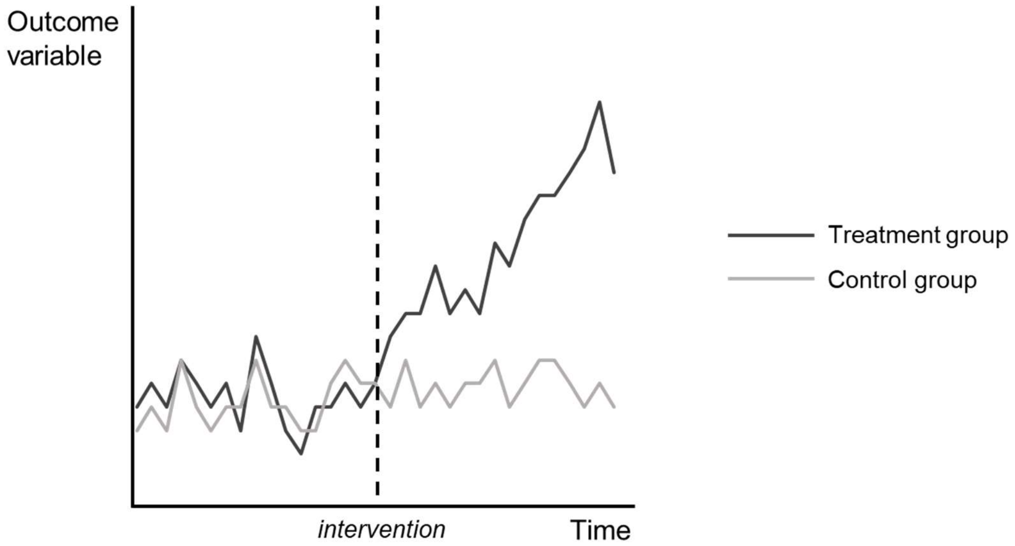

The differences-in-differences model can be estimated by OLS:

$$y_{it} = a_1 + a_2 treated_{it} + a_3 post_{it} + a_4 treated \times post_{it} + e_{it} \tag{4}$$

where $treated$ is a binary variable that takes the value 1 if individual i belongs to the treatment group and 0 if she belongs to the control group. The binary variable $post$ indicates if the outcome y is observed at a point of time t after the treatment. Accordingly, the interacted term $treated \times post$ relates to individuals after they received the treatment. The coefficient of interest is a_4, which measures the change in the outcome variable due the treatment relative to the control group. Thus the coefficient provides the average treatment effect on the treated.

Barone, D'Acunto, and Narciso (2015) investigate how the switch from analog to digital TV in Italy affected voters. The switch to digital TV increased the number of TV channels aired by networks with no ties to Berlusconi or to the government. The authors exploit the fact that the switch-off dates randomly varied between 2008 and 2012 across Italian regions, which makes it possible to compare the voting behavior of individuals that received the "digital TV treatment" with those that had not. The differences-in-differences approach allows to isolate the effect of exposure to digital TV, which the authors estimate to have decreased the vote share for Berlusconi's coalition by between 5.5 and 7.5 percentage points.

Peukert, Claussen, and Kretschmer (2017) use a differences-in-differences approach to study the effects of online piracy of movies on box office revenues. For that purpose, the authors exploit the sudden shutdown of Megaupload, a popular file hosting platform at the time. Between 2005 and 2012, Megaupload could be used to download or stream pirated movies. In January 2012, the website was deactivated after their owners were arrested and company premises raided. Using movies as observation units, the authors estimate the effects of the shutdown on box office revenues by comparing movies previously hosted on Megaupload to those movies that were not available on the platform. They find that only movies circulated in many cinemas benefited from the closure, while the overall effect on revenues was negative.

Martin and McCrain (2018) employ a differences-in-differences strategy to investigate the effects of the 2017 takeover of numerous local US TV stations by the Sinclair Media Group. The authors use data on broadcast transcripts to evaluate how the content of the acquired stations changed in comparison to stations that were not acquired by Sinclair. Their results indicate that the takeover led to an increase in coverage of national politics, whereas coverage of local politics declined. Since viewership decreased in response to this change, they conclude that the takeover was mainly supply driven.

As these examples illustrate, the differences-in-differences approach is well-suited to investigate the effects of policy changes, changes in the law, changes in technology, or other shocks that change certain conditions at a specific point of time. Since the data requirements are not extensive, the approach may be a reasonable alternative when a controlled experiment is not possible or an instrumental variable not available. If the common trends assumption plausibly holds in given context, the approach accounts for unobserved confounders, especially the possibility of another shock that might have affected the outcome variable at the same time as the treatment of interest. A disadvantage is that it can be difficult to find individuals or other units of observation that are as good as randomly assigned to the treatment and control groups. In many cases, there is the possibility of pre-selection, or that the assignment of the treatment follows some deliberate considerations. In addition, conventionally computed standard errors can be misleading in most differences-in-differences models. The reason is that the outcome variable usually varies at a much finer level than the treatment variable, which deflates the standard errors if there is correlation within units of observation (Bertrand, Duflo, and Mullainathan, 2004). However, most statistical software allows to correct for the deflation by computing standard errors clustered by observation unit.

4.3 Regression discontinuity

The regression discontinuity approach is another method to estimate treatment effects in non-experimental settings. The approach was first proposed by Thistlethwaite and Campbell (1960) and can be applied if the treatment depends on some assignment variable. More precisely, individuals or other units of observation receive the treatment if the assignment variable exceeds a certain threshold. The approach relies on comparing the outcome variable for observations just below and just above this threshold. The are two main assumptions for the approach to provide valid estimates. The first one is that units of observation on either side of the threshold are similar in all observed and unobserved characteristics, except for their treatment status. The second main assumption is that individuals cannot precisely manipulate the assignment variable in the proximity of the threshold. If both assumptions are met, the treatment is as good as randomly assigned.

Figure 2 shows the distribution of a hypothetical outcome variable around a treatment threshold. The variable continuously increases for higher values of the assignment variable, but it exhibits a discontinuity at the cut-off. In a regression discontinuity design, only observations close to this cut-off are evaluated, as illustrated by the shaded areas.

Figure 2: Hypothetical discontinuity

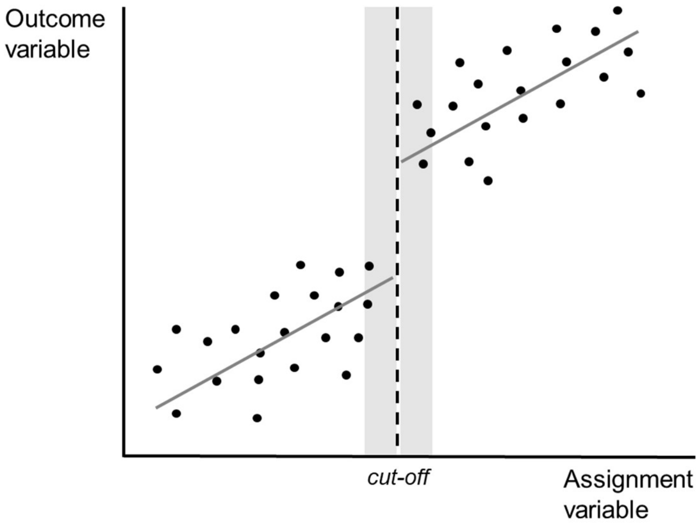

A basic regression discontinuity can be modeled by estimating an OLS regression of the following form:

$$y_i = a_1 + a_2 treated_i + a_3 X_i + e_i \quad (5)$$

where y is the outcome, $treated$ is a binary variable that indicates if the unit of observation i received the treatment, and X is the assignment variable. The treatment dummy takes the value 1 if the assignment variable exceeds the known threshold. When estimating Equation (5), the data are restricted to a certain bandwidth (i.e., observations just above and just below the threshold). Selecting the bandwidth is subject to a trade-off: A larger bandwidth increases the precision of the estimates, as more observations are used; a narrower bandwidth decreases the bias in estimating the treatment effect, because the units of observation are more similar in their other characteristics. The optimal bandwidth can be determined by using selection criteria, as discussed by Lee and Lemieux (2010). With this restriction, the coefficient a_2 captures the treatment effect.

Using historical data from the US, Gentzkow et al. (2015) investigate the effect of party control of state governments on the partisan composition of the press. They implement a regression discontinuity design to account for the reverse direction of causality, that the press influences voters and thus the chances that one or another party controls the government. Specifically, the authors use close elections, which allow for a comparison of situations in which parties receive a vote share just below the majority threshold with vote shares just above this threshold. The candidates and conditions are assumed to be similar in both situations, except that the parties won some elections by a close margin and lost others. Their results suggest that the number and content of Republican and Democratic newspapers in US states are usually not affected by who is in charge of the government.

Anderson and Magruder (2012) study the effects of consumer reviews published by Yelp.com on restaurant visits. The authors focus on the implications that arise from the forum's 1- to 5-star rating scale. Yelp.com calculates each restaurant's average rating but displays the value rounded to the nearest half-star. As a consequence, two restaurants with similar actual ratings (e.g., 3.24 and 3.26) would be shown to be of a different quality (e.g., 3-star vs. 3.5-star). Exploiting these discontinuities allows the authors the address endogeneity issues in the relationship between review quality and restaurant quality, which helps them to identify the positive effect that consumer reviews have.

Kaniel and Parham (2017) investigate the relationship between media attention to investment funds and consumer investment behavior. The authors compare funds mentioned in a Wall Street Journal "Category Kings" ranking with funds who just missed being mentioned in the ranking. Both types of funds are assumed to be similar in their performance and other characteristics, but the ones mentioned in the ranking are found to be more popular investment choices by consumers. Here, the regression discontinuity accounts for the possibility that funds are mentioned in the ranking only because they are already in great demand.

Compared to other quasi-experimental methods, the regression discontinuity approach requires less restrictive assumptions. For example, it is not necessary to assume that the treatment is generally as good as randomly assigned, nor is it required for the exclusion restriction to hold as in the instrumental variable approach. It is merely necessary that the units of observation cannot control their assignment to treatment near the cut-off, and that the discontinuity is not present in confounding variables. Despite these relatively weak assumptions, the regression discontinuity design can be interpreted and analyzed similarly to randomized controlled experiments. However, restricting the data to some bandwidth around the threshold involves a loss of observations. Thus it is crucial to have a large enough sample size, so that the parameters of interest can be accurately estimated. Another disadvantage is the limited external validity. The estimation results only pertain to the observations close to threshold, they could be different for other data points.

4.4 Field experiments

Field experiments resemble those conducted in the laboratory. They also involve a randomization of subjects and the experimenter has precise control over the desired intervention. However, with a field experiment, the treatment is evaluated in a natural environment rather than the laboratory. This difference is an important reason why these experiments have grown more popular in media-related research over the past decades. For instance, when testing for media effects in laboratory settings, it is difficult to account for people's tendency to selectively expose themselves to certain sources or items of media (e.g., partisan selective exposure). It is unclear how to interpret tests that do not account for this kind of self-selection; i.e., if subjects are confronted with information that they would rarely encounter in the real world.

Gerber, Karlan, and Bergan (2009) address this issue by conducting a field experiment before the 2005 Virginia gubernatorial election. They randomly assigned subjects to two treatment groups and a control group. Subjects in the treatment groups either received a free subscription to the Washington Post or the Washington Times. The authors do not find any effects on political attitudes or knowledge, perhaps because subjects could not be forced to read their subscribed newspaper.

King, Schneer, and White (2017) recruited 48 US media outlets and asked them to publish articles on selected but randomized topics and dates. The randomization allowed the authors to causally track the repercussions of the articles on the national news agenda and public debate. In this setting, the results are unlikely to be contaminated by recipients' selective exposure. Reading the articles was voluntary, including recipients' choice of which outlet(s) to use. The results of the experiment indicate a positive effect of the articles on the website pageviews of the outlets. In addition, the overall amount on Twitter posts on the topics addressed by the articles increased after the intervention.

Dertwinkel-Kalt, Kerkhof, and Münster (2018) conducted a field experiment shortly before the 2017 federal elections in Germany. The authors sent different versions of a letter to the editor to a large set of newspapers, randomizing the subject of the letter (i.e., Angela Merkel vs. the main challenger Martin Schulz) and the evaluation of the subject (negative vs. positive). The authors tested for different types of media bias by monitoring which version got published relatively more often. According to their results, German newspapers exhibited a form of incumbency bias, with the version about chancellor Merkel being printed more often.

Thus studying subjects in the field has the big advantage of avoiding the artificial environment that most laboratory settings involve, which may increase the external validity. Field experiments also grant more control over the randomization and intervention than the quasi-experimental methods discussed in Sections 4.1 to 4.3. The biggest disadvantage are the costs, as field experiments require enormous logistical efforts. In addition, it might be difficult or impossible to obtain access or permission to run an experiment in a certain environment. When conducting a controlled experiment outside the laboratory, there are usually greater concerns about contamination by third variables. These concerns can be addressed by recruiting a larger number of subjects than required in a laboratory setting, and by measuring and controlling for potential confounders. However, such actions usually produce additional costs and make the replication of a field experiment more difficult.

5. Outlook

Recent quantitative research on media and communication has shown a strong trend towards analyzing observational data – especially from online sources – with methods that support causal inference. This trend is a likely result of the so-called credibility revolution in empirical work (Angrist and Pischke, 2010). As a consequence, the chances of academic studies being used as an input in the decision-making process by companies, regulators, and policy makers have possibly increased. The role of web data in media-related research can be expected to continue to grow. The increasing availability of structured information from online sources and decreasing costs of processing these data will generally open new research opportunities. In addition, it will be much easier to conduct research in (less developed) countries that have been previously neglected because of the lack of traditional data. The continuing development of automated text processing and growing computational power will likely increase the use and usefulness of text-based measures. Automated content analysis of photo and video material has not been used much so far. Researchers may want to allocate resources in this direction, as these forms of communication likely have larger effects on recipients than text messages. Boxell (2018) provides an example of how machine-based content analysis can be applied in the field. The author uses over a million images from various US websites to construct an index of non-verbal media slant. This task involves an algorithm that identifies images containing a face of a politician, as well as facial recognition software to characterize different emotions, such as happiness, anger, disgust, and sadness. Analyzing such an amount of data is hardly possible with human content analysis. As researchers continue to develop and apply automated methods that can master complex coding tasks, open research gaps will be filled though.

References

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

Allcott, H., Gentzkow, M., & Yu, C. (2018). Trends in the Diffusion of Misinformation on Social Media. Working Paper.

Anderson, M., & Magruder, J. (2012). Learning from the Crowd: Regression Discontinuity Estimates of the Effects of an Online Review Database. *Economic Journal*, 122, 957–989.

Angrist, J. D. (1990). Lifetime Earnings and the Vietnam Era Draft Lottery: Evidence from Social Security Administrative Records. *American Economic Review*, 80, 313–336.

Angrist, J. D., & Krueger, A. B. (1991). Does Compulsory School Attendance Affect Schooling and Earnings? *Quarterly Journal of Economics*, 106, 976–1014.

Angrist, J. D., & Pischke, J. (2010). The Credibility Revolution in Empirical Economics: How Better Research Design is Taking the Con out of Econometrics. *Journal of Economic Perspectives*, 24, 3–30.

Baker, M. J., & George, L. M. (2010). The Role of Television in Household Debt: Evidence from the 1950's. *B.E. Journal of Economic Analysis & Policy*, 10, 1–36.

Barone, G., D'Acunto, F., & Narciso, G. (2015). Telecracy: Testing for Channels of Persuasion. *American Economic Journal: Economic Policy*, 7, 30–60.

Beattie, G. (2017). Advertising, Media Capture, and Public Opinion: The Case of Climate Change. Working Paper.

Beckmann, K. B., Dewenter, R., & Thomas, T. (2017). Can News Draw Blood? The Impact of Media Coverage on the Number and Severity of Terror Attacks. *Peace Economics, Peace Science and Public Policy*, 23, 1–16.

Berger, J., & Milkman, K. L. (2012). What Makes Online Content Viral? *Journal of Marketing Research*, 49, 192–205.

Bertrand, M., Duflo, E., & Mullainathan, S. (2004). How Much Should We Trust Differences-In-Differences Estimates? *Quarterly Journal of Economics*, 119, 249–275.

Boxell, L. (2018). Slanted Images: Measuring Nonverbal Media Bias. MPRA Working Paper 89047.

Budak, C., Goel, S., & Rao, J. M. (2016). Fair and Balanced? Quantifying Media Bias Through Crowdsourced Content Analysis. *Public Opinion Quarterly*, 80, 250–271.

Bursztyn, L., & Cantoni, D. (2016). A Tear in the Iron Curtain: The Impact of Western Television on Consumption Behavior. *Review of Economics and Statistics*, 98, 25–41.

Cagé, J. (2017). Media Competition, Information Provision and Political Participation: Evidence from French Local Newspapers and Elections, 1944–2014. CEPR Working Paper 12198.

Cagé, J., & Rueda, V. (2016). The Long-Term Effects of the Printing Press in sub-Saharan Africa. *American Economic Journal: Applied Economics*, 8, 69–99.

Campante, F. R., & Hojman, D. A. (2013). Media and Polarization: Evidence from the Introduction of Broadcast TV in the United States. *Journal of Public Economics*, 100, 79–92.

Dertwinkel-Kalt, M., Kerkhof, A., & Münster, J. (2018). Incumbency Dominance in Letters to the Editor: Field Experimental Evidence. *Political Communication*, forthcoming.

Durante, R., & Knight, B. (2012). Partisan Control, Media Bias, and Viewer Responses: Evidence from Berlusconi's Italy. *Journal of the European Economic Association*, 10, 451–481.

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

Eisensee, T., & Strömberg, D. (2007). News Droughts, News Floods, and U.S. Disaster Relief. *Quarterly Journal of Economics*, 122, 693–728.

Engelberg, J. E., & Parsons, C. A. (2011). The Causal Impact of Media in Financial Markets. *Journal of Finance*, 66, 67–97.

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

Falck, O., Gold, R., & Heblich, S. (2014). E-Lections: Voting Behavior and the Internet. *American Economic Review*, 104, 2238–2265.

Flaxman, S., Goel, S., & Rao, J. M. (2016). Filter Bubbles, Echo Chambers, and Online News Consumption. *Public Opinion Quarterly*, 80, 298–320.

Friebel, G., & Heinz, M. (2014). Media slant against foreign owners: Downsizing. *Journal of Public Economics*, 120, 97–106.

Garz, M., & Pagels, V. (2018). Cautionary Tales: Celebrities, the News Media, and Participation in Tax Amnesties. *Journal of Economic Behavior & Organization*, 155, 288–300.

Garz, M., & Sörensen, J. (2017). Politicians under Investigation: The News Media's Effect on the Likelihood of Resignation. *Journal of Public Economics*, 153, 82–91.

Garz, M., Sörensen, J., & Stone, D. F. (2018). Partisan Selective Engagement: Evidence from Facebook. SSRN Working Paper 3102249.

Garz, M., Sood, G., Stone, D. F., & Wallace, J. (2018). What Drives Demand for Media Slant? SSRN Working Paper 3009791.

Gatti, L., Guerini, M., & Turchi, M. (2016). SentiWords: Deriving a High Precision and High Coverage Lexicon for Sentiment Analysis. *IEEE Transactions on Affective Computing*, 7, 409–421.

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

Gentzkow, M., Petek, B., Shapiro, J. M., & Sinkinson, M. (2015). Do Newspapers Serve the State? Incumbent Party Influence on the US Press, 1869-1928. *Journal of the European Economic Association*, 13, 29–61.

Gentzkow, M., & Shapiro, J. M. (2010). What Drives Media Slant? Evidence from U.S. Daily Newspapers. *Econometrica*, 78, 35–71.

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

Gentzkow, M., Shapiro, J. M., & Sinkinson, M. (2011). The Effect of Newspaper Entry and Exit on Electoral Politics. *American Economic Review*, 101, 2980–3018.

Gentzkow, M., Shapiro, J. M., & Sinkinson, M. (2014). Competition and Ideological Diversity: Historical Evidence from US Newspapers. *American Economic Review*, 104, 3073–3114.

Gerber, A. S., Karlan, D., & Bergan, D. (2009). Does the Media Matter? A Field Experiment Measuring the Effect of Newspapers on Voting Behavior and Political Opinions. *American Economic Journal: Applied Economics*, 1, 35–52.

Groseclose, T., & Milyo, J. (2005). A Measure of Media Bias. *Quarterly Journal of Economics*, 120, 1191–1237.

Gurun, U. G., & Butler, A. W. (2012). Don't Believe the Hype: Local Media Slant, Local Advertising, and Firm Value. *Journal of Finance*, 67, 561–598.

Halberstam, Y., & Knight, B. (2016). Homophily, Group Size, and the Diffusion of Political Information in Social Networks: Evidence from Twitter. *Journal of Public Economics*, 143, 73–88.

Jetter, M. (2017). The Effect of Media Attention on Terrorism. *Journal of Public Economics*, 153, 32–48.

Jo, D. (2017). Better the Devil You Know: An Online Field Experiment on News Consumption. Working Paper.

Kaniel, R., & Parham, R. (2017). WSJ Category Kings: The Impact of Media Attention on Consumer and Mutual Fund Investment Decisions. *Journal of Financial Economics*, 123, 337–356.

King, G., Schneer, B., & White, A. (2017). How the News Media Activate Public Expression and Influence National Agendas. *Science*, 358, 776–780.

Latham, O. (2015). Lame Ducks and the Media. *Economic Journal*, 125, 1918–1951.

Lee, D., Hosanagar, K., & Nair, H. S. (2017). Advertising Content and Consumer Engagement on Social Media: Evidence from Facebook. *Management Science*, 64, 5105–5131.

Lee, D. S., & Lemieux, T. (2010). Regression Discontinuity Designs in Economics. *Journal of Economic Literature*, 48, 281–355.

Lelkes, Y., Sood, G., & Iyengar, S. (2017). The Hostile Audience: The Effect of Access to Broadband Internet on Partisan Affect. *American Journal of Political Science*, 61, 5–20.

Loughran, T., & McDonald, B. (2011). When Is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10-Ks. *Journal of Finance*, 66, 35–65.

Martin, G. J., & McCrain, J. (2018). Local News and National Politics. Working Paper.

Martin, G. J., & Yurukoglu, A. (2017). Bias in Cable News: Persuasion and Polarization. *American Economic Review*, 107, 2565–2599.

Miner, L. (2015). The Unintended Consequences of Internet Diffusion: Evidence from Malaysia. *Journal of Public Economics*, 132, 66–78.

Müller, K., & Schwarz, C. (2017). Fanning the Flames of Hate: Social Media and Hate Crime. Working Paper.

Peukert, C., Claussen, J., & Kretschmer, T. (2017). Piracy and Box Office Movie Revenues: Evidence from Megaupload. *International Journal of Industrial Organization*, 52, 188–215.

Puglisi, R., & Snyder Jr., J. M. (2015). The Balanced US Press. *Journal of the European Economic Association*, 13, 240–264.

Qin, B., Strömberg, D., & Wu, Y. (2017). Why Does China Allow Freer Social Media? Protests versus Surveillance and Propaganda. *Journal of Economic Perspectives*, 31, 117–140.

Qin, B., Strömberg, D., & Wu, Y. (2018). Media Bias in China. *American Economic Review*, 108, 2442–2476.

Sen, A., & Yildirim, P. (2016). Clicks Bias in Editorial Decisions: How does Popularity Shape Online News Coverage? SSRN Working Paper 2619440.

Simonov, A., & Rao, J. M. (2018). What Drives Demand for Government-Controlled News in Russia? Working Paper.

Snyder Jr., J. M., & Strömberg, D. (2010). Press Coverage and Political Accountability. *Journal of Political Economy*, 118, 355–408.

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

Szeidl, A., & Szucs, F. (2016). Media Capture through Favor Exchange. Working Paper.

Tausch, F., & Zumbuehl, M. (2018). Stability of Risk Attitudes and Media Coverage of Economic News. *Journal of Economic Behavior & Organization*, 150, 295–310.

Thistlethwaite, D. L., & Campbell, D. T. (1960). Regression-Discontinuity Analysis: An Alternative to the Ex Post Facto Experiment. *Journal of Educational Psychology*, 51, 309–317.

Ulbricht, D., Kholodilin, K. A., & Thomas, T. (2017). Do Media Data Help to Predict German Industrial Production? *Journal of Forecasting*, 36, 483–496.

Webbink, D., van Erp, J., & van Gastel, F. (2017). The Effect of Media Exposure of Suspects on Solving Crime. *Economic Journal*, 127, 547–570.

Wright, P. G. (1928). *The Tariff on Animal and Vegetable Oils*. Macmillan, New York.