THE UNIVERSITY *of* EDINBURGH

Essays on Media Reportage and Economic Behaviour

Jonathan Spiteri

Doctor of Philosophy

School of Economics

University of Edinburgh

2015

Declaration

I declare that this thesis was composed by myself and that the work contained therein is my own, except where explicitly stated otherwise in the text.

(*Jonathan Spiteri*)

Acknowledgements

During the course of this epic four-year journey I have benefitted from the advice, knowledge and company of several people. Firstly, I would like to thank my supervisor Kohei Kawamura for his tireless support and help; it has been an immense pleasure to work with and learn from him. I also thank José V. Rodrìguez-Mora, Michèle Belot, Andrew Clausen, Steven Dieterle, Ed Hopkins, Philipp Kircher, Tatiana Kornienko, József Sákovics and Tim Worrall for all their assistance and feedback at various points throughout this process, and for being so generous with their time. My gratitude also goes to my examiners Andreas Steinhauer and Johannes Münster for making the whole viva experience a great deal more pleasant and fruitful than I could ever have hoped for.

I am highly indebted to my fellow PhD students and office mates for making this such an enjoyable experience. I particularly would like to acknowledge the efforts of Athanasia Arnokourou, Rachel Forshaw, Gerdis Marquardt and Sarah Schröder in helping me to perfect the art of productive procrastination.

Special thanks must also go to my contingent of Maltese friends in Scotland, namely Nicholas Micallef, Renzo Cordina, Chris Mangani, Daniel Borg, Michael Camilleri, Laura Cioffi, Clara Galea and Christine Barbara, for helping me explore the city's eclectic selection of watering holes (as well as the fantastic house parties). I also thank Glenn Fenech, Brian Camilleri and Ramon Bondin for never failing to remind me how warm the weather is back home.

Finally, I am eternally grateful to the people closest to my heart - my parents, my brother Nicky, and Paula Aquilina for their unconditional support and encouragement.

Abstract

This thesis looks at the economics of mass media from a variety of perspectives. The main aim is to analyse the key factors that influence media reporting behaviour, and in turn the impact of reportage on individual decision-making processes.

The first chapter provides a brief summary of the contextual background of this thesis, by presenting the main points tackled in the subsequent chapters as well as a concise overview of the main contributions across various fields of study.

The second chapter explores the relationship between advertisers and the media using a simple model of horizontal and vertical product differentiation in a duopolistic setting. In this framework, when a news story is published one firm will benefit in terms of higher consumer demand and profits, while the other will suffer. Firms can influence the media's decision to publish the news story or withhold it via advertising expenditure. The main result shows that in equilibrium when news signals conform to people's prior beliefs, extreme or strong stories will be withheld from publication by the media. This is because strong stories will result in a drastic decline in profits for one firm, thus providing it with an incentive to switch over and change its production process to mimic the other (beneficiary) firm, thereby eliminating vertical product differentiation. Therefore, the beneficiary firm would have an incentive to ensure that the news story is withheld to prevent this increase in competition and the subsequent erosion of its profit margins. The results provide an alternative rationale to explain recent evidence on under-reporting by the U.S. media in relation to various issues like climate change and the nutritional content of food.

The third chapter looks at the responsiveness of individual private behaviour to media coverage of a particular news story. Survey data on charitable gift-giving in the U.S. are used in order to analyse the impact of newspaper coverage of the 2004 Indian Ocean tsunami on both the likelihood and magnitude of monetary disbursements towards the relief effort. The identification strategy employed in this paper exploits differences in county-level growth rates of violent crime in order to account for the variation in newspaper coverage of the tsunami, thus circumventing potential endogeneity problems. The results show that media coverage only had a modest effect on people's decision to donate or not, but conversely had a significant and non-trivial impact on the amount of money donated. Furthermore, this impact was larger for young adults within the 25-34 age bracket and individuals who had undertaken some form of voluntary work in the previous year. These results hold even after the implementation of various robustness tests,

and serve to highlight the growing influence of the media on people's behaviour.

The final chapter analyses the impact of media reports on electoral outcomes, and in particular the extent to which soft or sensationalist news reportage influences voting. Survey data on individual voting behaviour during the 2000 U.S. Presidential election is used, together with a novel dataset on the amount of coverage afforded to the Monica Lewinsky scandal over the period January 17, 1998 to August 31, 2000. We first show that Lewinsky coverage was not driven by the newspapers' political bias, but rather by other factors including tabloid journalism. This independence enables us to focus solely on the impact of media reports on voting, in contrast to the rest of the literature which deals with the electoral influence of politically-biased media outlets. We then look at how newspaper coverage of the Lewinsky scandal influenced voting patterns in the 2000 U.S. Presidential election. To account for potential endogeneity issues we use county-level variation in the number of deaths caused by extreme weather events as an instrument for Lewinsky articles. We find that media coverage of the scandal had a positive and statistically significant impact on the likelihood of voting for George W. Bush, and conversely a negative influence on the probability of voting for Al Gore: this pattern is visible among both Democrats and Republicans. The results are robust to various tests, and raise several questions regarding the media's role within the democratic process.

Contents

Chapter 1

Preface

The principal driving force behind this thesis is the simple fact that the media, in all its forms and guises, is everywhere in today's world. We are constantly being bombarded with information, be it television programmes, newspapers, online portals and social media, all available at the touch of a button or the swipe of a screen, 24 hours a day. News consumption across all forms of media has never been higher - this is seen in Figure 1.1 (based on a Pew Research Center 2012a survey in the U.S.) which reports that the average number of minutes dedicated to news consumption per day has increased markedly over the last 15 years across most age groups. Given this ubiquity, a couple of crucial questions immediately come to mind; this thesis aims to tackle both of them, given that they are intimately related with the first one having a significant impact on the other:

1. What type of news content is the media publishing, and what influences its reporting practices?

2. What influence (if any) does such reportage have on people's behaviour?

1.1 News Content and Media Reportage

The first question is related to media reportage and content, and is often framed within the context of biased news. A proper understanding of what influences media reporting behaviour is crucial since the media has traditionally been seen as a gatekeeper in society (Williams and Carpini, 2000), chanelling crucial information to the masses in order to keep politicians in check and ensure that policies are in line with people's expectations. In other words, the information published as news by the media plays a critical role in the proper functioning of the democratic process.

A crucial assumption underpinning the above discussion is that the media is independent from any external forces or interest groups which may somehow influence the level as well as tone of news content. In reality, there has been much

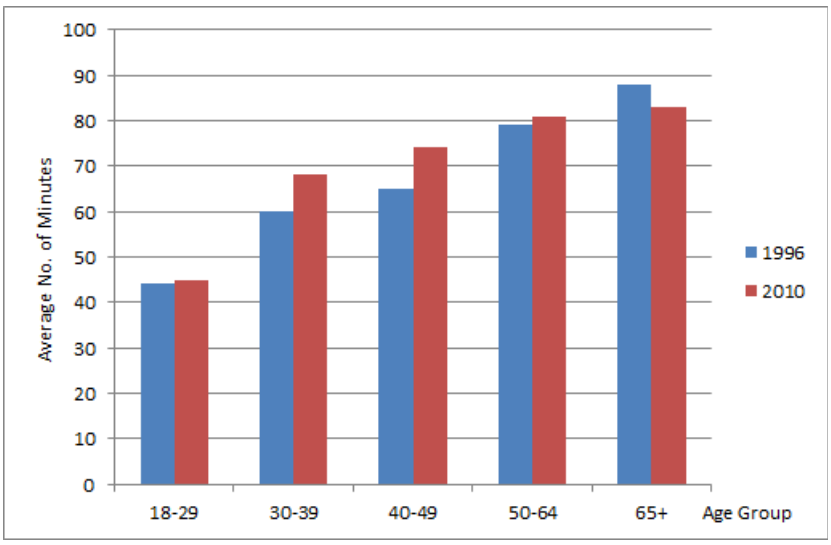

Figure 1.1: Minutes Spent on News Consumption per day in the U.S. (Source: Pew, 2012a)

debate about the media's independence in recent decades, which has called into question its gatekeeper role. In a recent survey by the Pew Research Center (2013) on people's perceptions regarding the media in the U.S., 76% of respondents said that news organisations tend to be biased towards one political party or another, while 75% maintained that media reportage is influenced by "powerful people and organisations".

Most of this work is related to the literature on media agenda-setting behaviour, which promulgates the view that the media's reporting practices help to shape public discourse and hence the importance that people attribute to particular topics or issues. For example, Bennett, Lawrence, and Livingston (2008) state that the U.S. media played a key part in helping to raise public awareness regarding the devastation caused by Hurricane Katrina in New Orleans in 2005, and in particular the shortcomings of the government's response efforts.

Therefore, agenda-setting behaviour can be thought of as a natural consequence of the media's work. An independent, free press would tailor its reporting practices according to public interest (although this notion is somewhat vague); in the absence of independence then media reportage would cater to the interests of other external forces (Iyengar and Kinder, 2010). This is where the idea of media bias comes into play, a concept that in recent years has been the subject of fierce debate among academics and observers from across the world. In the next few paragraphs, I shall attempt to organise these arguments in a logically consistent manner, since in reality the debate on media bias is deceptively diverse.

The first and most prominent type of media bias is political or partisan slant, whereby a media outlet either covers news stories that are favourable to one side of the political spectrum (or stories that harm the other), or else wilfully slants its coverage to the benefit of one party. Over the years, several authors have attempted to determine whether the news media favours one political ideology or another, with typically varying results. For example, Groseclose and Milyo (2005)

report that the U.S. print media has a clear liberal/left-wing bias in terms of its news reportage, a result contradicted by Niven (2003) who finds no systematic bias in either direction in the U.S. media.

Various factors have been proposed to explain the incidence of political or partisan bias in the news media. The most commonly cited argument is that the media tailors its reportage to suit the opinions or political beliefs of its readers in order to boost circulation figures. This demand side factor, which has been used by the likes of Gentzkow and Shapiro (2006) and Mullainathan and Shleifer (2005), is the result of a well-known fact in psychology that people tend to prefer reading information that conforms to their prior beliefs or expectations, rather than contradictory news (Lord, Ross, and Lepper, 1979). Therefore, under this scenario partisan slant occurs due to readers' desire for confirmatory information, rather than any external or outside forces directing the media's reportage.

A number of other determinants of political media bias have also been proposed, which can all broadly be classified as supply side factors. These include the political leanings of the newspaper editor, the beliefs of the owners, as well as the influence of politically driven advertisers. The common thread with these supply side factors is that they all involve some sort of deliberate attempt at manipulating the news to suit the requirements of a particular interest group, as opposed to simply tailoring reportage to audience preferences in order to boost market shares. The evidence of such bias is mixed, in part reflecting the difficulty involved in determining both the existence and source of biased news. For example, Gentzkow and Shapiro (2010) find little to no evidence that political bias in news reportage is being driven by ownership beliefs in the U.S. print media, instead contending that this is mainly the result of demand driven concerns. Conversely, Tella and Franceschelli (2011) find a negative correlation between the level of monthly government advertising in the four major national newspapers in Argentina and front-page coverage of government corruption scandals.

The second type of media bias is advertiser driven or commercial media bias, whereby media reportage is influenced by advertiser interests. In essence this means that the media would, at the behest of its advertisers, avoid covering news stories which would somehow be detrimental to the advertiser, and instead ramp up coverage of stories that paint the advertiser in a positive light. For example, a media outlet would avoid covering news of any oil spills caused by one of its oil producing advertisers, and instead focus on more positive elements like technological innovations or philanthropic initiatives.

An interesting observation to be made is that commercial media bias can be thought of as a by-product of the media's political independence. Petrova (2011) relates how in the nineteenth century, most newspapers in the U.S. were affiliated with one political party or another due to the fact that a large proportion of their revenue was coming from political sponsorships. It was only with the advent of advertising from private firms, which started to increase towards the latter part of the century, that newspapers emerged as politically independent and thus capable of fulfilling their gatekeeper role in society without any interferences. In fact, Petrova (2011) shows that places with higher local advertising rates were

more likely to have newspapers that were independent from political parties.

Therefore, in many respects advertising has been crucial in allowing media outlets to acquire independence from political interference. However, many (e.g. Chomsky and Herman, 1988) contend that although the media may technically be free from political influence, in reality the rise of advertising has only led to a shift in influence from the government to the boardroom, with private interests dictating public discourse via the media.

The existence or otherwise of commercial media bias is open to debate, once again reflecting the general difficulty involved in analysing such phenomena empirically, as well as more general arguments that question its validity. Sutter (2002) states that advertising cannot significantly affect media reporting behaviour since the promotion of pro-industry information (or conversely the suppression of anti-industry news) is a public good that requires large-scale coordination across firms in order to succeed. Otherwise, the incentive to free ride would prevent such activities from occurring. In addition, news stories that are beneficial to a particular firm are likely to have a negative impact on another firm, which in turn may lead the latter firm to promote its own agenda, with the net result leaving the original status quo unchanged.

Nonetheless, there is some evidence supporting the existence of commercial media bias in certain situations. For example Fletcher (2003) finds that editorial positions in the leading North American medical journals are partially influenced by the requirements of their advertisers, which may lead to conflicting preferences since many of the leading North American journals within this field only accept advertisements for medical products in their publications.

The third and final (in this context) type of media bias is sensationalism or the "dumbing-down" of news content. This is related to the idea that over the last few decades there has been a marked shift in media content away from factual or 'hard' news in favour of more soft or vacuous content. Thus for example the media would, under this idea, increasingly focus its coverage on entertainment or personal interest news stories (like sex scandals or gossip) at the expense of political news or international affairs.

Various arguments have been cited in order to support this dumbing-down hypothesis of news content. Recall Figure 1.1, where it was shown that news consumption across all media has increased over time. This should technically bode well in terms of how well-informed the public is with regards to current affairs, which in turn should lead to improved electoral outcomes. However, Figure 1.2 contradicts this assertion, showing that knowledge of public affairs in the U.S. (based on a Pew Research Center (2007) survey) is actually shrinking. This supports the claim that although the news is more accessible and readily available than ever before, the quality of the news content is on the decline due to this increased focus on soft news at the expense of hard news.

In addition, Hamilton (2004) finds a distinct shift from hard political news to entertainment stories in U.S. network evening news programmes over the period 1969 to 1998. The author cites a number of reasons to explain this change in content, including deregulation of media markets in the U.S., higher concentration

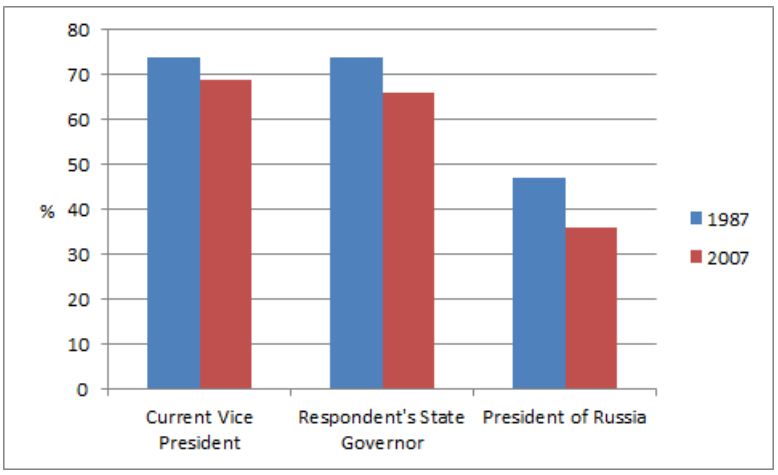

Figure 1.2: Knowledge of Current Affairs in the U.S. (Source: Pew, 2007) [1]

of media ownership and advertising. In fact, advertising has garnered increased attention within the scope of dumbed-down news content. Bagdikian (2000) argues that the media's main tradeable service is not the provision of news but rather the provision of potential consumers for its advertisers. Therefore, the media will avoid any news that may somehow antagonise its advertisers and instead focus on uncontroversial content, i.e. soft news. Hence, the dumbing-down of news content can be considered as a subtle subset of commercial media bias.

To summarise, the concept of biased media reportage has generated a vibrant debate among academics and the general public alike, and is set to continue in the future due to the media's ever-growing presence. Recent advances in automated content analysis and computational linguistics, such as those described in Hansen et al. (2014), coupled with improvements in online news archiving, should assist in the derivation of more reliable media bias measures (across all forms of bias). Although plenty of attention has been afforded to the media's supposed partisan/political slant, there are still several aspects of media bias that are ripe for further research. In particular, much work still needs to be done in terms of understanding the complex relationship between the media and its diverse array of advertisers, often with conflicting publication preferences, both from a theoretical and empirical perspective. In addition, a more thorough treatment is required of the underlying causes and existence of dumbed-down news content across all media forms.

1.2 Media Influence and Individual Behaviour

So far we have dealt with the media's reporting practices and the main determinants of media bias. I now turn to the second question listed above, namely the impact of media reportage on individual behaviour. The idea is that although

[1] In each case, survey participants were specifically asked to name the current Vice-President of the U.S., their state governor and the current Russian President.

it is important to assess the reporting behaviour of the media, such an exercise would be rendered almost futile unless such reportage has any visible effect on people's actions or decisions.

The study of media influence, both from a theoretical and empirical perspective, has a long history across several academic fields. Perhaps the first proper studies within this context are those by Lazarsfeld et al. (1944) and Berelson et al. (1954) who look at the effects of media communications on the 1940 and 1948 U.S. Presidential elections. In fact, a large proportion of studies dealing with media influence invariably focus on its effects on voting and electoral outcomes. Authors like Bartels (1993) and Shields, Goidel, and Tadlock (1995) have all found that media reportage plays a key role in shaping public opinion regarding electoral candidates and has a significant impact on final voting decisions.

Over the years, other authors have looked beyond the confines of electoral outcomes to analyse the effect of media reportage on other types of individual behaviour. This has created a vast literature on a wide array of topics including health, violence and government policy. Within the sphere of public health, Agliata and Tantleff-Dunn (2004) find that men who were exposed to media messages or advertisements espousing ideal body images developed higher levels of muscular dissatisfaction and became significantly more depressed, relative to a control group who were only exposed to neutral images. Similarly, Villani (2001) finds that children and adolescents exposed to violence, alcohol and tobacco abuse, sexual activity and other so-called "risky behaviour" through various forms of mass media communications are more likely to engage in such behaviour themselves.

At this point one might think that the media is the root of all evil in today's world due to the rather depressing portrayal of media influence depicted thus far. However, it is important to note that a substantial portion of this literature has also been dedicated to demonstrating the ability of the media to act as an agent of positive change within society. For example Jaramillo (2001) reports that mass media communications related to tuberculosis (TB) awareness and education in Cali, Colombia resulted in a significant increase in the number of TB smear tests undertaken as well as a 52% increase in the number of new TB cases diagnosed, although these benefits were relatively short-lived. Similarly, Stromberg (2004) shows that increased access to to the media leads to higher government responsiveness and accountability, providing evidence that during the 1930s the U.S. government allocated higher levels of relief funds as part of the Depression-era New Deal programme to those counties with greater access to radio.

Although there is substantial evidence showing that media reportage does have a tangible and significant effect on behaviour, it is worth noting that this is by no means the general consensus. Indeed, several authors have sought to downplay or even dismiss the idea that the media can have any impact on people's decisions or opinions. One of the first studies in this vein was by Trenaman and McQuail (1961) who looked at the influence of television broadcasts on UK voters in the 1959 General Election, finding no evidence of any change in voter attitudes towards each political party, although it did increase their general knowledge of

party politics. In addition, Mermin (1997) finds that media coverage of the 1992 civil war in Somalia did not mobilise U.S. government intervention, contrary to popular opinion. Rather, it was politicians in Washington who set the agenda for foreign policy, which in turn resulted in greater media coverage of the crisis.

This discussion raises several interesting points regarding the exact nature of media influence and the varying forms that it can assume. Firstly, and perhaps most obviously, does the media have the same impact on all people across all situations? This question is critical since it seems unreasonable to assume that the media would have a homogenous, blanket influence regardless of people's socio-economic background or the issue being reported. In fact, during the height of public concern regarding media influence, largely driven by the expansion of television across the globe, Klapper (1960) commented on how the media must not be considered in isolation, but rather within its proper social context and in particular the way in which it interacts with socio-economic forces and institutional characteristics.

Therefore, it is important to look at the extent to which media influence varies across individuals and issues. Iyengar and Kinder (2010) state that traditionally the bulk of this literature has focussed on heterogeneity of influence on the basis of education, whereby it is asserted that people with a higher level of educational attainment are less susceptible to media influence due to their increased cognitive abilities and skills (collectively referred to as cognitive sophistication). However, the authors cite various alternative theories that have been proposed over the years, most of which are related to differences in media consumption patterns.

One such theory is by MacKuen and Combs (1981), whose so-called "attentiveness model" states that people who are highly motivated or engaged by a particular topic will actively seek out such stories in the media, and will therefore be more susceptible to their influence. Thus for example, an individual who is concerned by matters related to obesity or fitness is more likely to read media stories on nutrition or body image, which in turn may affect his/her behaviour. Another theory is by Erbring, Goldenburg, and Miller (1980) whose "issue sensitivity" hypothesis explains that people are more likely to be affected by stories which deal with issues that they consider (both consciously and subconsciously) to be salient or relevant to their present situation. For example, according to this hypothesis unemployed people's perception of the current government may be significantly influenced by media reports on any cuts or increases in welfare benefits.

Hence, it is entirely plausible that the extent of media influence depends on a whole host of factors related to both the individual and the nature of the news story. A second important question to consider is which media are more influential in terms of altering people's behaviour. This question has cropped up consistently throughout the last century with the introduction of new media like radio and television, and is particularly relevant in the current context given the rise of online news portals and social media which have to some degree revolutionised the way in which the news is consumed. In a global survey of media consumption patterns and the relative importance that people attach to each

medium, McRoberts and Terhanian (2008) find that the internet is almost twice as influential (in terms of daily time spent on each medium as well as influence over purchases) as television and approximately ten times more influential than newspapers or magazines. Needless to say, these results have important implications for future research within this field, and highlight the potential sensitivity of any analysis on media influence to the type of medium under consideration.

One final question to consider is related to the impact of biased media reportage (as described above) on people's behaviour. So far I (and indeed, most of the literature) have avoided commenting on the actual content of media reports when discussing their influence on people. However, in light of the extensive discussion presented above on media bias, this question assumes significant importance since it deals with people's susceptibility to propaganda (Chomsky and Herman, 1988) and their ability to carefully consider and filter the information being presented to them by the media.

In this regard, the literature on the extent to which media bias influences behaviour is somewhat limited. Certainly, there is a clear paucity of studies on people's susceptibility to commercial media bias as well as soft news reportage. This reflects both the difficulty involved in actually gauging the existence or otherwise of these two types of biases, as well as identifying the exact causal relationship between biased coverage and behaviour in such cases. Rather, this literature is mainly focussed on whether politically biased news reportage has any impact on people's voting decisions.

The evidence regarding the electoral implications of partisan media bias is mixed. On the one hand, Druckman and Parkin (2005) find that the editorial slant of newspaper electoral campaign coverage has a significant impact on voting patterns. Similarly, Martin and Yurukoglu (2014) report that watching four additional minutes of Fox News per week increases the probability of voting Republican by almost one percentage point, echoing the results obtained earlier by DellaVigna and Kaplan (2007).

On the other hand, a number of authors have found no evidence of changes in people's voting induced by politically slanted news reportage. For example Gerber, Karlan, and Bergan (2009) conduct a field experiment in Virginia to assess the impact of exposure to politically biased newspapers (i.e. the Washington Post versus the Washington Times) on voting in the 2005 gubernatorial election, finding no significant change in voting patterns across the two treatments. This fits in with other empirical findings related to the so-called "filter hypothesis" (Schmitt-Beck, 2003) which postulates that people's interactions and conversations with others (family, friends, colleagues, etc.) are crucial to their general acceptance and interpretation of political information derived from the media, which could potentially help to filter out the effects of biased news.

The exact nature of the media's influence over individual behaviour is still up for debate, providing significant scope for further work within this field. Perhaps the most glaring gap in the literature is related to the impact of news reportage on economic choices, which has for some reason been largely neglected. In addition, more work still needs to be done in order to determine people's susceptibility to

biased news coverage, across all forms of media bias as defined earlier, particularly with regards to more recent phenomena like commercial media bias and soft news reportage.

1.3 Overview of Thesis

Chapter 2 deals with the material covered in Section 1.1 above, by looking at how advertisers may influence media reporting behaviour. The focus is specifically on the advertiser-media relationship since advertising revenue accounts for the majority of media organisations' annual income, thereby indirectly bestowing significant clout and influence on these firms. In this chapter, I develop a simple model whereby a monopolist media firm receives a noisy signal regarding the true state of the world, and must decide whether to publish the news or withhold it. We then have two horizontally and vertically-differentiated advertisers who have conflicting incentives with regards to the publication or not of the news signal - one firm benefits, while the other would be worse off.

This represents a clear departure from the literature on commercial media bias, and mirrors the real-world situation where each media firm has a plethora of advertisers with differing preferences regarding which stories to publish and which to withhold. Similarly, consumers are heterogeneous in terms of their horizontal preferences for each advertiser's goods, but share a common prior belief regarding the true state of the world, whose realisation will in turn influence the relative utility they derive from either firm's product offering. Each firm may influence the media's reporting decision depending on the size of the advertising fee offered. The main result is that extreme or strong news signals that conform to consumers' prior beliefs will be withheld from publication by the media, while only moderate stories will be published.

This is because the advertiser who, at face value, would benefit from publication actually has an incentive to conceal the news story since the strength of the signal would be sufficient to induce its competitor to switch over and eliminate vertical product differentiation. In other words, strong news signals lead to more direct competition for the ex-ante beneficiary firm, which leads to lower profits than in the non-publication scenario. Thus, our simple model predicts the possible existence of advertiser-driven media bias, which in turn may help to explain several documented cases of biased reportage in the news, including newspaper coverage of climate change issues.

Having looked at one of the key determinants of media reporting practices, we then proceed to analyse the influence of reportage on people's choices. In chapter 3 we tackle one particular aspect of individual behaviour by looking at the impact of media coverage on individual economic choices. More specifically, the focus is on U.S. newspaper reports on the devastating Indian Ocean tsunami which hit Southeast Asia on December 26, 2004, and the influence of such coverage on both the likelihood and size of charitable donations by private U.S. citizens towards the subsequent relief effort.

One of the main issues when dealing with this sort of analysis in the literature is the inevitable endogeneity issues which are likely to be present and which limit the scope for any causal relationships to be derived from these results. Therefore, to circumvent such issues I use an instrumental variables approach (IV), where I instrument for tsunami coverage using county-level growth rates in average violent crime. The results, which are robust to various additional tests and specifications, show that newspaper coverage of the tsunami had a modest impact on the likelihood of donating money, but conversely had a non-trivial influence on the amount of money donated.

In fact according to the results one additional tsunami-related news story per week published by the average American newspaper, over the 6-month period immediately after the disaster, would have boosted average donations by around 7%, translating to approximately US$131 million. In addition, further analysis shows that this impact was larger for young adults within the 25-34 age bracket, consistent with the marketing literature on the responsiveness of different demographics to advertising messages. The results therefore serve to highlight the growing influence of the media over people's economic choices.

Finally, chapter 4 looks at the political economy aspect of media influence by analysing the extent to which soft or sensationalistic news coverage affects voting. The focus is on soft news since this refers to entertainment or personal interest stories that provide limited information regarding the suitability of electoral candidates or policy efficacy, meaning that it is unclear whether media coverage of such stories should have any impact on voting. To this end, I look at U.S. newspaper coverage of the infamous Monica Lewinsky scandal and individual voting patterns during the 2000 U.S. Presidential election between Al Gore and George W. Bush.

Firstly, I take a closer look at what influenced the level of Lewinsky coverage across newspapers, concluding that this is uncorrelated with the individual newspaper's political slant or indeed the beliefs of their target audience. This helps to eliminate the possibility that any relationship between coverage and voting is merely a reflection of the newspapers' political beliefs, instead enabling us to focus solely on the impact of (soft) media reportage. I then move on to the main part of the analysis, namely the impact of Lewinsky coverage on the outcome of the 2000 election. Once again, I circumvent any potential endogeneity issues by employing an instrumental variables approach, using cross-county variation in the number of deaths caused by extreme weather events as an instrument for Lewinsky coverage.

The results show that Lewinsky coverage had a positive and significant impact on the likelihood of voting for Bush, and a negative effect on the likelihood of a Gore vote. More importantly, this pattern was consistent across both Democrats and Republicans, as well as those who had voted for Bill Clinton in the previous election. In fact, a 5% increase in newspaper coverage of the Lewinsky scandal could have increased the likelihood that a non-Republican voted for Bush in the 2000 Presidential election by around 1.9%, with a similar decline reported in the probability of voting for Gore. Although small, these magnitudes were sufficient

to swing the election given the tightness of the actual results, and raise several questions regarding the media's role in the democratic process as a conduit of information.

Therefore, this thesis incorporates ideas from several fields within the broader economics literature, like industrial organisation, public economics and political economy, and applies them to the economics of media reportage and its influence over individual behaviour. The study of media economics has grown substantially in recent years, and has spawned a diverse and intellectually stimulating literature. I hope that this thesis continues in this tradition and serves as a modest contribution to this burgeoning field.

Chapter 2

When is No News Good News? A Model of Information Disclosure and Commercial Media Bias

2.1 Introduction

The media's reporting behaviour has come under intense scrutiny in recent years, in line with its growing presence within modern society. Increasingly, the focus has been on the potential influence of advertisers on media reportage (Chomsky and Herman, 1988), with several cases used to highlight these concerns, from the impact of the pharmaceuticals industry on editorial positions in medical journals (Fletcher, 2003) to the alleged influence of automotive advertising on the under-reporting of climate change issues in the mainstream media (Oreskes and Conway, 2010).

The mounting evidence of advertiser-driven media bias seems somewhat odd given that the media relies on a wide variety of advertisers, who in turn often have conflicting interests in terms of which stories to publish or suppress. For example although traditional fossil fuel vehicle manufacturers may have an incentive to conceal or distort evidence of man-made climate change, other manufacturers that produce hybrid or alternative fuel vehicles would actually benefit from full disclosure. One may argue that the influence of certain interest groups may be larger due to their relative size and clout; however this argument fails to account for the persistence in biased news coverage documented in various cases. For example Toyota, the leading hybrid car brand in the world, is one of the top ten advertisers in the U.S. with an annual expenditure of well over US$860 million (Kantar, 2014). Therefore, the continued under-reporting of several news stories like climate change is surprising given the existence of key industry players who would, at least at face value, benefit from their publication.

This paper contributes to the ever-growing literature on commercial media

bias by considering how the media's decision to publish a news report is influenced by the nature of the story as well as the nature of its advertisers. More specifically, this paper recognises the fact that the media typically has several advertising relationships with various firms, and that a news story may thus have contrasting implications for different advertisers. Using a simple model of horizontal and vertical product differentiation, we look at a duopolistic situation where each firm produces a similar yet differentiated good, and where consumer tastes are heterogeneous.

The novelty of this paper is that consumers also share a common belief regarding a particular state of the world (for example the likelihood that global warming is being caused by human activities). The realisation of this state has an impact on consumer utility and hence the choice over which product to consume. We assume that the media receives a noisy signal regarding the state of the world, with varying degrees of intensity and strength, and that its decision to publish this signal is dependent on the advertising fee offered by either firm, given that both firms have divergent interests in terms of the publication or concealment of the story since one firm would nominally benefit from its publication while the other would not.

The main result is that for extreme or strong news signals that conform to people's prior beliefs (indexed by specific parameter values in the model), the media would always have a strictly positive incentive to withhold publication of the news story. The reason for this outcome is that the firm who benefits most from the publication of the story actually has an incentive to *conceal* strong signals since such stories would induce the other firm to change its product offering (i.e. remove vertical product differentiation) and compete directly with it (provided that the fixed costs associated with switching are sufficiently low). Hence as a result strong news that would be highly beneficial to one firm would be withheld, since (paradoxically) it stands to lose from publication due to increased competition. The interaction between consumer priors and the news signal received by the media (both in terms of direction and strength) plays a key role in our model, since different permutations of priors and signals alter the relative publication incentives of each firm substantially as well as the size of their advertising fees.

The model's results must be seen in light of the vital importance of advertising for the continued survival of private media firms, precipitated by the rise in online news media and free-to-air TV news channels. This is seen in Figure 2.1, which shows the composition of newspaper revenues in the U.S. over the period 1990 to 2009; as seen in the chart, advertising accounts for an average of 75-80% of newspapers' total revenue.

The model's results shed light on a new kind of commercial media bias, one that emerges due to competitive concerns even though the news may (at face value) seemingly favour the advertiser in question. The results underscore the complexity of the advertiser-media relationship and raise several questions regarding the potential effectiveness of any regulations in this regard. The model also suggests that increased media competition and decentralisation of ownership

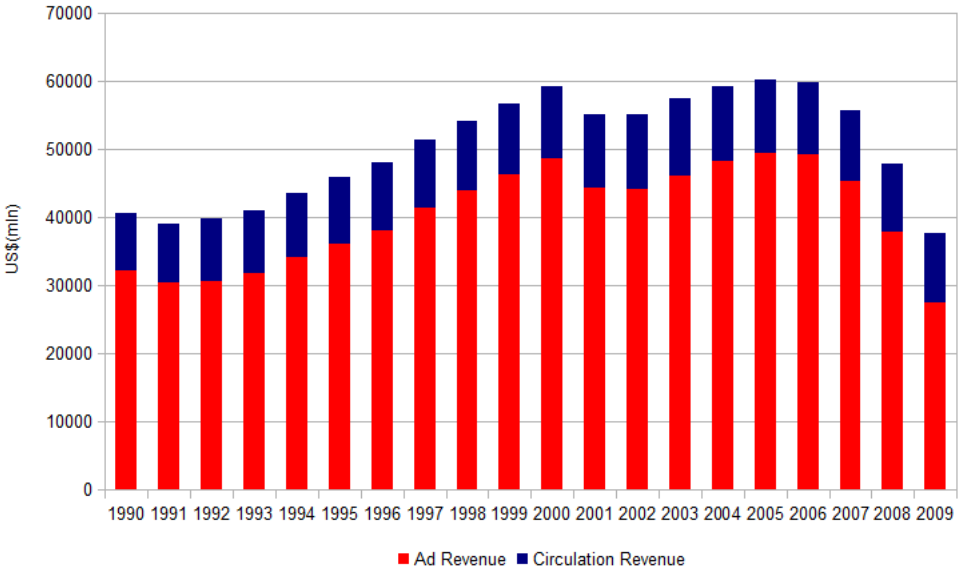

Figure 2.1: Newspaper Revenue in the US, 1990-2009 (Source: Newspaper Association of America)

may assist in mitigating this form of commercial media bias, provided that the media's reputation concerns are sufficiently high, in contrast with recent policy efforts by the U.S. government relaxing media merger and ownership regulations.

This paper fits in with the growing economics literature on the determinants of media bias, which can broadly be divided into two categories - those driven by demand side considerations, and others primarily driven by supply side factors.[1] In demand side models media firms tailor their reports in order to suit the beliefs or opinions of their target audience (Mullainathan and Shleifer, 2005, and Gentzkow and Shapiro, 2010), given that consumers are more likely to read news reports that are in line with their prior beliefs. On the supply side, biased reporting emerges from various other sources that act as inputs in the provision of news, including rent-seeking journalists with career concerns (Baron, 2006) and advertising (Germano and Meier, 2010; Gal-Or, Geylani, and Yildirim, 2012 and Blasco, Pin, and Sobbrio, 2015).

Within this context, a closely related paper to ours is by Ellman and Germano (2009), who propose a model of advertising whereby the media must strike a balance between consumer preferences and advertiser interests in order to maximise its profits. The results show that the media will under-report stories that may be damaging to its advertisers, although this bias disappears as the size of the advertising increases. As mentioned earlier, this paper differs in that we introduce the idea of a news report that may be beneficial to one party and harmful to another, thereby leading to diverging incentives for advertising firms in terms of their preferences for publication/non-publication. This leads to a distinct (yet

[1]For a detailed survey of the relevant literature, refer to Prat and Stromberg (2011) and Blasco and Sobbrio (2012).

complementary) channel through which commercial media bias can arise, namely via competitive concerns.[2]

2.1.1 Media Reportage and Advertiser Incentives: Motivating Examples

To shed light on the potential impact of advertiser-media relations on news reporting, we introduce two real world examples which will also help illustrate the main results derived in this paper. One of the most cited cases of media bias concerns the under-reportage of the scientific evidence on anthropogenic climate change in the U.S. (e.g. Boykoff and Boykoff, 2004). Many (e.g. Oreskes and Conway, 2010) have linked this bias to the relative influence of the U.S. automotive industry in terms of the volume of advertising undertaken each year - Borell Associates estimate that total advertising expenditure by the U.S. automotive sector totaled US$22.6 billion in 2011, the second largest amount by industrial sector in the country.

However, it is also true that there are several firms and industries who would benefit directly from the publication of climate news stories, most notably those who specialise in alternative energy solutions. In fact, the U.S. is the leading market for hybrid vehicles in the world, with total sales exceeding 2 million units over the period 1997-2011 (Wards, 2011), with several U.S. auto manufacturers like Ford and Lincoln introducing their own hybrid models to compete with the likes of Toyota and Honda.

Thus, within this context it would seem as though the publication of news stories confirming the existence of anthropogenic global warming would be beneficial for a growing number of automakers. And yet according to the Pew Research Center, only 58% of Americans believe that there is solid evidence that global warming is occurring, while coverage of climate change issues in the U.S. print media has plummeted by over 80% since 2006 (Pew, 2012b).

Another example is related to the 'pink slime' scandal that rocked the U.S. beef industry in March 2012. Briefly, pink slime or 'lean, finely-textured beef' as it is known within the industry, is a mixture of beef-related trimmings that is processed and used as an additive in various beef products.[3] Pink slime was originally intended for use in dog food, however it has also been used in burger patties, mince meat and other products since the 1990s (even though it was only approved for human consumption by the USDA in 2001). Given that meat manufacturers are not required by law to disclose the pink slime content of their beef products on any labeling, the vast majority of consumers were unaware of the existence of pink slime or its ubiquity in beef products.

[2]The modeling environment also differs significantly in our paper, since in Ellman and Germano (2009) the setup closely follows the two-sided market specified in Rochet and Tirole (2003), while in this paper we develop a model of horizontal and vertical product differentiation.

[3]ABC News reported that, just before the story broke in March 2012, around 70% of ground beef sold in U.S. supermarkets contained pink slime.

Nonetheless, it was only after a social media campaign which kicked off after the airing of an episode of *Jamie Oliver's Food Revolution* in April 2011 that consumer awareness regarding pink slime started growing, culminating in a March 2012 exposé by ABC News. In the aftermath of this news report, Harris Interactive (2012) found that over 76% of American adults were 'at least somewhat concerned' with the inclusion of pink slime in everyday beef products, while major supermarkets and food chains like McDonald's and Taco Bell ceased using pink slime in their products. On the other hand, several firms are expected to have benefited from the publication of this news story[4], in particular organic beef manufacturers since their products are clearly labeled by the USDA (vegan and vegetarian food producers would also have benefited from this scandal). Despite all this, the story was only exposed some 20 years after the first known use of pink slime in meat products.

2.2 The Model

There are two firms (hereafter denoted as Firm A and Firm B respectively) operating in a duopolistic market where each firm produces and sells a horizontally differentiated product. Thus, although firms engage in price competition with one another, they are essentially monopolists within the specific class of product attribute that they produce. For example, this scenario could reflect the automotive industry where although each firm produces private automobiles, Firm A is a European manufacturer while Firm B is a U.S automaker. This means that each firm will capture a specific portion of the automotive market depending on consumer preferences and beliefs regarding the automobile's characteristics, tastes, etc.

There is a continuum of consumers who are heterogeneous in terms of their preferences for each respective firm. More specifically, consumer tastes are denoted by the variable v, where v is uniformly distributed with support $[-1, 1]$. Without loss of generality we state that consumers with $v \in [-1, 0)$ prefer the good offered by Firm B, while the remaining consumers with $v \in (0, 1]$ prefer Firm A's product complement. Furthermore, consumers' willingness to pay for each product is strictly increasing in the absolute value of v for both goods.[5]

However, the relative utility derived from each firm's goods also depends on the realisation of a binary random variable Y, where $Y \in \{a, b\}$ denotes the state of the world. The relative payoffs associated with each state of the world given the choice of product is summarised in the matrix below.

[4]These include a wide variety of firms, from the fast-food chain Wendy's right down to local butchers (ABC, 2012).

[5]Although this is akin to the standard Hotelling setting whereby each firm is located at opposite ends of the 'linear city' to maximise differentiation, the key difference is that in our case there are only two types of goods, meaning that firms cannot position themselves at any point along the taste space.

		Product Choice	
		Firm A	Firm B
State Y	a	1	−1
	b	−1	1

Consumers have a common prior belief regarding the true state Y, where $Pr(Y = a) = \theta$; this is common knowledge to all participants. Thus if $\theta > \frac{1}{2}$, in the absence of the taste parameter v Firm A would capture the entire market, and vice-versa for $\theta < \frac{1}{2}$. So for example, if we continue the automotive industry analogy then the taste parameter v may refer to various idiosyncratic preferences like brand loyalty, country of origin or design, whilst Y may refer to the true likelihood of climate change due to anthropogenic pollution sources like motor vehicles, which if high may persuade consumers to switch to hybrid cars. We can now formally set out the consumers' utility functions for each product:

$$EU_A(v, \theta) = v + \gamma E(Y \mid \theta) - P_A \tag{2.1}$$
$$EU_B(v, \theta) = -v + \gamma E(Y \mid \theta) - P_B, \tag{2.2}$$

where γ $(0 < \gamma \leq 1)$denotes the relative importance that the consumers attribute to the realisation of the state of the world, and P_i denotes the price of good $i = \{A, B\}$ respectively. We assume that each consumer can at most buy one good from either firm.

In addition to the above players, there is also a monopolist media firm who, with a probability α (where $\alpha \in [0, 1)$; this is common knowledge to all participants), receives an unbiased yet noisy signal regarding the true state of Y. More specifically, we parameterise this noisy signal as $\beta \in [0, 1]$, which denotes the probability that $Y = a$ (and, conversely, where $1 - \beta$ indicates the probability that $Y = b$). Thus, we can denote the news report (if published) whenever $\beta \in (\frac{1}{2}, 1]$ as $s = \hat{a}$, and likewise the report when $\beta \in [0, \frac{1}{2})$ as $s = \hat{b}$. This is because, by construction, when $\beta \in (\frac{1}{2}, 1]$ then the media firm's news report (if published) would indicate that the true state of the world is likely to be a; conversely, it would be equal to b when $\beta \in [0, \frac{1}{2})$.

The actual value of β reflects the relative strength of the signal received by the media (in either direction). So for example if the media receives quasi-incontrovertible evidence that $Y = a$ then the value of β would be close to 1, whereas for a weak signal that $Y = a$ then β would be closer to $\frac{1}{2}$.[6]

For example, Y may refer to the rise in global temperatures and rising sea levels caused by global warming, which may induce consumers to switch to alternative energy vehicles like hybrids over more traditional diesel or gasoline options.

[6]We depart from the traditional framework employed in the literature (Gentzkow and Shapiro, 2010) by assuming that the value of β is not fixed and/or dependent on media firm quality, but rather can vary according to the strength of the signal received, since within this context it helps to distinguish between a weak signal and a strong signal, which exist regardless of the media's reporting capabilities. Perhaps a better measure of a media firm's capabilities is α, which reflects the frequency with which a media firm receives signals from its sources.

Y could also represent the presence of pink slime in various beef products, which if true would harm those beef manufacturers who use pink slime to keep costs down, whilst also benefiting those producers who use pure beef ingredients, like organic food stores.

We postulate that the value of β is fully-revealed by the media when it decides to publish a news report $s \in \{\hat{a}, \hat{b}\}$, and thus serves as a basis through which consumers can update their beliefs regarding the true state of Y. Furthermore, the information received by the media is 'hard' as defined by Milgrom (1981), meaning that it cannot be misrepresented or tampered with, although it is possible to withhold its publication. This highlights the fact that, within the context of this paper, the term 'media bias' refers to the deliberate concealment of information by the news media as opposed to the outright garbling of news signals. We also assume that consumers do not receive any ex-post feedback regarding the true realisation of Y, and hence cannot assess the veracity of the news report published by the media. This assumption is consistent with various kinds of news stories, ranging from the suitability of a political candidate to run the country to the severity and/or cause of climate change (see Anderson and McLaren, 2012).

In this paper, the media's only source of revenue is through advertising from either firm. This assumption broadly reflects the modern media landscape where the advent of digital news sources and declining newspaper circulation have meant that media firms must increasingly rely on advertising receipts in order to maintain profit margins. In fact, as highlighted in the introduction, the Newspaper Association of America (2011) estimates that on average advertising revenue constitutes around 75-80% of total newspaper turnover, with this figure rising substantially for other kinds of media like TV news stations and online news portals.

More importantly, in this model advertising will also determine whether the news story is published or not, depending on the relative payoffs under each scenario for both firms. So for example if Firm A believes that it stands to benefit from the publication of a news story, it will advertise in the media in order to ensure that the story is indeed published (and vice-versa for Firm B). The monopolist media firm will then decide whether to publish or withhold the news story on the basis of which firm offers the higher advertising fee. In case of a tie, or if neither firm decides to advertise, then the media would be indifferent between publishing and withholding the news story: we assume, without any loss of generality, that in such instances the media will opt to publish the story.[7]

For the remainder of this paper we shall proceed, without any loss of generality, with the case where consumer prior $\theta \in [\frac{1}{2}, 1]$. In other words, we consider the case where consumer prior regarding the true state of Y is skewed in favour of $Y = a$, and thus Firm A's product offering. This is done since, due to the symmetric nature of the game, the results derived in this paper are equally applicable to the case where $\theta \in [0, \frac{1}{2}]$ by simply re-labelling the notation (e.g. in this case

[7]In this model we largely ignore the informative/persuasive aspect of advertising on consumers, since the main focus is on the impact of advertising on media reports. Nonetheless, we can easily extend the model to include such elements without compromising the validity of our results. For further discussion regarding the role played by advertising in this model, refer to Section 2.4.2.

a news signal $s = \hat{b}$ conforms to people's prior belief, rather than contradicts it as in the $\theta \in [\frac{1}{2}, 1]$ case). Thus, all the conclusions derived hereunder regarding the equilibrium reporting strategies for news signals that confirm or contradict people's prior remain unchanged.

Both firms A and B can fully observe the signal received by the media, and on the basis of this signal they must formulate their optimal advertising strategies. The media's decision to publish or not will depend on whose advertising fee is higher. We use Perfect Bayesian Equilibrium as our equilibrium concept in this paper. The timing of the game is as follows:

1. Nature chooses the state of the world $Y \in \{a, b\}$;
2. With probability α the media receives an imperfect signal $s \in \{\hat{a}, \hat{b}\}$ regarding the true state of Y; this is also observable by Firms A and B, although consumers cannot directly observe it;
3. Firms A and B decide individually whether to advertise in order to ensure that the news signal is published or withheld by the media;
4. The media firm decides which advertisement to accept on the basis of the fee proposed, and hence decides whether to publish the news report or not;
5. Consumers observe whether a news report has been published or not, and update their beliefs regarding Y according to Bayes' Rule (wherever possible);
6. Firms A and B set their respective produce prices P_A and P_B, and consumers decide which good to purchase;
7. Payoffs are realised.

2.2.1 Consumers

As highlighted above, consumers will use the information inferred from the publication or otherwise of a news report by the media in order to update their beliefs regarding the true state of the world Y, which will assist in their purchasing decision from either firm A or B. If a news report is published such that $s \in \{\hat{a}, \hat{b}\}$, then from the firms' perspective consumers will update their beliefs regarding Y according to Bayes' Rule:

$$\psi = \frac{\theta\beta}{\theta\beta + (1-\theta)(1-\beta)} \ ,$$

where ψ represents the consumers' (common) posterior probability that $Y = a$. Hence, the expected utilities derived from Firm A or B can be rewritten as:

$$EU_A(v, \psi) = v + \gamma(2\psi - 1) - P_A \tag{2.3}$$
$$EU_B(v, \psi) = -v + \gamma(1 - 2\psi) - P_B. \tag{2.4}$$

If a news report is not published, then consumers will still update their prior on the basis of what they know regarding the optimal advertising strategies of each firm, weighted by the probability $(1-\alpha)$ that the media did not receive any news story. We denote, by assumption, this updated belief by ω:

$$\omega = (1-\alpha)\theta + \frac{\alpha\theta\tau}{\theta\tau + (1-\theta)(1-\tau)}, \tag{2.5}$$

with τ representing the implied value of β that consumers associate with the deliberate non-disclosure of information by the media, which in turn depends on the firms' incentives to publish/withhold the news story. Thus, with probability $(1-\alpha)$ consumers will believe that no news story has been published since the media has not received a news signal s: in this case, the consumers' prior will be equivalent to their posterior. Conversely, with probability α consumers believe that news signal s is being purposely withheld from publication, inferring a value of τ for the strength of the withheld signal.

Hence in determining the value of ω, consumers will evaluate the conditions under which each firm would seek to get any potential news story published or withheld. Beginning with Firm A, it is clear that any news report s that results in an upward revision in the consumers' prior belief would be beneficial, thus creating an incentive to advertise. More formally, this incentive to publish would arise whenever

$$\psi \geq \theta \Rightarrow \beta \geq \tfrac{1}{2} \ .$$

Hence if $\beta \geq \frac{1}{2}$ Firm A would have a positive incentive to advertise in order to ensure the publication of the news story, for any value of θ; this condition will hold for any $s = \hat{a}$ (the opposite would be true for any $s = \hat{b}$). Similar arguments can also be made for Firm B; in this case, B would always have an incentive to *withhold* the news report when $\beta \geq \frac{1}{2}$ since this would result in a decline in consumer demand for B's product (once again, the opposite also holds for $s = \hat{b}$).

Therefore, both firms are in direct conflict with one another, with the ultimate winner dependent on who offers the highest advertising fee, based on the relative strength of their respective incentives to publish/withhold news. We shall return to the consumers' updating in the non-disclosure case later on after analysing the firms' strategy profiles. Using the above expressions we derive the condition under which a consumer would be indifferent between Firm A and Firm B, regardless of whether the news report is published or not. Setting $EU_A = EU_B$, we can characterise the indifferent consumer type as

$$v^* = \frac{P_A - P_B + 2\gamma - 4\gamma\Theta}{2}, \tag{2.6}$$

where $\Theta = \psi$ if a news report is published, and $\Theta = \omega$ if no news report is published. Therefore, given the distribution of v, we can derive the quantity demanded for each good:

$$D_A = \frac{1}{2} - \frac{P_A - P_B + 2\gamma - 4\gamma\Theta}{4} \tag{2.7}$$

$$D_B = \frac{P_A - P_B + 2\gamma - 4\gamma\Theta}{4} + \frac{1}{2}. \tag{2.8}$$

2.2.2 Advertiser Strategies

We now proceed to analyse the firms' respective advertising strategies which shall form the basis of the media's decision to publish a news report or withhold it. Both firms are assumed to be profit maximisers, where as usual profit π_i $(i = \{A, B\})$ is a function of price P_i, quantity demanded D_i and a constant marginal cost which we normalise to zero. For exposition purposes we also set $\gamma = 1$ and proceed accordingly.

2.2.2.1 Firm A

We start by analysing Firm A's strategy profile. In this case, the firm seeks to maximise profits, $\pi^A = P_A D_A$, with D_A corresponding to (2.7) above. Thus, as expected Firm A's strategy also takes into account Firm B's actions. When news story $s \in \{\hat{a}, \hat{b}\}$ is published, then the resulting price, quantity demanded and profit functions associated with this strategy are shown below

$$\begin{aligned} P_A^K(\psi) &= \tfrac{4(1+\psi)}{3} \\ D_A^K(\psi) &= \tfrac{1+\psi}{3} \\ \pi_A^K(\psi) &= \tfrac{4(1+2\psi+\psi^2)}{9} \ , \end{aligned}$$

where the superscript K is used to denote the firm's price, demand and profit functions given the publication of the news story $s \in \{\hat{a}, \hat{b}\}$. From the above equations, it is evident that $P_A^K\prime(\psi) > 0$ and $D_A^K\prime(\psi) > 0$, while the profit function π_A^K is strictly convex in ψ ($\pi_A^K\prime(\psi) > 0$, $\pi_A^K\prime\prime(\psi) > 0$). Thus, the stronger the signal $s = \hat{a}$ received by the media (corresponding to a high level of β), the higher will Firm A's profits be since consumers will update their prior belief by a larger degree, resulting in higher quantity demanded and a higher price.

If on the other hand the news report is not published, Firm A's profit function (which we denote by $\pi_A^{\emptyset}(\omega)$) is given by

$$\pi_A^{\emptyset}(\omega) = \tfrac{4(1+2\omega+\omega^2)}{9} \ ,$$

where ω is used to denote the non-publication or withholding of the news story. As seen above, $\pi_A^{\emptyset}(\omega)$ is almost identical to $\pi_A^K(\psi)$ above, with the only difference being that now profit is a function of posterior belief ω, which is formulated on the basis of no media report being published. Comparing the two profit functions, it is straightforward to see that $\pi_A^K(\psi) \geq \pi_A^{\emptyset}(\omega)$ whenever $\psi \geq \omega$. The intuition

behind this result is simple - whenever the media publishes a news report s with associated signal strength β that raises consumers' posterior belief ψ above the level that would have been attained under no publication (ω), then it is in Firm A's best interest to ensure that the news story gets published. Given our assumptions regarding the value of β, this will happen when $s = \hat{a}$. Hence, whenever $\psi \geq \omega$ Firm A will advertise in order to ensure that the news story is published. Whether the news report gets published or not will also depend on the actions taken by Firm B, given its own incentives to conceal information that may have a detrimental impact on its profits.

2.2.2.2 Firm B

We now move on to Firm B's strategy profile. Given the conclusions reached in the previous section, Firm B knows that, regardless of its actions, whenever $\psi \geq \omega$ Firm A's (weakly) dominant strategy is to advertise in order to encourage the publication of the news story.

If Firm B decides to advertise, then the media will compare each firm's maximum willingness-to-pay for advertising and select the highest offer, followed by a decision to publish or not depending on which offer is accepted. Conversely, if Firm B decides not to advertise then the news report will be published via Firm A's advertising activities. Firstly though, it is necessary to establish if Firm B would advertise when a news report is either published or withheld. The reasoning here is analogous to that employed for Firm A, with B's strategy taking Firm A's choices into account. If the news report is withheld, then Firm B's price, quantity demanded and profit functions are given by

$$\begin{aligned} P_B^{\emptyset}(\omega) &= \tfrac{4(2-\omega)}{3} \\ D_B^{\emptyset}(\omega) &= \tfrac{2-\omega}{3} \\ \pi_B^{\emptyset}(\omega) &= \tfrac{4(4-4\omega+\omega^2)}{9} \ . \end{aligned}$$

We can observe that in this case both $P_B^{\emptyset}$ and $D_B^{\emptyset}$ are decreasing in ω, whilst $\pi_B^{\emptyset}$ is convex in ω. On the other hand, if the news report is published then we obtain Firm B's profit function

$$\pi_B^K(\psi) = \tfrac{4(4-4\psi+\psi^2)}{9} \ .$$

Once again, comparing the two profit functions it is clear that the optimal outcome for Firm B will largely depend on the value of ψ and ω. More specifically, $\pi_B^{\emptyset}(\omega) \geq \pi_B^K(\psi)$ if and only if $\omega \leq \psi$. The rationale in this case is also fairly straightforward - if the consumers' posterior regarding the probability that $Y = a$ is higher if the news report is published, then Firm B would have a clear incentive to ensure that the story is withheld. Since it knows that under this scenario Firm A will advertise in order to ensure publication, then Firm B's best response would also be to advertise, albeit to encourage nondisclosure of information by the media. Conversely, if $\psi < \omega$ then Firm B would prefer the publication of the story since it would result in a lowering of consumer beliefs in B's direction.

2.2.3 Publication of News

We now have a full characterisation of the two firms' optimal strategies. In equilibrium, whenever $\psi \geq (\leq)\omega$ Firm A would advertise in order to ensure the publication (withholding) of the news story, whereas Firm B would advertise in order to induce the concealment (publication) of the signal. This means that the ultimate decision regarding whether to publish or not lies with the media, who will decide on the basis of which firm offers the highest advertising fee R. For each firm, R reflects the maximum additional benefit that the firm would accrue if it decides to advertise, relative to the no advertising alternative, given their opponent's strategy. It is therefore possible to derive the maximum advertising fee payable R by each firm

$$R_A = \pi_A^K(\psi) - \pi_A^{\emptyset}(\omega) \tag{2.9}$$
$$R_B = \pi_B^K(\psi) - \pi_B^{\emptyset}(\omega). \tag{2.10}$$

It is useful to define $\Delta_{A,B} = R_A - R_B$ which captures the difference between the two maximum advertising fees of firms A and B respectively, where $\Delta_{A,B} > 0$ denotes that Firm A's advertising fee exceeds that offered by B. By comparing R_A and R_B, we arrive at the following result:

Proposition 1. *When the news report received is $s = \hat{a}$ and with signal strength $\beta \in [\frac{1}{2}, 1]$, the media publishes the report. Conversely, when $s = \hat{b}$ and $\beta \in [0, 1-\theta]$, there exists a cut-off value β^L such that for any $\beta \in (\beta^L, 1-\theta]$ the media withholds publication of the report; for any $\beta \in [0, \beta^L]$ the media publishes the news report.*

Proof. See Appendix □

The result follows directly from our original specification of the news signal β as well as the convexity of R_A and R_B in consumer posterior beliefs ψ and ω. This implies that when the news signal received is $s = \hat{a}$, the consumers' prior belief is already skewed in Firm A's direction ($\theta > \frac{1}{2}$), meaning that the advertising fee offered by A in order to ensure publication will always exceed that offered by its opponent. When $s = \hat{b}$, initially Firm A still holds the upper hand for weaker news signals $\beta > \beta^L$, which will be suppressed since they would shift consumers' belief in Firm B's direction if published. Below β^L, then due to convexity the posterior belief under publication ψ would be low enough to make Firm B's advertising fee higher than that offered by Firm A, leading once more to publication of the news signal.

To summarise, Proposition 1 predicts mixed results in terms of the existence of commercial media bias, and that this largely depends on people's prior belief regarding state Y. Briefly, when the news report conforms to people's prior belief, then it will always be published by the media. Conversely, when signals contradict people's prior, then only 'strong' news stories will get published as parameterised by β. This seems somewhat at odds with the real-world experience, particularly

the examples described in Section 2.1.1 where in the case of both anthropogenic climate change and the pink slime debacle the news stories were either toned down or concealed for several years. For example, if we accept the findings from Proposition 1 then any pro-climate change story received by the media should be published since it would conform to people's prior, given that as mentioned earlier 57% of Americans believe that global warming is occurring (Pew, 2012b).

Clearly, one element which is missing from the model is the fact that one party may be larger or more influential than the other in terms of the number of firms involved and hence the advertising revenue on offer. In the model we assume that there are only two firms operating at either end of the spectrum - for example one gasoline automaker and one hybrid car manufacturer - but in reality things may be rather more lop-sided. For example in the automobile case although, as stated earlier, the number of hybrid/electric car manufacturers is on the rise, they are still heavily outnumbered by the more traditional gasoline/diesel producers (very often an automaker would have several gasoline/diesel car models and only one hybrid model).[8] The same can also be said with regards to the US beef industry, since as seen earlier an estimated 70% of ground beef products in supermarkets contained pink slime (as at March 2012).

And yet despite this evident imbalance, in both of these cases there is still enough clout, both in terms of influence and indeed advertising potential, to combat this perceived bias to withhold in media reporting. For example as mentioned earlier, Toyota is one of the leading advertisers in the U.S. with over US$860 million spent annually, and the second largest automotive advertiser (Kantar, 2014). In the next section we shall take a closer look at why media bias persists in such situations by analysing Firm B's incentives to switch over and compete more directly with Firm A in the face of a potentially damaging news story.

2.3 Good News Gone Bad

We now extend the basic framework of the model in order to account for the possibility that a firm may decide to eliminate vertical product differentiation (based on the realisation of state Y) in light of a potentially damaging news story being published. More specifically, we allow for the possibility that whenever $s = \hat{a}$, Firm B has the option to overhaul its processes and eliminate vertical product differentiation, in order to avoid the backlash caused by dwindling demand and prices as a result of the publication of the news story. Similarly, it is also entirely possible for Firm A to switch over whenever $s = \hat{b}$ is published by the media.

Note that although the vertical differentiation element of the two firms' products would be completely eliminated, they still retain the horizontal differentiation element characterised by the parameter v. For example, if Firm B decides to switch over when $s = \hat{a}$ is published, the consumers' utility functions would change slightly, such that

[8] In fact, LMC Automotive (2012) estimate that electric/hybrid cars only constituted around 3.1% of total vehicle sales in the U.S. in 2011.

$$EU_A(v, \psi) = v + \gamma(2\psi - 1) - P_A \tag{2.11}$$
$$EU_B(v, \psi) = -v + \gamma(2\psi - 1) - P_B. \tag{2.12}$$

Similarly, if Firm A were to switch over in response to the publication of a news story $s = \hat{b}$, consumer utility from the consumption of either good would now be equal to

$$EU_A(v, \psi) = v + \gamma(1 - 2\psi) - P_A \tag{2.13}$$
$$EU_B(v, \psi) = -v + \gamma(1 - 2\psi) - P_B. \tag{2.14}$$

There are several ways to interpret the nature of this switch over. Considering the automotive case, this would be akin to a gasoline automaker deciding to produce hybrid cars in light of a considerable change in public perception of anthropogenic global warming (represented by θ) following a serious nationwide news story about the negative impact of automotive emissions on human health and climate change. In this case, although the switch in production to hybrid cars eliminates the vertical product differentiation aspect (i.e. environmental-friendliness), the two cars would still be horizontally differentiated in terms of brand, country of origin, etc., which could all be parameterised by v.[9]

This fact is perhaps even more stark in the beef industry example, since although McDonald's decided to cease using pink slime in its products (thereby regularizing its position in relation to organic food producers), the overall nature of its product offering did not change. McDonald's is still nominally a fast food restaurant that caters to a specific target market which is manifestly distinct from that targeted by organic beef burger manufacturers.

Therefore, under this modeling environment both firms have the option to switch their production processes and eliminate vertical product differentiation. This switch over entails a fixed cost, denoted by F^{SW}, reflecting the fact that overhauling a firm's production process requires a non-trivial investment in new technology, research and development, and other setup costs.[10]

Alternatively, F^{SW} could also capture the extent to which one firm's production processes are safeguarded by patents and other Intellectual Property Rights (IPR) protection. In any case, the purpose of F^{SW} is to show that switching over is not a costless process, and must be factored in the switching firm's profits before deciding to proceed. Once again, variable costs are assumed to be constant and equivalent across both firms, and are thus normalised to zero to facilitate the exposition.

[9]When Ford decided to produce hybrid and electric cars it still retained its brand and U.S.-made characteristics; the same can be said for other manufacturers like Lincoln, Chevrolet, etc.

[10]We assume that both firms face the same (fixed) switching cost F^{SW} . The value of F^{SW} is common knowledge to all participants

At this point it is necessary to consider the conditions under which either firm would decide to switch over after a news signal s has been received by the media. In case of a switch over then the price, quantity demanded and profit for each firm are now

$$P_A^{SW} = P_B^{SW} = 2$$
$$D_A^{SW} = D_B^{SW} = \tfrac{1}{2}$$

$$\pi_i^{SW} = \begin{cases} 1 - F^{SW} (\text{for firm that switches over}) \\ 1 (\text{for other firm}), \end{cases}$$

where $i \in \{A, B\}$. Firstly, we look at a firm's switching decision when the news story s is published by the media. In this case, we compare the profit earned by the firm if it did not switch over to the profit earned above if it decided to proceed with the switch over. Denote the difference between π_A^{SW} and $\pi_A^K(\psi)$ by Σ_A (for Firm A), and likewise Σ_B for the difference between π_B^{SW} and $\pi_B^K(\psi)$:

$$\Sigma_A = \frac{5 - 8\psi - 4\psi^2}{9} - F^{SW} \tag{2.15}$$
$$\Sigma_B = \frac{16\psi - 7 - 4\psi^2}{9} - F^{SW}. \tag{2.16}$$

A simple manipulation of the expressions yields the following result:

Lemma 1. *When news report $s = \{\hat{a}, \hat{b}\}$ is published by the media, there exists a cut-off value for the strength of the media signal, denoted as β^{SW}, where:*

$$\beta^{SW} = \frac{(1-\theta)(4-3\sqrt{1-F^{SW}})}{2(2-3\theta)-3\sqrt{1-F^{SW}}(1-2\theta)} .$$

When $s = \hat{a}$, Firm B switches over if and only if $\beta > \beta^{SW}$ and $F^{SW} < \frac{5}{9}$. Conversely, when $s = \hat{b}$ Firm A switches over if and only if $\beta < 1 - \beta^{SW}$ and $F^{SW} < \frac{5}{9}$.

Proof. See Appendix □

Hence, for a sufficiently small F^{SW}, Firm B would eliminate the vertical product differentiation element previously present in the market whenever the media publishes a strong (high β) signal for $s = \hat{a}$. Similarly, when a strong (low β) contrary signal $s = \hat{b}$ is published Firm A would be induced to switch over and mimic its competitor, eliminating vertical product differentiation, provided that switching costs are low. The intuition behind this result is relatively simple - stronger news signals published in the media will induce a significant shift in consumer beliefs regarding the realisation of Y in one firm's favour, to the detriment

of the other, thereby leading the latter to switch over in order to earn higher profits.

However this is only one part of the story, since a firm may still opt to switch over even when no news report is published in the media. The reason is that in the absence of a published story, consumers may still update their beliefs sufficiently in the direction of one firm in order to render a switch over attractive to the other, relative to the status quo. When a story is withheld, a firm has two options - either stick to its current production process (earning a payoff of $\pi_i^{\emptyset}$) or switch over (resulting in a profit of π_i^{SW}). If we compare the two payoffs it is easy to show that a firm will always (at least weakly) prefer to switch over whenever (denoting posterior belief under no publication by ω)

- When $s = \hat{a}$: Firm B would switch over when $\omega_{s=\hat{a}}^{SW} \geq \frac{4-3\sqrt{1-F^{SW}}}{2}$;
- When $s = \hat{b}$: Firm A would switch over when $\omega_{s=\hat{b}}^{SW} \leq \frac{-2+3\sqrt{1-F^{SW}}}{2}$,

for a sufficiently-low value of F^{SW}. Again, the logic behind this result is clear - if consumers update their beliefs regarding Y very strongly in one firm's favour, even in the absence of any published news signal, then it would make sense for the other firm to switch over and eliminate vertical product differentiation.

Since we allow for switching over, we require a different expression to represent the posterior belief under no publication ω to the one used in the previous section. This new posterior belief under no publication, depending on prior belief θ, can be expressed as

$$\omega_{\theta \geq \frac{1}{2}} = (1-\alpha)\theta + \alpha\{\theta(\tfrac{\theta+1}{2}) + (1-\theta)(\tfrac{\frac{3}{2}-\omega}{2})\} \ .$$

Briefly, when no news report is published consumers believe, with probability $(1-\alpha)$, that this is due to the fact that no news report has been received by the media, in which case consumers' posterior belief will simply be equal to their prior. With probability α consumers believe that the media is purposely withholding the news signal: with a probability of θ this withheld signal is such that $s = \hat{a}$, and with probability $(1-\theta)$ the signal is $s = \hat{b}$.[11].

Therefore, we can find the conditions under which each would prefer to switch over even when no news story is published. The following result is obtained:

Proposition 2. *Given that no news story is published in the media, there exists a unique threshold prior belief $\theta^{SW}(\alpha, F^{SW}) \in [\frac{1}{2}, 1]$, such that when $\theta > \theta^{SW}$ Firm B switches over and eliminates vertical product differentiation, for any switching cost $F^{SW} \in [0, \frac{5}{9})$.*

Proof. See Appendix □

[11]When $\theta < \frac{1}{2}$, posterior belief under no publication would be equal to $\omega_{\theta<\frac{1}{2}} = (1-\alpha)\theta + \alpha\{\theta(\frac{\frac{3}{2}-\omega}{2}) + (1-\theta)(\frac{\theta}{2})\}$

What this result shows is that for a sufficiently large prior belief $\theta \in [\frac{1}{2}, 1]$ such that $\theta > \theta^{SW}(\alpha, F^{SW})$, when no news story is published consumers will update their beliefs regarding Y such that Firm B is compelled to switch over in order to avoid a drastic decline in demand and profits. The idea behind the result is that if consumer prior belief is already skewed towards Firm A's product offering then, based on our formal definition of $\omega_{\theta>\frac{1}{2}}$ above, the non-publication of news will only serve to strengthen these slanted beliefs in Firm A's favour, to the detriment of Firm B's current status quo. This would therefore induce Firm B to mimic Firm A and eliminate vertical product differentiation.[12]

Notice that θ^{SW} is strictly decreasing in α (the probability that the media receives a news signal), since for higher levels of α consumers would become increasingly suspicious that the media is purposely withholding information, resulting in a higher updating of prior and increasing the likelihood of exceeding the $\omega_{\theta\geq\frac{1}{2}} = \frac{4-3\sqrt{1-F^{SW}}}{2}$ threshold. Conversely θ^{SW} is strictly increasing in the value of F^{SW} since a higher switching cost reduces the likelihood that Firm B would undertake the switch-over, thus lowering the probability that the news story (if received) was purposely withheld in order to prevent the switch-over (which would occur if $\beta \geq \beta^{SW}$, where β^{SW} is also increasing in F^{SW}).

Clearly, a switch-over by one firm would have an impact on the other firm's profits. A simple comparison of $\pi_i^K(\psi)$ (the profits earned if the story is published, under no switch over) with π_i^{SW} (where, as usual, $i \in \{A, B\}$) reveals that $\pi_i^K(\psi) \geq \pi_i^{SW}$ whenever

- For $s = \hat{a}$ and Firm B switches over: $\psi \geq \frac{1}{2}$;
- For $s = \hat{b}$ and Firm A switches over: $\psi \leq \frac{1}{2}$,

which, by definition (given our assumptions regarding the value of β in either case) will always hold true. Therefore, whenever a news story $s = \hat{a}$ is published, Firm A is always better off when Firm B does not switch over, since the increase in direct competition brought about by the switch-over results in an inevitable decline in A's profits (and vice-versa for $s = \hat{b}$).

Next, we look at the non-publication scenario, and more specifically we compare a firm's profits under no publication ($\pi_i^{\emptyset}$) to the profit level when the news story is published and the firm's competitor decides to switch over (π_i^{SW}). Denote R_i^{SW} as the difference between these two profit levels, where once again it is straightforward to show that $R_i^{SW} \geq 0$ whenever

- For $s = \hat{a}$: $R_A^{SW} \geq 0$ if $\omega \geq \frac{1}{2}$;
- For $s = \hat{b}$: $R_B^{SW} \geq 0$ if $\omega \leq \frac{1}{2}$.

[12]Note that when $\theta < \frac{1}{2}$, by symmetry we can derive an equivalent condition for $\omega_{\theta<\frac{1}{2}}$, whereby if prior belief is such that $\theta < 1 - \theta^{SW}$ then Firm A would prefer to switch over even in the absence of a published signal.

Based on our definition of $\omega_{\theta \geq \frac{1}{2}}$ given above, it is easy to see that since consumer prior $\theta \geq \frac{1}{2}$, it follows that the first condition will always hold, whereas the second condition will, by definition, never hold. Once again this is due to the fact that when no news report is published, consumers are uncertain as to the true value of Y meaning that their posterior belief in the face of such uncertainty will invariably skew in the direction of their prior ($\theta \geq \frac{1}{2}$). Hence, an immediate corollary that emerges from this result is that when $\theta \geq \frac{1}{2}$ Firm B would always weakly prefer the news story $s = \hat{b}$ to be published, even if this leads to Firm A's switch over. This is because when $\theta \geq \frac{1}{2}$ consumer prior belief is slanted in favour of Firm A's product, meaning that if no news report were published then consumer posterior belief will still remain (due to Bayesian updating) in A's direction, to the detriment of Firm B, meaning that publication (and Firm A's switch-over) would be preferable to no publication.

Combining these results, we can state without proof the following lemma (assuming, as always, that switching costs are $F^{SW} \in [0, \frac{5}{9})$):

Lemma 2. *When a news report $s = \hat{a}$ with signal strength $\beta > \beta^{SW}$ is received by the media, Firm A advertises in order to conceal the news report for any $\theta \in [\frac{1}{2}, \theta^{SW})$. Conversely, if $s = \hat{b}$ and $\beta \in [0, \frac{1}{2})$ Firm B advertises in order to ensure that the news story is published.*

The rationale behind this lemma is clear. Whenever a published news signal $s = \hat{a}$ is strong enough (i.e. $\beta > \beta^{SW}$) to induce Firm B to switch over, Firm A has an incentive to ensure that the news story is withheld from publication via advertising, even though the news story $s = \hat{a}$ is at face value favourable for Firm A. Note that when $\theta > \theta^{SW}$ it would make no sense for Firm A to attempt to conceal the news story, since regardless of whether the story is published or not Firm B will always switch over as consumer updating will skew heavily towards Firm A's product offering, inducing the switch-over by B. On the other hand, when $s = \hat{b}$ then Firm B will always have a positive incentive to ensure publication, even when signals are strong enough to induce a switch-over by Firm A, since otherwise consumers' posterior belief under no publication would skew in A's favour.

We have already seen the conditions under which a firm would be induced to switch over: now it is worth analysing the potential switcher's advertising strategies. In this case, we compare the profit level under no publication and no switching ($\pi_i^{\emptyset}$) with profits under publication and switching (π_i^{SW}). We denote the difference between the two by R_i^{SW*}. It can be shown that, for $F^{SW} \in [0, \frac{5}{9})$:

- When $s = \hat{a}$: Firm B would prefer non-publication ($R_B^{SW*}(\omega) \geq 0$) for any $\theta \in [\frac{1}{2}, \theta^{SW})$;
- When $s = \hat{b}$: Firm A would also prefer non-publication ($R_A^{SW*}(\omega) \geq 0$) for any $\theta \in [\frac{1}{2}, \theta^{SW})$.

The intuition behind this outcome is also relatively straightforward. In the first instance, switching over entails both a fixed cost F^{SW} as well as a more

direct competitor in Firm A, both of which would lead to lower profits relative to the status quo where no news story is published and consumer updating under no publication does not skew too strongly in Firm A's favour. On the other hand, when $\theta \geq \theta^{SW}$ then Firm B would be indifferent between publishing or withholding the news story, since as seen from Proposition 2 consumer updating under no publication would still be strong enough to induce Firm B to switch over. For $s = \hat{b}$ the situation is more straightforward - Firm A would always prefer non-publication since in this instance consumers' posterior belief would still skew in A's favour, as opposed to the publication scenario.

2.3.1 Equilibrium Reporting

We can now utilise the results derived in the previous section in order to determine the media's equilibrium reporting behaviour. This will largely be determined by the interaction between each firm's optimal advertising strategies, based on their relative incentives to publish the story or not, as well as consumers' posterior beliefs under each scenario.

The first result presented here is in relation to strong news signals that conform to people's prior belief - i.e. where $\beta \geq \beta^{SW}$ when $s = \hat{a}$ since $\theta \in [\frac{1}{2}, 1]$ - and can be summarised as follows:

Proposition 3. *When switching costs $F^{SW} \in [0, \frac{5}{9})$, if the news signal $s = \hat{a}$ is such that its strength $\beta > \beta^{SW}$, then for consumer prior $\theta \in [\frac{1}{2}, \theta^{SW})$ the media withholds the signal from publication.*

This result, which is one of the key findings in this paper, follows directly from the results derived in the previous section. In essence, given the conditions set out above, we showed how both Firms A and B would always have at least a weakly positive incentive to ensure that strong news stories are not published by the media, meaning that both firms would advertise in order to conceal the story. Therefore the outcome in this case is unambiguous, since regardless of which firm offers the higher advertising fee the result would be the withholding of strong news signals.[13]

It is worth highlighting that the above only holds true for $\theta \in [\frac{1}{2}, \theta^{SW})$. Whenever the consumers' prior belief is such that $\theta \geq \theta^{SW}$ then both firms would be indifferent in terms of the publication or withholding of the news signal, since the consumers' posterior belief would still induce a switch over by Firm B in either case, meaning that there is no incentive for either firm to advertise in order to ensure any particular outcome.

Continuing with our analysis of confirmatory (in terms of people's prior) news stories, we now move onto more moderate signals where $\beta \in [\frac{1}{2}, \beta^{SW})$. Once again, the focus is on prior $\theta \in [\frac{1}{2}, \theta^{SW})$. In this case, as stated earlier Firm B would not switch over, implying that the profit levels under publication for firms A and B

[13] In the case when $\theta \leq \frac{1}{2}$ and $s = \hat{b}$ the result is reversed such that whenever $\beta \leq 1 - \beta^{SW}$ these signals are withheld from publication (where a low β is equivalent to a strong signal in B's favour).

are π_A^K and π_B^K respectively, and $\pi_A^\emptyset$ and $\pi_B^\emptyset$ under no publication. As discussed earlier, each firm's preference over the publication or withholding of the news story depends on the relative values of ψ and ω. With these factors in mind, we can formally state the following result:

Proposition 4. *When a news signal $s = \hat{a}$ is received such that signal strength $\beta \in [\frac{1}{2}, \beta^{SW})$ and consumer prior $\theta \in [\frac{1}{2}, \frac{3}{4})$, there exists a unique value of β, denoted as $\beta^T(\alpha, \theta) \in [\frac{1}{2}, 1]$. If $\beta^T > \beta$ the media withholds publication of the news report. Conversely, for $\beta^T < \beta$ the media publishes the news report.*

Proof. See Appendix □

What this result implies is that weak news signals that conform to people's prior may, under certain conditions, be withheld from publication. The reason for this is that in the absence of any news report consumer posterior would actually be higher relative to the case where the (weak) signal is published. This over-updating of beliefs mainly comes about due to the result described in Proposition 3 earlier, where extreme or strong confirmatory news signals would be withheld from publication, meaning that posterior belief under no publication would now have to take this possibility into account.

We now turn our attention to news signals that contradict people's prior belief - in this case, $s = \hat{b}$. As shown earlier, when $\theta \in [\frac{1}{2}, 1]$ and $s = \hat{b}$ we have a misalignment of incentives for both firms since Firm A's dominant strategy would be to withhold the story whereas Firm B would always prefer to have it published. We therefore have a similar situation to that described in Proposition 1, which leads to the following result:

Proposition 5. *When a news signal $s = \hat{b}$ is received by the media with signal strength $\beta \in [0, 1-\theta]$, and with consumer prior $\theta \in [\frac{1}{2}, \theta^{SW})$, there exists a cut-off value β^L such that for any $\beta \in (\beta^L, 1-\theta]$ the media withholds the news report from publication: conversely, for any $\beta \in [0, \beta^L)$ the media publishes the news report.*

Proof. See Appendix □

Therefore, whenever the news signal contradicts consumer prior, then weaker stories (where $\beta \in (\beta^L, 1-\theta]$) will not get published, while stronger signals ($\beta \in [0, \beta^L)$) are published. As with Proposition 1, the result is driven by the convexity of advertising fees R_A and R_B in consumer posterior belief, which means that for weak signals the 'losing' firm (in this case Firm A) has a stronger incentive to withhold the signal that the winner (Firm B), with the opposite holding true for stronger signals. A direct implication of this result is that, contrary to to the situation where the consumers' prior is aligned with the media's news signal (e.g. $\theta > \frac{1}{2}$ when $s = \hat{a}$), in this case only stronger news signals ($\beta < \beta^L$) are published in the media.

We now have a complete characterisation of the equilibrium advertising strategies adopted by both firms (given consumer belief regarding Y) as well as the

media's equilibrium reporting behaviour. The key findings can be summarised as follows:

- For news stories that conform to people's prior belief (e.g. $s = \hat{a}$ and $\theta \in [\frac{1}{2}, 1]$), strong news signals will be withheld from publication by the media, while moderate signals will be published;
- For stories that contradict people's prior (e.g. $s = \hat{b}$ and $\theta \in [\frac{1}{2}, 1]$), weak signals will be withheld from publication, while stronger signals will be published.

These results provide an alternative rationale for the continued bias and/or under-reporting in the news media, since they suggest that even the advertiser who stands to benefit most from the publication of the news story has an incentive to keep the report under wraps. Therefore in the automotive case considered in Section 2.1.1, hybrid/electric automakers would seek to withhold any strong evidence of anthropogenic climate change from publication in the news media, since this may result in even more automakers introducing rival hybrid/electric models, eroding the incumbents' profit margins. The same can also be said for the US beef industry and the pink slime scandal, since 'pure' or organic beef producers/sellers would, according to this paper, have benefited from keeping the story hidden given that this would have prevented other mass-market firms from switching over to 100% beef products, hence enabling them to maintain their leadership within their specific market segment.

2.4 Discussion and Extensions

In this section we look at how the model's main findings relate to the real-world, with a particular focus on U.S. media coverage of global warming issues over the last few years. We then consider a number of extensions to the model in order to assess the robustness of our results, including a discussion on the role played by advertising as a source of information/persuasion for consumers, the impact of media competition, as well as reputational concerns of media firms.

2.4.1 U.S. Newspaper Coverage of Climate Change Issues

In this section, we show how the main theoretical findings predicted by our model can be used to explain real-world media reporting patterns. Specifically, we compare the model's predictions regarding climate change reportage to data from the U.S. print media. Figure 2.3 below plots annual reporting on climate change issues in the leading U.S. newspapers over the period 2004-2010, and U.S. public opinion regarding the man-made origins of climate change over the same period. The survey data is taken from the 2011 Environment Poll by Gallup (2011).[14] For

[14]The exact question posed in this survey was:"And from what you have heard or read, do you believe that increases in the Earth's temperature over the last century are due more to the

newspaper climate coverage, we collect data on the total number of articles related to the scientific evidence on global warming published by a sample of 144 leading U.S. newspapers, accounting for around 60% of total newspaper circulation in the country (NAA, 2012), using the NewsLibrary online database and the ProQuest library as our main sources.Full details regarding the data collection process are provided in Appendix B.

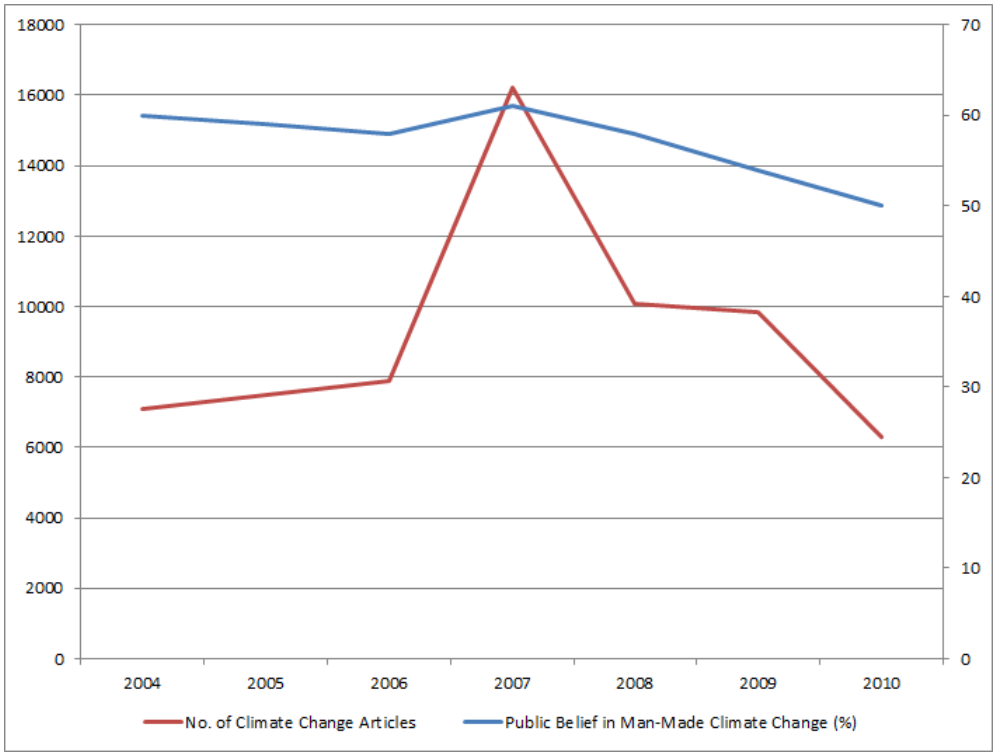

Figure 2.2: U.S. Public Belief in Anthropogenic Origins of Climate Change and No. of Published Newspaper Articles on Global Warming, 2004-2010 (Sources: Gallup (2011), NewsLibrary, ProQuest)

As seen from Figure 2.2, public perception regarding the anthropogenic origins of global warming has been fairly steady over the last few years, with an average of around 55% of Americans agreeing with the assertion that human activities are the main contributors to climate change. Based on this, our model predicts that coverage of climate issues in the media should mainly consist of stories that include moderate evidence supporting this view, or conversely of stories that strongly contradict this belief. This is partially supported by the findings in Boykoff and Boykoff (2004) as mentioned earlier, with the vast majority of climate-based news reports presenting a relatively balanced view regarding the man-made nature of global warming. Clearly, the idea of news 'strength' cannot be derived from our coverage dataset, since it only reports the number of published articles on climate-related evidence. Nonetheless we can see that, with the anomalous exception of 2007 aside (which was mainly driven by reports regarding

effects of pollution from human activities, or natural changes in the environment that are not due to human activities?".

Al Gore's 'An Inconvenient Truth'), climate coverage in the leading U.S. newspapers was largely unchanged over the period under review, and in fact dipped slightly towards the end of the decade.

What is also clear from the graph is that in 2010 both coverage of climate evidence and public opinion dipped. This may have been due to the eruption of the so-called climategate scandal in 2009, when leaked email exchanges among leading climatologists from the University of East Anglia's Climate Research Unit purportedly showed that global warming data had been manipulated and that scientific evidence which downplayed the extent of rising temperatures and other implications of global warming were wilfully suppressed. These allegations were strongly rebutted by the scientists involved, who claimed that the emails were purposely taken out of context to distort the truth.[15] Hence, this type of news story would constitute a 'strong' contradictory signal in terms of people's prior belief, which our model predicts would be published in the media due to the beneficiary's strong incentives. In fact, Oreskes and Conway (2010) state that the mainstream media preferred to focus on the initial furor surrounding the scandal with very little coverage afforded to the subsequent exoneration, pointing the finger at the closeness of the relationship between the fossil fuel industry (as advertisers) and the media.

This seems to be supported by the data, since apart from the dip in evidence-based news reports on climate change in 2010 we also observe a significant drop in public opinion regarding the human origins of global warming, from 54% to 50%. In addition, a quick search for news stories related to climategate over the period 2009 to 2010 shows that in total around 300 articles were published on the subject, which equates to approximately 5% of total newspaper climate coverage over the period. This figure is significant, particularly when considering the fact that by mid-2010 all investigations had been completed and any allegations of impropriety or fraud summarily dismissed. Nonetheless, it appears as though the climategate story may have had the desired effect in terms of tempering public belief in anthropogenic climate change.

2.4.2 Informative Advertising

In this paper the only role played by advertising is to influence the media's publication decision. We therefore ignore any other direct impact of advertising on consumers as described in the literature. In our setup, advertising does not provide any information regarding product attributes (Telser, 1964 and Dukes, 2004), nor does it directly persuade consumers to purchase it (Bloch and Manceau, 1999 and Johnson and Myatt, 2006). Similarly, in the model advertising has no signalling value whereby higher advertising fees are reflective of higher product

[15]Furthermore in the aftermath of the scandal various investigations were initiated in both the U.S. and the UK in order to establish whether the leaked emails demonstrated any sort of fraud or misconduct on the part of the scientists involved, with each investigation dismissing the allegations.

quality (Nelson, 1974 and Milgrom and Roberts, 1986).[16] We omit such considerations since the primary aim of this paper is to look at the impact of advertisers on media reporting behaviour. The only information that consumers can glean from advertising is when no news report is published, since this would be indicative of a purposely withheld news report, although s/he would still not be able to accurately tell whether the withheld signal is $s = \hat{a}$ or $s = \hat{b}$.

Nonetheless, a couple of points are worth mentioning. Firstly, the model can easily be generalised to incorporate informative advertising. More specifically, we can readily think of the media report s as being an advertisement that indicates some product attribute which affects consumer payoffs. For example, in our automotive case we can think of Y as being the probability that Firm A's cars are more environmentally-friendly or energy efficient than Firm B's offering, with $s = \hat{a}$ being an advertisement highlighting the impressive green features of Firm A's vehicles (e.g. hybrid technology, fuel efficiency ratings, etc.). In this case, Firm A's incentive to withhold this signal stems from its desire to prevent Firm B from switching over and adopting its now-advertised technology or product features (provided that consumers' prior is already skewed in A's favour).

Similarly, we can easily incorporate more formal informativeness/persuasiveness features of advertising in our baseline model without altering the results obtained. For example, it is possible to allow for advertising to directly augment the utility derived by consumers from purchasing either good depending on who decides to advertise. There are several potential ways of doing this. Denote the 'persuasive' value of advertising by either firm as λ, where $\lambda > 0$. The most straightforward way of incorporating this effect would be to include λ additively in both (1.3) and (1.4), interacted with a binary indicator variable M_i (M_A in (1.3) and M_B in (1.4)) such that $M_i = 1$ when firm i advertises, and 0 otherwise. Thus, when Firm A advertises, this would raise the expected utility that consumers obtain from product A by λ.

Another way of doing this would be to interact λM_i with the taste parameter v in (1.3) and (1.4), in order to capture the idea that advertising will only raise the expected payoff of consumers whose tastes are already skewed in favour of the product's horizontal attributes. For example, an advertisement by Firm A will only raise the expected utility derived from consuming product A for those people whose taste parameter $v \in (0, 1]$, as shown below:

$$EU_A(v, \psi) = (1 + \lambda M_A)v + \gamma(2\psi - 1) - P_A \ .$$

In this setup the advert simply serves to underscore the product's attributes (e.g. the automobile's American origins), which may appeal to some people and not to others. Notice that when $v \in [-1, 0)$ advertising would actually have a *negative* impact on expected utility from consuming Firm A's product. This reflects the fact that consumers may derive disutility from advertising, particularly if the information being highlighted in the advert does not conform to people's

[16] For a detailed treatment of the extensive advertising literature in economics, see Bagwell (2007)

tastes (a refinement of the 'nuisance' effect of advertising as described in Peitz and Valletti, 2008 and Crampes et al., 2009).

Alternatively, it is also plausible to assume that a firm's advertising will only increase expected utility if it conforms to people's beliefs regarding the state of the world Y. For example, Firm A's advertising would raise expected utility from consuming A's products if and only if posterior beliefs ψ (under publication) or ω (under no publication) $\in (\frac{1}{2}, 1]$, as shown below:

$$EU_A(v, \psi) = v + \gamma(2\psi - 1) + \lambda M_A \psi - P_A \ .$$

Once again, the basic idea is that advertising that highlights the product's vertical qualities (e.g. its environmentally-friendly credentials) would only raise expected payoffs if consumers actually value these attributes: if not, then once again advertising will lead to a reduction in expected utility due to nuisance factors. Hence, based on the above discussion, we have shown how we can extend the basic model described in this paper to incorporate the informativeness/persuasiveness of advertising as traditionally described in the literature. More importantly, it is also straightforward to see that including any (or a combination of) these elements in our analysis would not have any meaningful impact on the main results derived in this paper, since the *relative* incentives to publish or withhold news stories would remain largely unchanged, yielding the commercial media bias patterns described above.

Finally, one may argue that our current specification closely resembles a model of bribery or corporate lobbying, since in effect the firms are paying in order to influence the media's reporting behaviour without any of the other canonical features of advertising cited in the literature. Although this is valid, in reality it would be more appropriate to interpret such payments as advertising rather than bribery. Firstly, commercial bribery is illegal in several countries including the UK (the 2010 Bribery Act) and the U.S., where it is punishable as a felony in 36 out of 51 states and forms part of the definition of 'aggravated felony' under federal immigration law. Therefore, in many countries the most straightforward (and legal) way that firms have of influencing media reportage would be through advertising, particularly since as mentioned earlier such revenues are crucial for the continued survival of private media firms.

Secondly, there are numerous real-world cases of this kind of advertiser-influenced media reportage. For example Warner, Goldenhar, and McLaughlin (1992), in a study on magazine reportage of tobacco news stories over the period, state that magazines that did not contain any cigarette advertisements were circa 40% more likely to contain news stories that related the health hazards of smoking to readers compared to other magazines that contained cigarette advertising. There are also several well-documented cases of firms ending their advertising relationships with certain media organisations in response to the publishing of critical or potentially-harmful (yet factual) news reports. For example Bagdikian (2000) relates how in 1957 tobacco companies ceased their advertising activities in Readers' Digest following the publication of an article on the negative effects of smoking

on health. Therefore, based on the above it seems reasonable to interpret the payments described in this paper as advertising fees.

2.4.3 Media Competition

In this section we relax the assumption of a media monopolist in order to assess whether competition affects the conclusions derived so far in this paper. We start off with the case where media firms are identical to each other in terms of the magnitude of α, meaning that media firms are of equivalent quality in terms of their likelihood of receiving a news signal.

In this case, an immediate observation is that, for any $N > 1$ media firms in the market then this would ensure the publication of weak news signals (corresponding to $\beta \in [\frac{1}{2}, \beta^{SW})$) that conform to people's prior even when $\theta \in [\frac{1}{2}, \theta^{SW})$. This is because with competition Firm A cannot afford to suppress weak signals in all media outlets and take advantage of consumer over-updating whenever $\beta^T > \beta$. For example, if $N = 2$ and we have a duopolistic media market then Firm A can, at most, suppress such signals in one media firm. This is because it is relatively easy to show that R_A, which is the advertising fee offered by Firm A, is always less than double the advertising fee offered by Firm B (R_B), for $\beta \in [\frac{1}{2}, \beta^{SW})$. Hence, although one media firm may have an incentive to withhold publication of the news story, the other would accept Firm B's advertising fee and publish the story. A similar argument can also be made for contradictory news signals $s = \hat{b}$ such that $\beta \in (\beta^L, 1 - \theta]$, since although in this case Firm A's advertising fee would exceed B's, the existence of $N > 1$ media firms will ensure its publication.

Therefore, media competition ensures the publication of weak news stories, both when $s = \hat{a}$ and $s = \hat{b}$. However, in terms of extreme or strong signals where $\beta \in [\beta^{SW}, 1])$ then things are unchanged - these stories would still not be published whenever $\theta \in [\frac{1}{2}, \theta^{SW})$. The reason is that both firms have the incentive to withhold the news story, meaning that both firms will advertise towards this end. In practice this will simply mean that each firm will split their advertising fees accordingly across each of the $N > 1$ media firms operating in the market. The end result would be that strong news signal would still be withheld from publication, regardless of the level of media competition.

Similar results are also achieved when considering heterogeneous media firms that differ in terms of the likelihood of receiving a news signal, α. Once again, the existence of more than one media outlet breaks the suppression of weak signals, and yet strong signals are still withheld from publication due to the congruence of each advertiser's interests. Nonetheless, a couple of additional interesting insights are worth noting.

For exposition purposes we consider the duopolistic media market situation with one high-quality media outlet, denoted by α^H, and a low-quality outlet with α^L, where $\alpha^H > \alpha^L$. For simplicity, we also assume that despite the quality difference, both media firms have received an equivalent news signal $s = \hat{a}$. If we assume that each outlet has a fixed number of consumers who cannot read the

other outlet's news, then when media signals are weak, such that $\beta \in [\frac{1}{2}, \beta^T)$, it is always optimal for both Firms A and B to advertise in the high-quality media outlet (α^H) rather than the low-quality outlet (α^L). The reason is that since β^T is increasing in α, Firm A knows that when the high-quality media outlet withholds a news signal its consumers will over-update their prior more than the low-quality media outlet's consumers, who will tend to rely more on their prior (hence $\tau^T \to \frac{1}{2}$).

Therefore Firm A will have a strong incentive to advertise with the high-quality media outlet to withhold the news story, while Firm B will also opt for the same media outlet, except in this case the aim would be to ensure publication. In this scenario, the model predicts that the high-quality media outlet will withhold the news signal, while the low-quality outlet will publish it. Next, we consider the case of moderate news signals where $\beta \in [\beta^T, \beta^{SW})$. Now, it is optimal for both Firms A and B to advertise with the low-quality (α^L) media outlet. The rationale is similar to that expressed before, in that Firm A knows that consumers of the high-quality media outlet will still update their prior more than those for the low-quality outlet, given that β^T is increasing in α.

2.4.4 Reputation Concerns

So far we have omitted any demand-side media considerations from the market for news, focusing instead on how advertising (a supply-side factor) influences media reporting behaviour. We now introduce reputation effects into our model, reflecting the fact that the willful suppression of information by media outlets is detrimental to consumers. The broad idea in this case is that whenever the media does not publish a news story s, consumers (regardless of their taste parameter v) can 'punish' the media, either by reducing their actual consumption of news or else by reducing their trust in the media. Regardless of the exact nature of the punishment, the non-publication of a news story has a negative impact on the media's profits.

To keep things as general as possible, assume that when no news story is published, the media incurs a disutility of $L_R = L(\alpha)$, where L_R denotes the reputation cost associated with non-publication. Note that $L'(\alpha) > 0$ since the higher the value of α then it is even more likely that a non-publication is the result of a deliberate withholding of information rather than the lack of an actual signal.

Once again, we assume that $\theta \in [\frac{1}{2}, \theta^{SW})$. The key to determining whether the media would still opt to withhold news signals lies in comparing the revenue accrued from non-publication to the reputation cost L_R. Since the media's sole source of revenue is from advertising, then whenever

$$R_i \geq L(\alpha) \ ,$$

provided that R_i is being paid to withhold the news signal, the media will always prefer to do so. In our main results, the media had at least a weakly positive

incentive to withhold the news report $s = \hat{a}$ for $\theta \in [\frac{1}{2}, \theta^{SW})$ in two instances, namely when signals are either weak ($\beta \in [\frac{1}{2}, \beta^T)$) or strong ($\beta \in [\beta^{SW}, 1]$). In addition, when $s = \hat{b}$ the media would also have an incentive to withhold publication whenever $\beta \in (\beta^L, 1-\theta]$. Clearly, the presence of the reputation cost reduces the likelihood of the suppression of weak signals when $s = \hat{a}$ since this renders Firm A's advertising fee less attractive relative to Firm B's, who in this case would prefer the story to be published. Therefore, if $R_A - L(\alpha) < R_B$ then the news signal $\beta \in [\frac{1}{2}, \beta^T)$ will be published by the media. Similarly, reputation costs also render Firm A's advertising fee less attractive in the $s = \hat{b}$ case where $\beta \in (\beta^L, 1-\theta]$, which in turn reduces the probability that such signals will be withheld.

Matters are a little different when it comes to the strong signal $s = \hat{a}$ case, i.e when $\beta \in [\beta^{SW}, 1]$. This is because of the congruence of both advertisers' incentives, who both wish to withhold the news signal. Therefore, in order for the news report to be published, then the reputation cost must exceed the higher of the two advertising fees. It is straightforward to show that $R_A > R_B$ for $\theta \in [\frac{1}{2}, \theta^{SW})$: therefore the publication of strong news signals $s = \hat{a}$ depends on whether $R_A < L(\alpha)$.

Both situations clearly demonstrate that even in the presence of some form of reputation effect for the media, commercial media bias as described in this paper may still be present, particularly with regards to the withholding of strong news signals. Note that our current general specification for reputation costs may easily be incorporated in a more complicated two-sided market framework as described in Ellman and Germano (2009) whereby advertising yields additional benefits to Firms A and B in terms of raising consumer's willingness to buy their products, and where the effectiveness of this type of advertising is strictly-increasing in the media's reputation (decreasing in L_R). The results would be similar to those described above.

It would also be interesting to combine two of the extensions discussed so far, namely media competition and reputation concerns, given their natural complementarity. If we start with the simple duopolistic media setting, then although $R_A \geq L(\alpha)$ it may still be the case that the strong news signal is published, so long as $R_B < L(\alpha)$. This is because although Firm A would successfully advertise in order to withhold the news signal in one media outlet, Firm B's advertising fee would be rejected by the other media outlet in favour of publication, in order to avoid the negative reputation cost L_R. It is straightforward to observe that when the number of competing media outlets N increases, the minimum threshold for the reputation cost L_R required in order to induce the publication of strong signals $\beta \in [\beta^{SW}, 1]$ falls.

Hence competition among media outlets reduces the likelihood of withheld or suppressed news signals due to the individual reputation concerns of each media outlet. It is interesting to note that in Gentzkow and Shapiro (2006) reputation concern is the main driver behind their form of demand-side media bias where reports are skewed towards the consumers' prior: in this paper, such concerns, in tandem with increased media competition, reduce the likelihood of advertiser-

driven media bias.

2.5 Concluding Remarks

This paper has sought to analyse the advertiser-media relationship within the context of conflicting incentives to publish/withhold a news story. Using a simple model of horizontal and vertical product differentiation in a duopolistic setting, we showed that in equilibrium when news signals conform to people's prior beliefs, only intermediate or moderate news signals are published in the media; the more extreme or 'strong' news stories are withheld from public consumption. This result is brought about because more extreme news stories result in a significant shift in consumer demand across the duopoly, meaning that the firm who suffers as a result of publication would have an incentive to switch over and eliminate the vertical product differentiation element, thereby competing on a more direct basis with the other firm (for sufficiently-low levels of switching costs). Hence, the beneficiary firm would experience a decline in its profits due to this increased competition, thus providing an incentive to ensure that the extreme news story is withheld via advertising.

This result provides an alternative insight into the various empirical and anecdotal evidence related to the under-reporting of news stories within the media (e.g. Boykoff and Boykoff, 2004) in relation to numerous issues like anthropogenic climate change and the US beef industry's reliance on additives like pink slime. The paper also shows that this type of media bias is eroded as competition between media firms increases and concerns for reputation are introduced in the model.

Therefore, under certain conditions this paper predicts that media bias due to advertiser influence will indeed arise, which can have an impact on public perceptions regarding important issues like climate change and the nutritional value of food. An interesting outcome of this paper is that, unlike previous efforts in the literature (e.g. Ellman and Germano, 2009) the policy prescriptions in terms of regulating the influence and/or role of advertisers are somewhat more ambiguous since in this simple model media bias (partially) arises in a more subtle manner due to the beneficiary firm's incentive to avoid increased competition.

One clear policy conclusion that emerges from this paper is that effective IPR and patenting protection may help to eliminate this form of media bias by raising the fixed costs associated with switching over, thereby ensuring that regardless of the strength of the news story the beneficiary firm would still continue to enjoy a dominant role in the market. Furthermore, apart from encouraging increased investment in research and development, this paper also implies that stringent IPR also promotes the diffusion of new ideas and knowledge via a wide variety of media (not simply newspapers/TV/radio/etc.) since innovators would know that their creations or discoveries would be protected from illicit replication by third parties.

Evidently, such considerations must be weighed against the relative merits

of encouraging increased competition among firms, since this paper has nothing to say in this regard; for example, increased competition may lead to increased investment in research and development to improve the product and re-establish vertical product differentiation.

Another interesting point that emerges from this paper concerns the role played by media competition in reducing the likelihood of advertiser-driven media bias. This outcome is broadly in line with a number of other studies related to media competition and media bias. For example, Kerkhof and Muenster (2015) report that increased competition increases the welfare gains from a cap on advertising introduced by the government, thereby reducing the extent of commercial media bias by enhancing the alignment between media content and viewer preferences.

Historically, the U.S. Federal Communications Commission (FCC) has strongly advocated the proliferation of competition in media markets, with regulations in place barring a single media entity from reaching more than 35% of U.S. households by TV. However, in recent years there has been a notable shift in FCC policy towards deregulation of media ownership laws and media mergers. Among these measures, in 2003 the FCC relaxed this limit to 45%, and in 2007 voted to eliminate the provision that previously forbade a single company from owning both a TV station and newspaper or radio station in the same city. Predictably, many of these policies have been met with stern criticism, since according to several authors (e.g.Bagdikian, 2000) they have led to increased concentration of media firms in the hands of a few large conglomerates, at the expense of media competition.

2.6 Appendices

Appendix A

Proof of Proposition 1

We begin by solving $\Delta_{A,B}$ for ψ and ω. It is straightforward to see that $\Delta_{A,B} \geq 0$ if and only if:

$$\psi \begin{cases} \geq \omega(s = \hat{a}) \\ \leq 1 - \omega(s = \hat{b}). \end{cases}$$

The rationale behind this result is as follows. When $s = \hat{a}$, if $\psi \geq \omega$ then Firm A has a clear incentive to ensure that the news story gets published, and will hence advertise towards this end. From the above expression, this willingness-to-pay on A's part (R_A) will exceed Firm B's advertising fee R_B, meaning that $\Delta_{A,B} \geq 0$ and the media would publish the news story. This occurs because when $\theta > \frac{1}{2}$, demand and profits are already skewed in Firm A's favour due to the magnitude of consumers' prior, which combined with the convexity of $\Delta_{A,B}$ in both posterior beliefs ψ and ω, results in Firm A's advertising fee exceeding

that offered by Firm B.

When $s = \hat{b}$, this means that the news signal is contrary to consumers' prior ($\theta > \frac{1}{2}$), which given the assumptions we made earlier regarding the strength of β would be sufficient to completely alter consumer beliefs and move them in Firm B's direction ($\psi \leq \frac{1}{2}$). In this scenario, due to convexity of $\Delta_{A,B}$, Firm A stands to lose more from the publication of $s = \hat{b}$ than Firm B gains, resulting in the withholding of the news story. This persists up to the point where the strength of the news signal β is such that $\psi \geq 1 - \omega$ (i.e. for relatively 'weak' news signals $s = \hat{b}$). Beyond this point (for 'stronger' signals $s = \hat{b}$) Firm B would now have more of an incentive to ensure publication than Firm A's willingness to withhold.

Hence, the decision to publish or not hinges on the values of ψ and ω. We now consider consumer's updated beliefs when no news report is published (ω). If no news report is published (as always, for $\theta > \frac{1}{2}$), then consumers know that this situation would arise if either no news report was received by the media (with probability $1 - \alpha$) or else if a news report were purposely withheld by the media (with probability α). It is important to note that ex-ante, consumers have no idea whether this allegedly-withheld news report is $s = \hat{a}$ or $s = \hat{b}$. In the former case, the willful withholding of news would only make sense if somehow Firm A believes that the posterior belief under no publication (ω) exceeds the posterior belief following publication (ψ). Since this implies that $\beta \in [\frac{1}{2}, 1]$, then by induction it would be optimal for consumers to pick $\tau = \frac{1}{2}$ as their imputed value of β whenever they believe that $s = \hat{a}$ is being purposely withheld, since this ensures that $\psi \geq \omega$, thus preventing the deliberate concealment of such signals.

On the other hand, whenever a signal $s = \hat{b}$ is being purposely withheld, consumers know that this may occur for any news story with signal strength β such that $\psi \geq 1 - \omega$. Therefore, when $\theta > \frac{1}{2}$ and no news report has been published, consumers' posterior belief can be characterised as follows

$$\begin{gathered} \omega = (1-\alpha)\theta + \alpha\{\theta^2 + (1-\theta)(\tfrac{\frac{3}{2}-\omega}{2})\} \\ \rightarrow \omega_{\theta>\frac{1}{2}} = \tfrac{2}{2+\alpha-\theta\alpha}(\theta + \alpha\theta^2 + \tfrac{3}{4}\alpha - \tfrac{7}{4}\alpha\theta) \ . \end{gathered}$$

All that remains is to check that the optimal advertising strategies specified earlier for each firm are consistent with such beliefs. When $s = \hat{a}$ this entails checking whether the posterior belief under publication (ψ) exceed that under no publication (ω). From the above expression it is evident that this holds true for any value of $\theta > \frac{1}{2}$ given that $\omega_{\theta>\frac{1}{2}} \leq \theta \forall \theta \in (\frac{1}{2}, 1]$, while for any $s = \hat{a}$, $\psi \geq \theta$. Similarly when $s = \hat{b}$ we know that since (by assumption) $\beta \leq (1-\theta)$ then $\psi \leq \frac{1}{2}$, whereas from the expression above $\omega_{\theta>\frac{1}{2}} \geq \frac{1}{2}$ whenever $\theta \geq \frac{1}{2}$.

Therefore, whenever the strength of the news signal β is such that $\psi > 1 - \omega$, the media will withhold the news signal provided that the consumers' prior is contrary to the received signal s.We can now derive the cut-off value of β such that $\psi = 1 - \omega$, where

$$\beta^L = \tfrac{(1-\theta)(1-\omega_{\theta>\frac{1}{2}})}{(1-\theta)-\omega_{\theta>\frac{1}{2}}(1-2\theta)} \ .$$

When $\beta < \beta^L$, the posterior belief under publication are high enough to raise Firm B's incentive to publish the story sufficiently to ensure that its advertising fee exceeds that offered by Firm A, resulting in the publication of the story.

Proof of Lemma 1

We can derive the following condition for ψ which satisfies $\Sigma_{A,B} \geq 0$:

$$\psi^{SW}_{s=\hat{a}} = \frac{4-3\sqrt{1-F^{SW}}}{2} \text{ (for Firm B switching over when } s=\hat{a})$$
$$\psi^{SW}_{s=\hat{b}} = \frac{-2+3\sqrt{1-F^{SW}}}{2} \text{ (for Firm A switching over when } s=\hat{b}) \ ,$$

where it is easy to see that $\psi^{SW}_A = 1 - \psi^{SW}_B$ due to symmetry.

Given that our consumers are assumed to be Bayesian, we can express ψ in terms of their updated beliefs, which enables us to derive the following expression for the cut-off value of β:

$$\beta^{SW} = \frac{(1-\theta)(4-3\sqrt{1-F^{SW}})}{2(2-3\theta)-3\sqrt{1-F^{SW}}(1-2\theta)} \ ,$$

where β^{SW} corresponds to the strength of the media signal $s=\hat{a}$ such that the posterior belief (when the news story is published) is equal to $\psi^{SW}_{s=\hat{a}}$. Similarly, $1-\beta^{SW}$ refers to the point where, if news signal $s=\hat{b}$ is published by the media, consumers' posterior belief would be equal to $\psi^{SW}_{s=\hat{b}}$.

It is easily verifiable that whenever $F^{SW}=0$ then $\Sigma_{A,B} \geq 0$ will always hold true provided that $\psi^{SW}_{s=\hat{a}} \geq \frac{1}{2}$ or, in the case of $s=\hat{b}$, $\psi^{SW}_{s=\hat{b}} \leq \frac{1}{2}$. As seen before, when $s=\hat{a}$ then $\beta \in [\frac{1}{2},1]$ meaning that $\psi^{SW}_{s=\hat{a}} \geq \frac{1}{2}$ will always hold true. Hence, Firm B would *always* have an incentive to switch over whenever a negative news story $s=\hat{a}$ is received by the media. Similarly, when $s=\hat{b}$ then we know that $\beta \in [0,1-\theta]$, ensuring that $\psi^{SW}_{s=\hat{b}} \leq \frac{1}{2}$ will always hold.

For $F^{SW}>0$, the only conditions that must be satisfied are that in the case of $s=\hat{a}$, $\psi^{SW}_{s=\hat{a}} \leq 1$, and in the case of $s=\hat{b}$ that $\psi^{SW}_{s=\hat{b}} \geq 0$. Therefore, setting:

$$\psi^{SW}_{s=\hat{a}} = \frac{4-3\sqrt{1-F^{SW}}}{2} \leq 1 \ \psi^{SW}_{s=\hat{b}} = \frac{-2+3\sqrt{1-F^{SW}}}{2} \geq 0 \ ,$$

we arrive at the stated result, whereby in order for both of the above to hold, $F^{SW} \leq \frac{5}{9}$.

Proof of Proposition 2

Firstly, note that since we are dealing with the case where the prior belief is such that $\theta \in [\frac{1}{2},1]$, then posterior belief under no publication will never go below $\frac{1}{2}$; i.e. $\omega_{\theta \geq \frac{1}{2}} \geq \frac{1}{2}$. Hence in this case we need not consider the possibility of Firm

A switching over, since posterior belief under no publication will never be low enough to induce Firm A to switch over. The reason behind this result is that when no news report is published, consumers' posterior belief will be skewed in the direction of their prior since they are uncertain as to the true realisation of Y as well as the value of β.

Therefore, we can focus solely on $\omega^{SW}_{s=\hat{a}}$. Firstly, note that $\omega_{\theta>\frac{1}{2}}$ is continuous in θ over the interval $[\frac{1}{2},1]$, while the critical cut-off point $\omega^{SW}_{s=\hat{a}} \in [\frac{1}{2},1]$ for $F^{SW} \in [0,\frac{5}{9}]$ is exogenous in θ. Therefore, by continuity it follows that θ^{SW} such that $\omega_{\theta>\frac{1}{2}} = \frac{4-3\sqrt{1-F^{SW}}}{2}$ exists and is unique. Setting $\omega_{\theta>\frac{1}{2}} = \frac{4-3\sqrt{1-F^{SW}}}{2}$, we obtain the following expression for our cut-off value of θ:

$$\theta^{SW}(\alpha, F^{SW}) = \frac{(\frac{7}{4}\alpha-1)+\sqrt{(1-\frac{7}{4}\alpha)^2-4\alpha(\frac{3}{4}\alpha-2+\frac{3\sqrt{1-F^{SW}}}{2})}}{2\alpha} \ .$$

Proof of Proposition 4

As discussed earlier, when $\theta \in [\frac{1}{2},\theta^{SW})$, we can formally specify consumers' posterior belief under no publication as follows:

$$\omega_{\theta\in[\frac{1}{2},\theta^{SW})} = \frac{2}{2+\alpha-\alpha\theta}\{\theta(1-\frac{5}{4}\alpha)+\frac{\alpha\theta^2}{2}+\frac{3}{4}\alpha\} \ .$$

We have already seen that when $s=\hat{a}$, Firm A's advertising fee exceeds that on offer by Firm B given that the consumers' prior belief is already skewed in A's favour. Furthermore, we showed that whenever $\psi \geq \omega$ then it is in Firm A's best interest to ensure that the story is published. The key here is therefore to determine whether this holds true given our specification of posterior belief under no publication $\omega_{\theta\in[\frac{1}{2},\theta^{SW})}$ above, and given that the news signal is weak or moderate ($\beta \in [\frac{1}{2},|\beta^{SW}$). We know that $\psi > \theta$ whenever $\beta > \frac{1}{2}$ for $\theta > \frac{1}{2}$, so in this case it suffices to show whether $\omega_{\theta\in[\frac{1}{2},\theta^{SW})} < \theta$. From the above expression it is straightforward to show that $\omega_{\theta\in[\frac{1}{2},\theta^{SW})} < \theta$ iff $\theta \in (\frac{3}{4},1)$, meaning that for any $\theta \in [\frac{1}{2},\frac{3}{4})$ it is possible that the posterior belief under no publication would actually exceed that under publication (ψ). In this case, Firm A would prefer to suppress the news signal due to this over-updating by consumers.

Therefore, it is necessary to show the conditions under which consumers with prior belief $\theta \in [\frac{1}{2},\frac{3}{4})$ would over-update their beliefs regarding state Y when the news signal is withheld relative to the situation where a news story is indeed published. To do this, we compare posterior beliefs under publication and withholding, as shown below:

$$\frac{\theta\beta}{\theta\beta+(1-\theta)(1-\beta)} \leq \omega_{\theta\in[\frac{1}{2},\theta^{SW})}.$$

Solving for β, we get the following expression:

$$\beta^T \leq \frac{\omega_{\theta\in[\frac{1}{2},\theta^{SW})}(1-\theta)}{\theta-2\theta\omega_{\theta\in[\frac{1}{2},\theta^{SW})}+\omega_{\theta\in[\frac{1}{2},\theta^{SW})}} \ ,$$

which indicates the cut-off point for signal strength β such that when $\beta \in [\frac{1}{2}, \beta^T)$ Firm A will prefer to withhold the news story rather than have it published due to over-updating of the consumers' prior. Thus, the media will withhold such news stories from publication.

Proof of Proposition 5

The proof is somewhat similar to (part of) the one provided in Proposition 1. When $\theta \geq \frac{1}{2}$, we can specify consumers' posterior belief under no publication $\omega_{\theta \geq \frac{1}{2}}$ as before, where:

$$\omega_{\theta \geq \frac{1}{2}} = \frac{2}{2+\alpha-\alpha\theta}\{\theta(1-\frac{5}{4}\alpha) + \frac{\alpha\theta^2}{2} + \frac{3}{4}\alpha\} \ .$$

Firstly, recall that since $\beta \in [0, 1-\theta]$ consumers' posterior belief under publication (ψ) would always be weakly less than $\frac{1}{2}$ when $\theta \geq \frac{1}{2}$. Secondly, from the above expression it is also easy to show that $\omega_{\theta \geq \frac{1}{2}} \geq \frac{1}{2} \forall \theta \in [\frac{1}{2}, 1]$. Hence whenever $s = \hat{b}$ and $\theta \geq \frac{1}{2}$, Firm B (A) would have an incentive to advertise in order to ensure (no) publication. By Proposition 1, we know that when $\theta \geq \frac{1}{2}$ Firm A's advertising fee is at least weakly greater than B's when the strength of the news signal β is such that $\psi \geq 1 - \omega$. As before, we denote this cut-off value of β as β^L, where:

$$\beta^L = \frac{(1-\theta)(1-\omega_{\theta \geq \frac{1}{2}})}{(1-\theta)-\omega_{\theta \geq \frac{1}{2}}(1-2\theta)}$$

with the only difference being that now our expression for $\omega_{\theta \geq \frac{1}{2}}$ is the one derived above, where firm switch-over is allowed.

Appendix B - Details regarding Number of Climate Change Articles in U.S. Newspapers, 2004-2010

Data on the annual number of published newspaper articles related to global warming issues (over the period 2004-2010) was collated using two main resources - NewsBank's NewsLibrary website, which is a searchable online repository of newspaper articles in the U.S. covering over 4,000 newspapers across the country, and ProQuest[17]. For the purposes of searching for relevant climate change news stories in each newspaper, an appropriate Boolean string-search protocol was employed, whereby articles containing the words "global warming", "climate change" or "greenhouse gas" in either the headline or lead paragraph were sought out, excluding those articles containing the words "climategate", "skeptic", "hoax",

[17]These sources have been widely-used in numerous other studies within the literature, like for example Gentzkow and Shapiro (2010).

"global cooling", "myth", "denier", "denial", "conspiracy" or "swindle" in the headline or lead paragraph, for each year over the period January 1, 2004 to December 31, 2010. The search protocol used can be justified on the basis of the following points:

- Restricting the search terms "global warming", "climate change" and "greenhouse gas" to just the headline or lead paragraph automatically eliminated several news stories which had no link to global warming issues;
- The adopted approach also minimised the number of climate-related news stories appearing as part of a 'News in Brief'-type section, where the actual story would only form a small part of the overall content. Furthermore, this also helped to minimise those articles which simply alluded to global warming briefly as part of some turn-of-phrase within the scope of an altogether distinct story;
- The rather exhaustive list of exclusion terms was used in order to eliminate any overtly climate-skeptic articles which question the scientific validity of global warming, since the aim of this paper is to gauge the level of coverage afforded to the scientific facts or developments on global warming over the period under review. This list of terms was compiled on the basis of content analysis of climate-skeptic news articles which showed up in our sample when the search protocol omitted the exclusion terms listed above.

To further ensure reliability and consistency, the search results obtained for each newspaper were then individually perused for any invalid articles which should have been excluded but which somehow evaded the search protocol, with the main focus being on climate-skeptic news stories. In reality the stringency of the string-search protocol employed meant that very few articles were excluded on this basis.

Figure 2.3 shows the top 50 newspapers on the basis of the number of climate-related news stories published in the period under review. From the diagram it is clear that the largest newspapers in the country dominate the list, with the Washington Post leading the way with 1,993 articles, followed by the New York Times (1,687 articles) and the Boston Globe (1,683 articles).

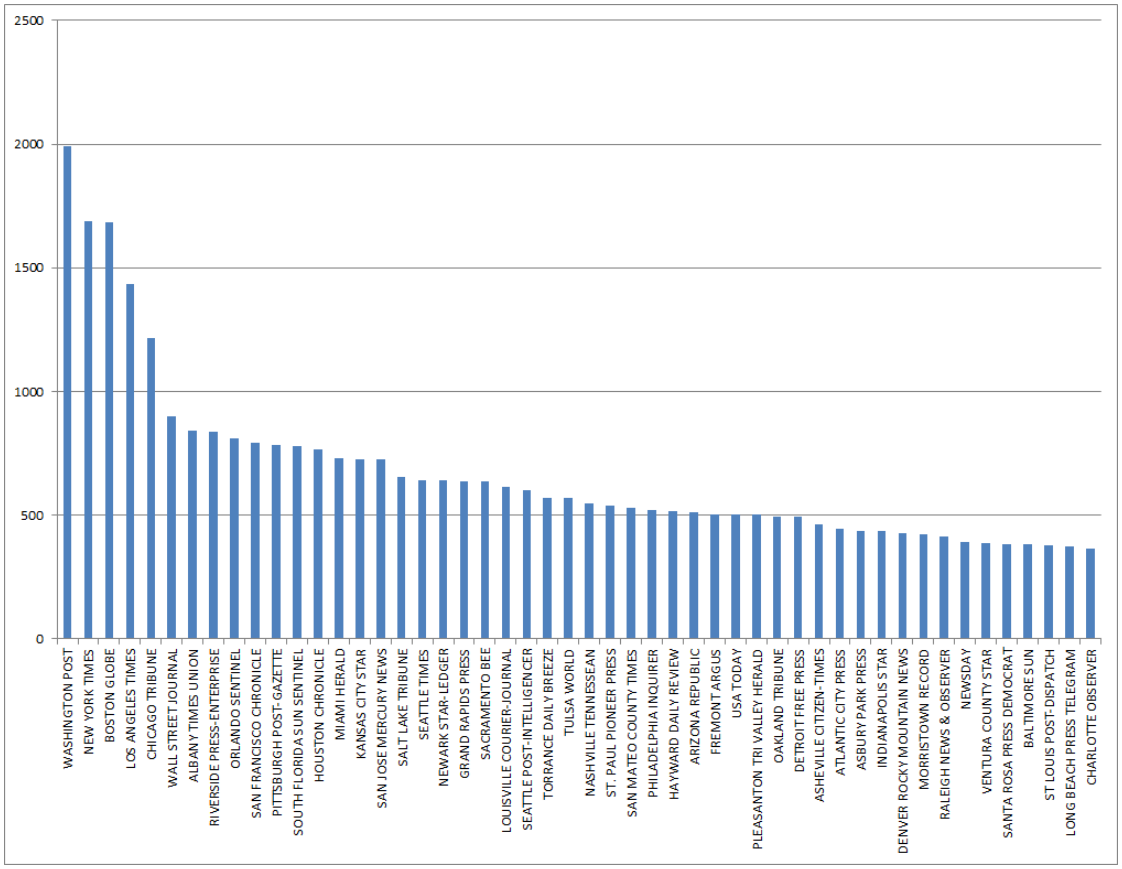

Figure 2.3: Top 50 U.S. Daily Newspapers by Number of Climate-Related News Stories Published from January 1, 2004 to December 31, 2010 (Source: NewsLibrary, ProQuest, 2014)

Chapter 3

(Non-)Random Acts of Kindness: Media Coverage of the 2004 Indian Ocean Tsunami and Charitable Donations in the U.S.

3.1 Introduction

The rise of the news media as the Fourth Estate has raised several questions regarding both its reporting activities as well as its influence on people's beliefs. The longstanding view of the media as a neutral supplier of information to the masses has regularly been challenged, with mounting evidence of slanted or biased reporting (e.g. Groseclose and Milyo, 2005). And yet none of this debate would have any importance to society unless the media exerted any sort of influence over its audience's decisions, something that has been shown in various contexts ranging from women's attitudes towards their ideal body shape (Park, 2005) to the level of homicides in the U.S. (Phillips, 1983). In this paper the primary focus is on how the level of media reportage affects individual economic behaviour, abstracting away from all considerations related to bias. More specifically, we study the causal link between media coverage of the 2004 Indian Ocean tsunami in the U.S. on American people's donations towards the tsunami relief effort, both in terms of the decision to donate or not as well as the magnitude of these disbursements.

In order to identify the underlying causal link between newspaper articles and charitable donations, we use county-level *growth rates* in violent crime for the year 2005 as an instrument for tsunami coverage. The basic idea behind this instrument is that a spike in violent crime in 2005 (relative to the previous year, or even the average of the previous 5 years) would have generated a surge in media coverage of such stories, at the expense of tsunami-related articles. The results show that the number of tsunami-related news stories published in U.S. newspapers had a statistically significant effect on both the probability and size of private donations. The estimates indicate that media coverage only had a

modest impact on the average U.S. citizen's final decision to donate money or not. By contrast, the results imply a non-trivial relationship between media coverage and donation levels since they suggest that if the average daily newspaper had published one additional tsunami-related news story per week over the 6-month period immediately after the disaster, this would have boosted average donations by around 7%, translating to approximately US$131 million (based on estimates for total private donations towards the tsunami relief effort - see Becker and Strom, 2005).

In addition we also provide evidence that the media had a larger impact (in terms of encouraging higher donations) on individuals who had previously undertaken some form of voluntary work in 2004 as well as people in the 25-34 age bracket. This result is of particular interest for media advertisers since it implies that young adults are more responsive to information derived from the media relative to other age groups and controlling for a whole host of respondent-specific characteristics.

The selection of the Indian Ocean tsunami was motivated both by the scale of the disaster and the extensive media coverage afforded, as well as the fact that 2005 Center of Philanthropy Panel Study (CPPS), which forms the basis of this paper, includes a special section on charitable donations related to the tsunami. We combine this with data on newspaper reportage of the tsunami in the six months after the disaster (covering 386 newspapers in counties from all across the 50 states and Washington DC). The results underscore the heightened importance of the media within today's society, as well as its influence over people's economic choices. The results may also be important for the design of future charitable fund raising campaigns as well as media advertising strategies, by enabling the careful targeting of specific population subgroups who would be more likely to respond to any information received via the media.

This paper forms part of the growing economics literature on the behaviour and impact of the media. For the most part, the main focus of the literature so far has been on media bias and how such reportage may influence political outcomes. For example DellaVigna and Kaplan (2007) find that the introduction of Fox News in the U.S. led to a significant increase in the proportion of Republican votes during both the Presidential and Senate elections, while Chiang and Knight (2011) look at how newspaper endorsements of political candidates in the U.S. affect voting behaviour. On the other hand, the non-political impact of the media has been afforded relatively scant attention in the literature, which is where this paper fits in. For example Dyck, Volchova, and Zingales (2008) look at how increased media coverage of corporate governance violations perpetuated by public limited companies in Russia can induce these companies to change their policies or attempt to redress the situation in order to mitigate the potential fallout with shareholders and consumers.

A closely-related paper to ours is by Eisensee and Stromberg (2007), who look at the U.S. government's relief efforts in response to various natural disasters over the period 1968 to 2002, concluding that such decisions are mainly driven by media coverage of these disasters. By contrast, this paper adds to the

existing literature by focusing on the impact of media reportage on private individuals' economic behaviour rather than voting/political choices or government interventions. Given the pervasiveness of the modern-day media, particularly via new channels like social media and online portals, such analysis may have significant repercussions in terms of consumer sentiment/preferences as well as the level of economic activity within a country (Doms and Morin, 2004). In addition, this paper also differs from the existing literature in that it adopts an instrumental variables strategy in order to determine the causal link between media reports and individual behaviour.

The paper also represents one of the first attempts at analysing the influence of media reportage on people's choices within the context of charitable gift-giving. The literature on voluntary contributions is relatively vast, and is mainly concerned with the theory of public goods and contribution design mechanisms (see for example Andreoni, 1990). Nonetheless, a number of empirical studies have focussed on identifying the main determinants of charitable gift-giving in the economics literature, including income and age (Kitchen, 1992), tax incentives (Randolph, 1995), religious beliefs (Andreoni, Payne, et al., 2011), as well as education levels and prior donations towards charitable causes (Brown, Harris, and Taylor, 2012). By contrast, despite the wealth of anecdotal evidence, the role of the media as a potential determinant for charitable gift-giving has been afforded quasi-negligible attention in the literature. Lobb, Mock, and Hutchinson (2012) analyse the relationship between coverage of the 2010 Haiti earthquake in traditional and social media and total donations towards the subsequent relief effort in the U.S., with results showing a positive and statistically significant correlation between the two. As highlighted earlier, this paper is concerned with identifying the causal impact of media coverage on both the likelihood and level of individual charitable disbursements.

3.1.1 The 2004 Indian Ocean Tsunami and Relief Effort

On December 26, 2004 an earthquake of magnitude 9-9.3 struck the Indian Ocean seabed, some 100 miles off the coast of Sumatra in Indonesia. The earthquake, which is officially the third largest recorded in human history, led to massive wave formations in the ocean, with some waves reaching 98 feet and traveling at significant speeds. The aftermath of the earthquake and subsequent waves (or tsunami) was devastating - a total of 185,000 people were confirmed to have died, although it is estimated that the true death toll was in reality closer to 250,000.[1]. Fatalities were recorded in no fewer than 14 countries across South East Asia and Africa, with the worst-hit being Indonesia, Sri Lanka, India, Thailand and the Maldives, while a significant number of foreign nationals perished due to the popularity of these countries as tourist destinations during the winter period. In fact, a total of 2,307 foreigners were reported as dead or missing in the aftermath of the tsunami from 47 different countries, including some 571 people from Sweden

[1]According to the US Geological Survey the earthquake vibrated the entire world by approximately 1cm, and created an 800 mile fault-line in the Earth's seabed.

and around 552 from Germany. In addition, the physical damages caused by the disaster in terms of destroyed homes, schools and infrastructure systems were catastrophic; it is estimated that over 1.6 million people were displaced as a result of the tsunami[2].

In the wake of the devastation caused by this disaster, an unprecedented international relief effort was launched by the UN in tandem with several countries and aid organisations. As seen in Figure 3.1 the U.S. was by far the largest donor with an estimated US$2.9 billion in aid, followed by Australia (US$1.3 billion - the highest amount relative to GDP) and Germany (US$1.1 billion). Note that the figures shown reflect amounts pledged by each country rather than actual disbursements - an important distinction since a number of countries were reportedly lagging behind in terms of honoring their pledged donations (UN, 2005). The U.S. relief effort, which was spearheaded by former Presidents George H.W. Bush and Bill Clinton, consisted of around US$350 million in donations by the Federal government (together with the deployment of military assistance to coordinate the distribution of resources), with the rest of the money coming from the private sector, boosted by the government's announcement that any private donations would be eligible for tax deductions[3].

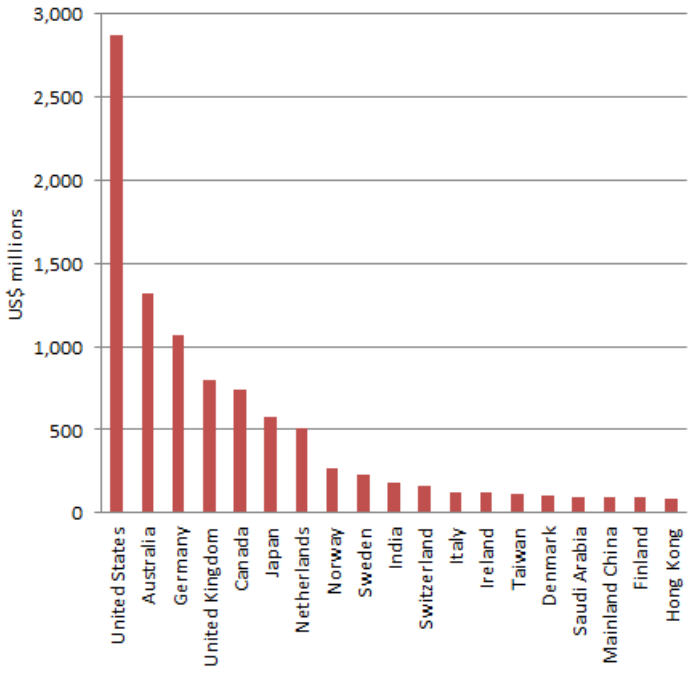

Figure 3.1: Top 20 Donor Countries for the Tsunami Relief Effort (Source: United Nations Development Program, 2005)

[2]The financial cost of such damages in the Maldives alone amounted to a reported US$470 million, or 62% of the country's GDP, according to the UN.

[3]The general response to the U.S.'s relief efforts by the international community was somewhat mixed, with the UN accusing it (and other Western countries) of being 'stingy', while on the other hand a number of organisations like the US Burroughs Wellcome Fund argued that the high levels of private donations (fueled by the aforementioned tax deductions) for tsunami victims was coming at the expense of other causes like African malaria aid.

3.2 Data

The data on charitable donations towards the tsunami relief effort were obtained from the 2005 Center of Philanthropy Panel Study (CPPS), which is a wide-ranging bi-annual national survey on charitable gift-giving conducted by Indiana University-Purdue University Indianapolis's Center of Philanthropy in conjunction with the University of Michigan Institute for Social Research's Panel Study of Income Dynamics (PSID). A unique feature of the CPPS is that it has been following the same households across the U.S. since 1968, with children and siblings added on over time; in fact the 2005 survey included a total of 8,041 households (this figure was reduced to 7,215 observations after accounting for data omissions and other shortcomings). The CPPS includes extensive data on charitable donations and volunteer work undertaken by each household in the preceding year across all categories of charities, including religious, environmental, international aid, medical, arts and culture and educational causes. The 2005 CPPS also included a separate section on charitable donations towards the 2004 tsunami relief effort; data on the dollar amount donated for such purposes were also included, and this forms the basis of the empirical work conducted in this paper. In addition, the CPPS also includes a detailed record of the households' socioeconomic and demographic characteristics, like educational background, age, and gender, and these are also used in the subsequent analysis as control variables. Figure 3.2 shows the average donation per household towards the tsunami relief effort in each of the 50+1 States according to the 2005 CPPS data. A cursory glance at the diagram shows that Delaware recorded the highest donation per household of US$150, followed by Kansas and Montana.

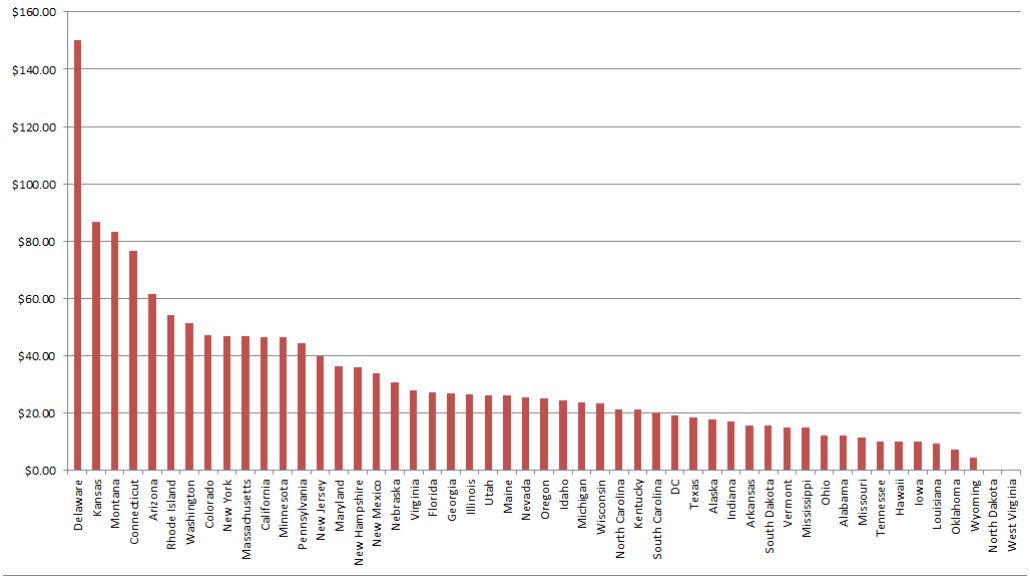

Figure 3.2: Average Tsunami Donations per Household by State (Source: CPPS, 2005)

Data on the number of published newspaper articles related to the tsunami in the 6-month period immediately after the disaster (from December 26, 2004 to

June 26, 2005) was collated using two main resources - NewsBank's NewsLibrary website, which is a searchable online repository of newspaper articles in the U.S. covering over 4,000 newspapers across the country, as well as the individual newspapers' print archives.[4] This dual approach was employed in order to maximise the reliability and accuracy of the data; in case of a discrepancy between the two sources, the source with the highest number of verifiable articles was selected. In practice this was rarely the case since: (i) a large proportion of newspapers use NewsLibrary as their official print archive database, and (ii) in most other cases (with a few notable exceptions, like all Gannett-owned newspapers) the individual newspapers' archives were either incomplete or inaccessible. This meant that, effectively, NewsLibrary was the main source of data for news articles in this study. For the purposes of searching for relevant tsunami news stories in each newspaper, an appropriate Boolean string-search protocol was employed, whereby articles containing the words "tsunami" or "earthquake" in either the headline or lead paragraph were sought out, excluding those articles containing the words "letter" or "letters", over the period December 26, 2004 to June 26, 2005. The search protocol used can be justified on the basis of the following points:

- Restricting the search terms "tsunami" and "earthquake" to just the headline or lead paragraph automatically eliminated several news stories which had no link to the events of the 2004 Indian Ocean tsunami (in fact many sports-related articles routinely use these words allegorically to emphasise a victory or defeat);

- The adopted approach also minimised the number of tsunami-related news stories appearing as part of a 'News in Brief'-type section, where the actual story would only form a small part of the overall content. Furthermore, this also helped to minimise those articles which simply alluded to the disaster briefly as part of some turn-of-phrase within the scope of an altogether distinct story (this was particularly valid in the case of editorials and opinion pieces, which regularly used the disaster to emphasise some point, like for example the then-President's perceived shortcomings);

- The exclusion of articles containing the words "letter" or "letters" was done to eliminate all letters to the editor, since it was felt that despite their ubiquity they do not constitute news stories within the context of media coverage as defined in this study.

To further ensure reliability and consistency, the search results obtained for each newspaper were then individually perused for any invalid articles which should have been excluded (such as news in brief or letters to the editor) but which somehow evaded the search protocol. This entire process was conducted for a total of 386 daily newspapers covering all 50+1 States, including national, state-wide and local newspapers. Figure 3.3 shows the top 50 newspapers on the

[4]NewsLibrary has been used as the principal source of news-related data in various other studies, such as Gentzkow and Shapiro (2010).

basis of the number of tsunami-related news stories published in the six month period immediately after the disaster.

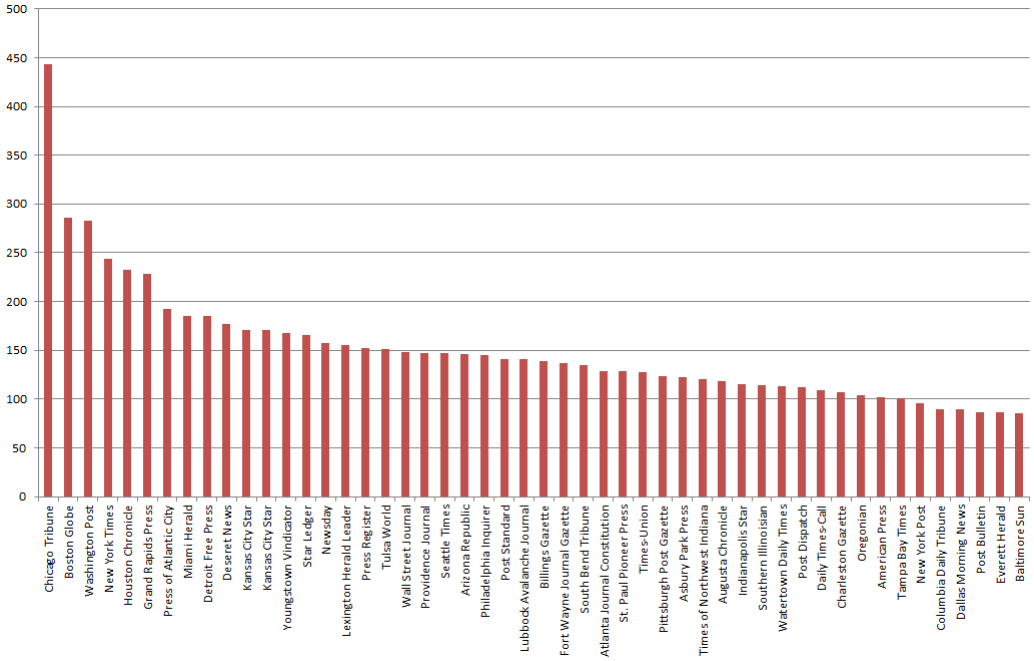

Figure 3.3: Top 50 US Daily Newspapers by Number of Tsunami-Related News Stories Published from December 26, 2004 to June 26, 2005 (Source: NewsLibrary, newspaper print archives, 2013)

The choice of newspapers was largely dictated by the demographic characteristics of the survey respondents in the 2005 CPPS survey, which records both the State and the Beale code associated to each respondent (city/county codes are omitted from the final survey due to data protection issues, meaning that for the purposes of this study it was necessary to focus on newspapers pertaining to each respondents' reported Beale code). The Beale coding system is a county-level geographical classification designed by the United States Department of Agriculture (USDA) used to identify the size and degree of urbanisation of the county in question. The code ranges from 1 to 9, with 1 denoting a central county within a metropolitan area with a population of over 1 million people, while 9 denotes a completely rural county with no adjacent metropolitan area nearby and a population of under 20,000 inhabitants. For each State, it is possible to identify which counties correspond to which Beale code using the USDA's Rural-Urban Continuum Codes (for the purposes of this study, the 2003 Continuum Codes were used) - so for example Los Angeles County pertains to Beale code 1 in California, whereas Woodruff County pertains to Beale code 9 in Arkansas (in general each Beale code contains around 5-8 individual counties per State). Hence for every single Beale code entry in the 2005 CPSS, each relevant county was listed together with (wherever possible) its corresponding highest-circulating daily newspaper using data compiled by the Alliance for Audited Media; this search process was conducted for each of these newspapers[5].

[5]In some cases (for example New York State) more than one top newspaper was listed per

To obtain the final figure for the number of news articles per Beale code (for each State) a weighted-average was taken based on the relative population of each county within each code. This was done since, in the absence of any data related to each respondent's newspaper of choice or city/county of origin, obtaining a weighted average for the number of tsunami-related articles per Beale code (per State) represented the most straightforward manner through which to estimate the level of tsunami coverage that each respondent was exposed to over the period under review. In order to ensure robustness, a second weighted average was also calculated, this time on the basis of the newspaper's daily circulation figures. Both weighted averages are used in the subsequent empirical analysis; in either case the results obtained are almost identical. An example will help to clarify matters. Consider the State of New Mexico, where Dona Ana (population of 174,682 inhabitants based on 2003 data), San Juan (113,801 inhabitants) and Santa Fe (129,292 inhabitants) counties are all classified as belonging to Beale code 3, defined as a "county in a metropolitan area of fewer than 250,000 population" (USDA, 2003). The top-circulating daily newspaper in each of these counties is the Las Cruches Sun-News (average daily circulation of 19,641 copies), the Farmington Daily Times (circulation of 113,801) and the Santa Fe New Mexican (22,000) respectively, and the number of tsunami-related news stories published in each of these newspapers over the period December 26, 2004 to June 26, 2005 was 6, 6 and 19 respectively. Hence if we consider the weighted average on the basis of population, the average number of tsunami-related articles pertaining to Beale Code 3 in New Mexico is 10 articles, whilst if the circulation-based weighted average is employed this figure rises to 11.

Table 3.1 presents a brief summary of the main descriptive statistics for each of the key variables described above. The mean donation to the tsunami relief effort of US$29.08 is somewhat misleading since this also includes a significant number of zero donations; in fact from the total sample of 7,215 respondents only 1,726 actually donated money (among those who did donate money, the mean donation was of US$121.55). The mean number of tsunami-related articles is 64.71 or 69.64 depending on the weighting method employed, which shows the relative congruence of the two approaches. Due to limited space, in subsequent sections we shall present the estimates derived using population-weighted tsunami articles as our explanatory variable of interest. Results using circulation-weighted tsunami articles have also been computed, and are quantitatively similar to those shown in the following sections (these are available on request).

3.3 Empirical Analysis and Results

In this section we describe the econometric strategy adopted in this paper. As alluded to earlier, we are interested in analysing both the impact of tsunami coverage on the individual's decision to donate money or not, as well as on the actual amount donated. We start with the first question. Let DY_i refer to a binary variable indicating whether respondent i donated any money or not as observed

county due to the size of the newspapers and/or the unreliability of circulation figures.

Table 3.1: Summary Statistics

Variable	Mean	Std. Dev.	Min.	Max.
Amount donated to tsunami (US$)	29.07789	142.2636	0	6000
No. of articles (by population)	64.70926	54.75623	1	283
No. of articles (by circulation)	69.64075	58.11173	1	283
N		7,215		

Notes: This table shows summary statistics for the two main variables of interest in this paper. Data on the amount donated towards the tsunami relief effort are from the 2005 Center of Philanthropy Panel Study (CPPS), whilst the number of tsunami-related articles were collated via the NewsLibrary website and newspaper print archives. The two different measures of tsunami articles relates to the fact that the CPPS codes each respondent's home county according to the level of urbanisation as specified in the USDA's Beale code system, thereby incorporating more than 1 county. Therefore for each Beale code we identified the top newspapers for each county and obtained the number of tsunami articles for each newspaper. To calculate the associated tsunami coverage for each Beale code we used two different weighting measures - one on the basis of the population size of each newspaper's home county, and the other based on each newspaper's average weekly circulation in 2005.

in our data (where $DY_i \in \{0, 1\}$), while DY_i^* is the corresponding latent variable based on the unobserved propensity to donate money towards the tsunami relief effort. We can therefore express the individual's decision to donate money or not in terms of the following latent variable formulation:

$$DY_i^* = \beta A_i + \mathbf{X}_i'\alpha + \varepsilon_i$$
$$DY_i = \begin{cases} 1 & \text{if } Y_i^* > 0 \\ 0 & \text{if } Y_i^* \leq 0 \end{cases}$$

where A_i is the number of tsunami-related news articles published by the top newspaper/s in respondent i's county over the period December 26, 2004 to June 26, 2005, $\mathbf{X}_i$ is a vector of covariates[6] and ε_i is a random error term. Under the assumption that ε_i is normally and independently distributed with zero mean and a constant variance, the above can be estimated consistently using a Probit model.[7]

We next turn to modelling the decision regarding the level of charitable disbursements, which is somewhat complicated by the fact that our donations data contains a significant number of zeros. Let Y_i denote the amount (in U.S.$) of money donated to the tsunami relief effort by respondent i as observed in our data, where $Y \geq 0$, while Y_i^* represents the corresponding (unobserved) latent variable for tsunami donations. We can express the amount of money that individual i decided to donate towards the relief effort by the following specification:

$$Y_i^* = \beta A_i + \mathbf{X}_i'\alpha + \upsilon_i$$
$$Y_i = \begin{cases} Y_i^* & \text{if } Y_i^* > 0 \\ 0 & \text{if } Y_i^* \leq 0 \end{cases}$$

[6] A full description of the control variables used, including summary statistics, is provided in Appendices A1 and A2.

[7] Logit estimates were also produced, yielding very similar results; these have been omitted due to space limitations.

where υ_i is a random error term. Under the assumptions that υ_i is homoskedastic and normally-distributed with mean zero, then applying a Tobit model to the above relationship would yield consistent estimates for our slope coefficients.[8]

The reduced-form equations which form the analytical core of this paper are thus specified as follows, where (3.1) is estimated using a Probit model while (3.2) is estimated using a Tobit model:

$$DY_i = \psi + \gamma A_i + \mathbf{X}_i'\delta + \upsilon_i \tag{3.1}$$

$$Y_i = \mu + \beta A_i + \mathbf{X}_i'\alpha + \varepsilon_i \tag{3.2}$$

Table 3.2 reports the initial results for Equations 3.1 and 3.2 respectively. Note that standard errors are clustered according to each respondent's Beale code within each state, corresponding to a total of 459 clusters. This specification is used throughout the rest of the paper. As seen from the results, media coverage is positively and significantly-related to both the likelihood of donating money (column 1) as well as the magnitude of charitable disbursements (column 2). It is interesting to note some of the other variables which manifest a positive and statistically significant relationship in both equations, namely whether the respondent made any charitable donations (above \$25) in the previous year, higher education, marital status, age and the respondent's participation in voluntary work in 2004.

Nonetheless, the estimates obtained must be treated with some degree of caution due to a number of econometric issues. The first problem relates to the potential reverse-causality that may exist between our dependent variables and the explanatory variable of interest, namely media coverage of the tsunami. It is entirely plausible to expect that although increased media coverage may influence the likelihood and extent of charitable gift-giving, the reverse relationship may also hold true. This is because if the newspaper realises that donations among its core readership are high (or that readers have a strong propensity towards such gift-giving), then it is likely that it will increase its coverage of the tsunami due to readers' apparent interest in the subject, in the hope that this will result in higher circulation figures (as alluded to in Gentzkow and Shapiro, 2010). This issue is further exacerbated by the fact that certain news stories detail fund-raising activities undertaken within the communities served by the newspaper in question, and which although may compel readers to donate/increase their donations, are primarily driven by the existing extent of charitable giving.

[8]In this case it would not be necessary to implement a double hurdle model as proposed by Cragg (1971) since as we shall see later on the key explanatory variable of interest, namely newspaper coverage of the tsunami, exhibits a statistically significant relationship with both the likelihood as well as the level of donations. In fact, this is also observable for most of the other explanatory variables used in both equations. This means that the factors related the decision to donate money or not are very similar to those related to the amount of money donated, which is implicitly-assumed in a Tobit specification.

In addition, it is likely that our estimates are further compromised by potential measurement errors which may bias our results downwards. This is because although we have taken various steps to maximize the validity of our tsunami coverage variable in terms of the content of each story, no information is available regarding the actual position of the individual news articles within the newspaper. The importance of a news article's position in a newspaper has been looked at by various authors (e.g. Bogart, 1989), both in terms of page number as well its relative position on the actual page, and the impact of such positioning on the likelihood that the article is read. This means that our tsunami articles variable may contain several stories which could have been overlooked due to their unfavorable position within the newspaper, thus limiting their influence on charitable giving within this context.

Table 3.2: Basic Regression Results

	(1) Donation > 0		(2) Amount donated to tsunami	
Tsunami Articles	0.000902***	(0.000372)	0.231**	(0.111)
More than $25	0.721***	(0.0328)	207.6***	(13.67)
Total family income	0.000000512***	(0.000000159)	0.000296***	(0.0000989)
Wealth	1.46e-08	(1.87e-08)	0.0000208***	(0.00000502)
Voluntary work	0.309***	(0.0404)	88.60***	(6.542)
High School	0.0779	(0.0644)	11.10	(20.91)
Undergrad	0.186**	(0.0918)	45.81*	(25.21)
Postgrad	0.254***	(0.0598)	77.69***	(23.06)
Catholic	0.0200	(0.0436)	0.137	(7.518)
Jewish	0.418***	(0.108)	118.1***	(23.85)
Protestant	-0.0239	(0.0587)	6.285	(11.49)
Other	0.161**	(0.0819)	101.8***	(34.09)
Between 25-34	-0.0871	(0.0627)	-30.59	(19.41)
Between 35-44	0.0965**	(0.0451)	28.13***	(10.73)
Between 45-54	0.0320	(0.0574)	13.98	(20.21)
Between 55-64	0.160***	(0.0393)	38.28***	(9.912)
Gender head	0.202***	(0.0580)	44.85**	(18.46)
Marital Status	0.139***	(0.0498)	29.25	(19.44)
Working status head	0.217***	(0.0561)	61.62***	(15.12)
Working Wife	0.0244	(0.0338)	15.81	(10.98)
Retired status head	0.177*	(0.0971)	54.14**	(25.89)
African American	-0.0694**	(0.0292)	-6.293	(11.23)
Hispanic	0.0341	(0.0910)	23.15	(21.57)
Northeast	0.304***	(0.0498)	82.68***	(16.96)
South	0.141***	(0.00964)	29.94***	(5.230)
West	0.174***	(0.0469)	70.95***	(15.80)
Constant	-1.982***	(0.102)	-645.9***	(34.80)
Observations	7208		7208	

Notes: Column 1 reports Probit coefficients where the dependent variable is a dummy which takes a value of 1 if the survey respondent donated money towards the 2004 Indian Ocean Tsunami relief effort. Column 2 reports Tobit coefficients where the dependent variable denotes the value in US$ individual donations towards the tsunami relief effort. In both cases the explanatory variable of interest is the number of tsunami-related articles published in the leading newspapers for each survey respondent's home county. All data are from the 2005 Center of Philanthropy Panel Study (CPPS), with the exception of the tsunami articles data which were collated from the NewsLibrary website and individual newspaper archives. Robust standard errors, clustered at the Beale code level, are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

3.3.1 Instrumental Variables Approach

In this section we develop an instrumental variables (IV) framework in order to account for the endogeneity problems cited above. For the purposes of this paper, the identification strategy employed is to exploit variations in county-level growth rates in violent crime[9]. The rationale behind the selection of this instrument is that a newspaper has a limited amount of space in which to publish news stories, with several potential articles vying for publication. Violent crime is one of the most widely-reported news topics in the U.S., and its ubiquity in American newspapers has been studied at length by the likes of Barak (1994) who contends that such stories help to sell a significant number of newspapers which is ultimately the primary aim of any news editor. The immense popularity of crime-related news stories, according to Barak (1994), is firmly rooted in the general public's continued fascination with the salacious and the macabre. In fact, some authors like Williams and Dickinson (2003) contend that in the pursuit of higher circulation figures the media often exaggerates the extent or magnitude of violent crime instances in order to render the news article more attractive to its audience.

It is therefore entirely likely that over the period December 26, 2004 to June 26, 2005 violent crime and tsunami-related news stories would have both been vying for column inches in U.S. newspapers. Assuming that there is some relationship between the incidence of violent crime in a particular county and violent crime news reportage, we can reasonably expect that newspapers would have focussed more on such stories if the rate of violent crime had increased relative to the previous period/s; this would have come at the expense of other news stories including tsunami-related articles. Therefore we expect a negative relationship between the growth in violent crime rates and media coverage of the tsunami, as newspapers would dedicate more space towards crime stories. It is also reasonable to expect that the selected instrument would not be correlated with either the likelihood or extent of charitable giving since we are considering the *change* in violent crime rates across counties in 2005 rather than the actual level of violent crime. The latter is more likely to be correlated with some omitted variable like neighbourhood affluence, educational attainment, etc. (which are controlled for in our analysis), while the former is mainly dependent on any idiosyncracies related to a given year (in this case, 2005).

Moreover, the use of growth rates in violent crime as our instrument allows us to directly tackle the problem of measurement error caused by the lack of information regarding each article's positioning in a newspaper. This is because as mentioned above, violent crime is a key news topic in the U.S. media, both in terms of actual coverage but also the prominence given to such stories in newspapers (Barak, 1994). Therefore, the use of this instrument will help us to identify only those tsunami-related stories that feature prominently in a newspaper, since only those stories would exhibit the kind of negative relationship described above; other non-prominent tsunami articles would simply not be in competition for col-

[9]As per the FBI's definition, violent crime includes all reported cases of murder, non-negligent manslaughter, rape, robbery and aggravated assault.

umn inches with violent crime stories. As a result, we expect the IV estimates to be higher than those obtained using simple Probit/Tobit, thus correcting for measurement error.

To further ensure the validity of this point, we use two distinct measures of violent crime growth rates: one which considers the annualised growth in violent crime in 2005 relative to 2004, and one which compares violent crime rates in 2005 to the average rate across the preceding 5 years (for each county). Data on county-level annual violent crime were obtained from the FBI's Uniform Crime Reporting Statistics for the years 2000 to 2005. Ideally, since we are only considering the six-month period after the tsunami it would perhaps have been more appropriate to compare violent crime rates between December 26, 2004 to June 26, 2005 with the rates in the preceding months; however only annual data pertaining to each calendar year are available.

Table 3.3 Panel B reports the first-stage results of the linear regression of media coverage on both average and annual change in county-level violent crime for the year 2005. As expected, growth rates in violent crime incidence are negatively and significantly related to the number of tsunami-centered news articles, irrespective of the instrument used or the weighting method employed. The results show that a one percentage-point increase in the growth rate of violent crime across counties is associated with a decline in tsunami-related media coverage by around 23 or 50 news articles, depending on the instrument used.

3.3.2 Results

Table 3.3 Panel A shows the results of our instrumental variables (IV) estimates for Equations 3.1 and 3.2 respectively, using average changes in violent crime rates in 2005 as our instrument. Due to limited space we only report results for the IV estimates in Panel A using changes in violent crime relative to the 5-year average; the results pertaining to annualised changes in violent crime are reported in Appendix B (the two sets of estimates are quantitatively-similar for both equations).

We start with the estimates reported in column 1, where the dependent variable captures the decision to donate money or not. Once again, media coverage is positively and significantly-related to the likelihood of donating money. The corresponding marginal effect of media coverage at the mean[10] is approximately 0.001, which implies that an additional news story on the tsunami would have, on average, increased the probability of donating money by 0.1 percentage points, which shows that despite the statistically significant relationship the actual impact of media coverage on an individual's decision to donate or not was relatively modest. Apart from various demographic characteristics like age and educational background, the results also reflect the strong correlation between prior gift-giving or voluntary work and the individual's decision to donate or not. [11]

[10]The marginal effects for all estimates in Table 3.3, evaluated at the mean, are provided in Appendix C.

[11]Consistent with the findings reported in Brown, Harris, and Taylor (2012).

In fact, the marginal effects obtained for both variables imply that donating more than US$25 to charitable causes in 2004, or undertaking some form of voluntary work, is correlated with a higher probability of donating towards the tsunami relief effort by 18.7 and 9 percentage points respectively.

As expected, the coefficient on family income is statistically significant, although the estimate for wealth is not. Also of note is the impact of geographical location, since the results imply that living in the South, North-East and West of the country was correlated with a higher likelihood to donate money than living in the Midwest[12]. This may be due to the fact that a large proportion of territories within these regions have coastal borders or are located in close proximity to the coast, meaning that people living in these places would be more likely to empathise with the plight of the tsunami survivors given their location and susceptibility to such disasters.[13]

In terms of the estimates reported in column 2, yet again we find that media coverage is positively and significantly related to the size of the donations made to the tsunami relief effort. The marginal effect, evaluated at the mean, of media coverage on the amount of money donated is 0.27, which implies that an additional tsunami-related news story published in a newspaper could have raised the average donation by around US$0.27. Unlike the previous case in column 1, the result shows that media coverage did have a relatively substantial impact on the magnitude of donations disbursed. Consider that, as mentioned previously, the average conditional donation as reported in the survey was of US$121.55; then by extension (assuming linearity) an additional tsunami-news story published in the average American newspaper per week (over the six-month period immediately after the tsunami) could have, according to the results, increased the mean value of donations by almost 7%, which in aggregate translates to around US$131 million, based on estimates for private donations towards the tsunami effort supplied in Becker and Strom (2005).

Once again, we observe the relative importance of both prior donations and prior voluntary work - the marginal effects estimates show that donating over US$25 to charitable causes in 2004 or undertaking voluntary work is correlated with an increase in average donations of some US$41 or US$19 respectively relative to others who did otherwise. Similarly, the role of geographical location on gift-giving is also significant, as are the personal income level and wealth of the individual concerned, which shows that although the decision to donate or not may have been motivated by non-monetary factors, the size of the donation was still related to financial considerations. It is also interesting to note the impact of religious beliefs as highlighted in Andreoni, Payne, et al. (2011), since in this case it seems that belonging to the Jewish faith and other religions (which includes

[12]Geographical regions have been allocated on the basis of the regional divisions provided by the U.S. Census Bureau's designated areas.

[13]In fact in the aftermath of the tsunami many of these coastal areas launched various initiatives like tsunami survival training, general tsunami public awareness information campaigns at local universities and tsunami-preparedness assessments. In addition, following the 2004 tsunami the West Coast and Alaska Tsunami Warning Center expanded its scope of coverage to include the U.S. and Canada's Atlantic coasts.

various religions like Buddhism and Hinduism) is correlated with higher average donations relative to Atheists, Catholics and Protestants.

To summarise, the instrumental variables approach (IV) estimates show that media coverage of the 2004 tsunami had a positive and significant impact on both the decision to donate money or not as well as the amount of money donated. In the former case, the results indicate that despite the statistical significance of the results, media coverage only had a modest effect on the individual's likelihood of donating money. On the other hand, the results in column 2 show that media coverage did indeed have an important and non-trivial impact on the size of charitable donations.

3.3.3 Interaction Terms

So far this paper has mainly focussed on how media coverage of the 2004 Indian Ocean tsunami influenced charitable donations from an aggregated perspective. However, it is also possible to assess how media coverage influenced the donating behaviour of different sub-groups within our dataset. This will enable us to determine which sub-groups or individuals are more responsive to media messages in terms of their observed behaviour in relation to the content of the news signal. Due to space limitations we shall only consider Equation 3.2 in the subsequent analysis; this is justified on the basis that as mentioned previously the estimates obtained for Equation 3.1 suggest that media coverage only had a minor impact on the likelihood of donating money.

The results are shown in Table 3.4, whereby each interaction term is included on a one-by-one basis in Equation 3.2, with the exception of those variables pertaining to the same broad class (for example the education interaction terms are all included in the same regression; the same applies for age, religion, race and geographical location). The estimates show that only two interaction terms have statistically significant coefficients, namely those pertaining to participation in voluntary work and the 25-34 respondent age group.[14] This implies that increased media coverage of the tsunami had the highest level of correlation with charitable disbursements for people who undertook some form of voluntary work in 2004 (relative to those who did not) as well as young adults in the 25-34 age bracket (relative to the under 25 and over 64 age groups). The slope coefficients for the voluntary work and 25-34 age group interaction terms are 0.443 and 0.5235 respectively, which in both cases is more than double the coefficient obtained on media coverage in the original model estimated in Table 3.2 (column 2).

Therefore, the inclusion of various interaction terms suggests a stronger correlation between media coverage of the tsunami and the level of charitable donations among people who undertook some sort of voluntary work in 2004 as well as individuals in the 25-34 age bracket (relative to the omitted age groups). The former effect is fairly predictable since one might expect people with volunteering tendencies to be more aware of, and possibly respond more emphatically to, any information related to a charitable cause, but the latter effect is some-

[14] Where the omitted reference categories are the under 25 and over 64 age groups

Table 3.3: IV Regression Results

Panel A: Second-Stage IV Results	(1) Donation > 0		(2) Amount donated to tsunami	
Tsunami Articles	0.00411***	(0.00106)	1.284***	(0.462)
More than $25	0.696***	(0.0322)	202.7***	(12.61)
Total family income	0.000000436***	(0.000000142)	0.000273***	(0.0000878)
Wealth	1.31e-08	(1.85e-08)	0.0000204***	(0.00000499)
Voluntary Work	0.310***	(0.0410)	90.47***	(6.259)
High School	0.0944	(0.0687)	16.81	(22.09)
Undergrad	0.172**	(0.0849)	42.14*	(23.48)
Postgrad	0.226***	(0.0708)	69.33***	(23.84)
Catholic	0.0139	(0.0456)	-1.675	(7.603)
Jewish	0.295***	(0.101)	80.22***	(26.51)
Protestant	-0.00423	(0.0612)	12.70	(13.44)
Other	0.147*	(0.0772)	98.00***	(33.86)
Between 25-34	-0.0850	(0.0631)	-30.23	(20.15)
Between 35-44	0.0955**	(0.0478)	28.27**	(11.61)
Between 45-54	0.0254	(0.0572)	11.95	(19.96)
Between 55-64	0.157***	(0.0398)	37.90***	(10.03)
Gender head	0.191***	(0.0512)	42.27**	(17.40)
Marital Status	0.144***	(0.0492)	31.31	(19.29)
Working status head	0.209***	(0.0559)	60.10***	(15.11)
Working Wife	0.0342	(0.0322)	19.36*	(10.31)
Retired status head	0.169*	(0.0952)	52.39**	(26.07)
African American	-0.142***	(0.0382)	-30.33*	(16.80)
Hispanic	0.00831	(0.0922)	14.83	(22.81)
Northeast	0.297***	(0.0511)	81.70***	(20.47)
South	0.229***	(0.0347)	59.53***	(10.90)
West	0.276***	(0.0492)	105.0***	(22.53)
Constant	-2.187***	(0.101)	-722.1***	(46.36)

Panel B: First-Stage IV Results	(1)		(2)	
Average change violent crime	-48.52***	(2.752)		
Yearly change violent crime			-21.35***	(1.078)
Constant	65.47***	(0.654)	66.46***	(0.668)
Observations	7215		7215	

Notes: Panel A shows the second-stage results of the instrumental variables (IV) regression analysis. Column 1 reports IV Probit coefficients where the dependent variable is a dummy that takes a value of 1 if the survey respondent donated money towards the 2004 Indian Ocean tsunami relief effort. Column 2 reports IV Tobit coefficients where the dependent variable denotes the value in US$ of individual donations towards the relief effort. In both cases the explanatory variable of interest (which is 'instrumented out') is the number of tsunami related articles published in the leading newspapers for each survey respondent's home county. All data are from the 2005 Center of Philanthropy Panel Study, with the exception of the tsunami articles which were collated from the NewsLibrary website as well as individual newspaper archives. Panel B shows the first stage IV results where in both columns the dependent variable is the number of tsunami articles and both explanatory variables are the instruments based on changes in violent crime. The first row shows results using growth rates in 2005 county-level violent crime relative to the average from the previous five years as the instrument (which is used as the instrument of choice in Panel A) while the second row utilizes growth rates in 2005 county-level violent crime compared to 2004 (used as the instrument of choice in Appendix B). Data for both instruments are from the FBI Uniform Crime Reporting Statistics 2000-2005. Robust standard errors, clustered at the Beale code level, are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3.4: Inclusion of Interaction Terms

	Amount donated to tsunami	
More than $25	0.0431	(0.2086)
Highschool	0.1441	(0.1292)
Undergrad	-0.0146	(0.2232)
Postgrad	0.6957	(0.8367)
Catholic	-0.0363	(0.3287)
Jewish	-0.5952	(0.5446)
Protestant	-0.2272	(0.2873)
Other	-0.4004	(0.4622)
Between 25-34	0.5235*	(0.2784)
Between 35-44	0.1358	(0.3412)
Between 4554	0.3202	(0.2758)
Between 5564	-0.0234	(0.2787)
Gender head	-0.2724	(0.2059)
Working status head	0.2039	(0.2167)
Retired status head	-0.2354	(0.2832)
Working wife	0.2448	(0.2212)
African American	-0.1746	(0.2109)
Hispanic	0.0529	(0.3534)
Marital status	0.0981	(0.1984)
Northeast	-0.2932	(0.4021)
South	-0.1342	(0.2251)
West	-0.4206	(0.4361)
Voluntary Work	0.4430*	(0.2322)

Notes: The table above estimates a standard Tobit model, where the dependent variable used in this table is the value in U.S.$ of individual donations towards the tsunami relief effort. The explanatory variables are obtained by multiplying each variable by the number of tsunami articles in order to obtain cross-partials. Each cross-partial is included separately or as part of its broader class (for example the age variables are all included together), as explanatory variables in the standard Tobit model estimated in Table 3.3. All estimates reported are Tobit coefficients. Robust standard errors are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

what more suggestive. The value of such analysis is twofold - firstly, it may assist charitable organisations in the design of appropriate donation solicitation campaigns by identifying precisely which types of individuals are more likely to increase their disbursements in response to information related to the charitable cause in question. So in this case, it appears that such information campaigns should be targeted at people with prior volunteering propensities (a somewhat self-evident conclusion) and, crucially, young adults. Secondly, the analysis also has wider implications for the study of the media's role in society as a whole since it enables us to identify which individuals or cohorts are more likely to be susceptible to news signals received from the media, although clearly one cannot infer any causal linkages from the above analysis. The importance of these findings go beyond the current framework of newspaper reports related to the 2004 tsunami, and extend to other issues like articles related to climate change, Obamacare and obesity, across a wide variety of media outlets. Various studies (primarily in the psychology literature) have shown the extent to which young adults are influenced by media reports in relation to various issues. For example, Harrison and Cantor (2006) find a positive and significant relationship between exposure to media content on ideal body shape and college women's negative self-perceptions and the incidence of eating disorder symptoms. This paper therefore complements this literature by suggesting that the young adult segment of the population (25-34) exhibited a stronger correlation between exposure to media content and their economic choices (as gauged by the increase in tsunami donations) relative to other age groups and controlling for a whole host of factors like education, religion, gender and prior donating/volunteering behaviour, although once again we must caution against any causal interpretations.

3.4 Robustness Tests

The preceding section showed how media coverage of the 2004 Indian Ocean tsunami influenced both the private individual's decision to donate money or not as well as the magnitude of such disbursements. These results hold under various conditions, including the use of alternative weighting methods for calculating media coverage in each U.S. county (i.e. by population vs. by circulation) and the use of annual change in county-level violent crime (from 2004 to 2005) as the instrumental variable of choice as opposed to the change in 2005 crime relative to the previous 5-year average. In this section we present a number of robustness checks designed to further test the validity of our results.

3.4.1 Additional Variables

The first set of robustness checks involve the inclusion of a number of additional variables which may somehow be correlated with our instrumental variable of choice as well as the likelihood of donating money and the amount donated.

Local community donations: The first variable we consider is related to the incidence and level of charitable donations towards community and neigh-

bourhood causes. It might be the case that a spike in violent crime during 2005 in a particular county may have fueled concerns regarding the well-being of local communities, which in turn may have led individuals to donate more money towards neighbourhood causes or other similar charities, at the expense of tsunami donations (and irrespective of the coverage afforded by the media to either topic). In this case there would still be the negative relationship between changes in violent crime rates and the likelihood/level of tsunami contributions as postulated above, albeit through a patently-different mechanism. Hence, the existence of such a relationship may lead to justifiable concerns regarding the reliability of our identification strategy as well as any results obtained in the previous section. It is therefore necessary to address such concerns via the inclusion of an additional variable that captures charitable donations towards neighbourhood causes in our IV regression models. Fortunately, data on the incidence as well as the level of donations towards neighbourhood and community improvement organisations in the U.S. for 2005 are readily available via the same 2005 CPPS dataset from where the tsunami donation data were derived, thus ensuring some degree of harmonisation and compatibility. We utilise this data to construct two variables - one is a dummy variable denoting whether the survey respondent donated any money towards neighbourhood/community causes in 2005, and another which reflects the amount of money donated by each respondent in 2005 towards such causes, with each variable included as part of the IV regressions estimated in Table 3.3, columns 1 and 2 respectively.

Population size: Secondly, we include the population size of each respondent's county of residence. The idea behind this inclusion is that people from large urban counties may be more likely to donate money to international aid causes such as the Indian Ocean tsunami. This is due to various factors like greater exposure to foreign cultures due to the cosmopolitan nature of large cities, greater awareness of international affairs and relief efforts due to increased media coverage, and an inherent predisposition to donate to such causes since urban dwellers may have innate characteristics that are compatible with such behaviour (Rudzitis, 1999). Therefore, it is likely that population size is correlated with both the amount of media coverage devoted to the tsunami as well as the likelihood/size of tsunami donations. More importantly, it is also possible that population size may be systematically correlated with changes in violent crime rates. Although the actual direction of the correlation may not be immediately clear ex-ante, a simple regression of changes in violent crime on population size in our data reveals a negative relationship whereby larger cities experienced the largest decline in violent crime rates in 2005; a trend that has largely persisted over the last 10 years (FBI, 2015) . It is therefore possible that population size may be driving our results even after the use of our violent crime instrument.

Asian population: The third additional variable is the proportion of people of Asian descent living in the survey respondents' county of residence. As mentioned earlier the tsunami's devastation was mainly concentrated in South East Asia, both in terms of fatalities as well as damages incurred. It is reasonable to expect that a higher proportion of Asian residents within a given county would have resulted in both a higher probability of donating money as well as higher

average donation levels. This is because such a disaster would, in all likelihood, have resonated to a larger extent among people of Asian origin, and thus increasing the likelihood and extent of any charitable disbursements. In addition, the presence of a large Asian community within a county may also have induced other residents to donate money or increase their contribution due to possible peer effects which could have led to a heightened awareness of the tsunami's impact. However, this variable may also have an acute impact on media coverage of the tsunami, since newspaper editors in such areas would have been expected to maintain a consistently high stream of tsunami-related news stories given the level of interest among the local community. We therefore introduce a variable based on US Census Bureau data from 2000, which measures the proportion of residents of Asian descent within each county.

International awareness: Finally, we also include a variable that seeks to capture respondents' awareness or concern with regards to international aid causes. The importance of innate non-monetary factors was evident in the results presented above; nonetheless it can reasonably be argued that prior donations and voluntary work are insufficient for the purposes of capturing respondents' international concern given that these two variables (particularly the voluntary work variable) may simply be reflecting domestic charitable activities. This international awareness or concern may be important not just in terms of its relationship with the dependent variable/s, but also due to the fact that a newspaper whose core readership holds an active interest in international relief issues would not diminish its coverage of the tsunami even in the face of rising violent crime instances. On the other hand, the opposite is also likely to hold true - a newspaper would be more likely to reduce its tsunami coverage and replace it with crime stories (which would be decidedly local affairs) if its readers are mainly concerned with domestic issues. To estimate respondents' international concern we use data on individual charitable donations towards international development agencies in 2005, which is collected as part of the CPPS survey. A central issue with this data is that it may also include donations made towards the tsunami relief effort, since technically it would still count as international aid. Hence we subtract any tsunami-related donations from the international aid figures, and construct a dummy variable which assumes a value of 1 if the net international aid donations (i.e. international aid less tsunami donations) are greater than US$0.

The IV results with the additional included variables are shown in Table 3.5. As seen from the results, both sets of estimates are largely unchanged from those obtained in the previous section. The coefficient on media coverage in column 1 is very close to the previous estimate, while the same coefficient in column 2 is still quantitatively close to 1 as before. Note that the coefficient on the neighbourhood/community donations dummy variable in column 1 is statistically significant. This indicates that people who donated money to neighbourhood/community charities were also more likely to donate money towards the tsunami relief effort, thus negating the alternative causal mechanism postulated in this section where higher violent crime fuels higher incidence of such contributions at the expense of tsunami-related disbursements. Conversely, the size of neighborhood/community donations had no impact on the level of tsunami do-

nations as seen in column 2. Similarly, population size had a negative influence on the likelihood of donating money (column 1), although the magnitude of the estimate itself is negligible, but had no impact on the amount of money donated (column 2). Note that neither the variable denoting the size of the Asian population in each county nor the dummy for international concern yield statistically significant coefficients in both of the models estimated.

3.4.2 Validity of Identification Strategy

In this section, we further assess the validity of our first-stage IV results (and the underlying assumptions underpinning the relationship observed) by running a similar regression, only this time using the number of sports-related stories as our dependent variable. The reliability of the IV strategy employed rests on the assumption that a sudden surge in county-level violent crime rates in 2005 relative to previous years led to a spike in crime-related stories within a newspaper's main circulating county, which in turn would have been published at the expense of more tsunami-related coverage. Hence, in essence our story is one where crime news stories, fuelled by a surge in violent crime, crowd-out newspaper coverage of the tsunami due to limited newspaper column inches.

In reality there may be other hypotheses which could reasonably explain this negative relationship between growth in violent crime and tsunami articles other than crowding-out, and which in turn may somehow be correlated with the decision to donate money towards the relief effort. In the previous subsection we sought to tackle some of these potential alternative stories through the inclusion of additional explanatory variables, like for example respondents' international concerns and the proportion of people of Asian descent living in each county. However, it is entirely possible that there are other potential explanations not taken into account in the previous set of robustness tests, and which may raise concerns regarding the IV strategy employed.

Therefore, we propose a more direct way of testing the validity of our first-stage results. This is based on the idea that if crime stories do indeed 'crowd-out' tsunami articles, then it must also be the case that they crowd-out other news stories too. We therefore test to see whether we observe the same qualitative relationship when using the number of sports-related news stories as our dependent variable. If this were not the case then it may easily be argued that the first-stage results reported above are either spurious or else evidence of some alternative identification story than the one proposed in this paper. For example, it could simply be capturing variations in the international outlook or openness of the newspaper/readers in question, which in turn may reasonably be correlated with the individual decision to donate money or not towards the tsunami relief effort.[15] By running the first stage regression on another, completely-different type of news story, namely sports, we may directly test whether the crowding-out story

[15]As mentioned earlier, we have already sought to address this issue through the inclusion of the international aid dummy, although in reality this may constitute a poor proxy for actual international awareness of respondents.

Table 3.5: IV Results with Additional Variables

	(1) Donations > 0		(2) Amount donated to tsunami	
Tsunami articles	0.00546**	(0.00233)	1.708**	(0.832)
More than $25	0.664***	(0.0390)	200.6***	(12.65)
Total family income	0.000000403***	(0.000000136)	0.000261***	(0.0000889)
Wealth	1.20e-08	(1.94e-08)	0.0000190***	(0.00000383)
Voluntary work	0.293***	(0.0469)	89.28***	(7.545)
High School	0.0852	(0.0712)	15.14	(22.47)
Undergrad	0.149*	(0.0871)	38.25*	(21.84)
Postgrad	0.194**	(0.0783)	63.64**	(24.84)
Catholic	0.0173	(0.0455)	0.791	(8.588)
Jewish	0.258***	(0.0943)	74.49**	(34.03)
Protestant	0.00103	(0.0629)	14.86	(14.95)
Other Religion	0.137*	(0.0738)	97.40***	(35.67)
Between 25-34	-0.0698	(0.0606)	-29.63	(19.95)
Between 35-44	0.111**	(0.0452)	30.56***	(10.52)
Between 45-54	0.0293	(0.0549)	11.88	(19.82)
Between 55-64	0.167***	(0.0386)	38.69***	(10.09)
Gender head	0.178***	(0.0513)	40.79**	(17.06)
Marital Status	0.146***	(0.0491)	32.00*	(19.28)
Working status head	0.207***	(0.0577)	60.83***	(16.16)
Working Wife	0.0403	(0.0284)	21.74**	(9.043)
Retired status head	0.166*	(0.0966)	54.73**	(27.61)
African American	-0.144***	(0.0390)	-31.06*	(17.89)
Hispanic	0.0110	(0.0796)	15.80	(20.19)
Northeast	0.227***	(0.0624)	69.93***	(26.28)
South	0.231***	(0.0577)	62.10***	(16.27)
West	0.343***	(0.103)	130.3***	(34.05)
Dummy community donations	0.254***	(0.0322)		
Level of community donations			0.0780	(0.0507)
Population	-5.91e-08**	(2.87e-08)	-0.0000138	(0.0000112)
Asian population	0.0168	(0.0118)	2.408	(2.849)
Dummy int aid	0.129	(0.102)	4.084	(20.93)
Constant	-2.242***	(0.144)	-744.9***	(58.04)
Observations	7208		7208	

Notes: Column 1 reports IV Probit coefficients where the dependent variable is a dummy which takes a value of 1 if the survey respondent donated money towards the 2004 Indian Ocean tsunami relief effort. Column 2 presents IV Tobit coefficients where the dependent variable denotes the value in US$ of individual donations towards the relief effort. In both cases the explanatory variable of interest is the number of tsunami-related articles published in the leading newspapers for each of the survey respondents' home counties. Once again we use growth rates in county level violent crime in 2005 relative to the previous five year average as our instrument for the tsunami articles variable. The key difference between this table and the results reported in Table 3.3 above is the inclusion of 4 additional variables - charitable donations towards neighborhood/local community causes (a dummy is included in column 1, which assumes a value of 1 if the survey respondent donated money to such causes; while in column 2 the US$ value of such donations is included), population size, the Asian population in each county, as well as a dummy denoting international aid donations. All data are from the 2005 Center of Philanthropy Panel Study, with the exception of the tsunami articles data which were collated from the NewsLibrary website and individual newspaper archives as well as the violent crime data, which are from the FBI Uniform Crime Reporting Statistics for the period 2000-2005. Robust standard errors, clustered at the Beale code level, are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

proposed above is valid or not.

The use of sports stories as our new dependent variable can be justified on several fronts.Apart from being completely different from violent crime stories, sports stories can largely be classified as local news rather than international affairs, meaning that we can directly observe whether a surge in violent crime crowded out all news stories, or simply those related to international incidents (like the tsunami).[16] Once again, we collect our sports news data using NewsLibrary as well as the individual newspapers' online print archives, over the period December 26, 2004 to June 26, 2005. In this case, we searched for articles containing the words "sports", "football', "basketball", "baseball", "soccer", "hockey", "golf", "tennis", "wrestling", "NASCAR", "IndyCar" or "UFC" in either the headline or lead paragraph, once again omitting articles containing the words "letter" or "letters". The rather exhaustive list of search terms was used in order to cover the entire spectrum of popular sports across the U.S.. This process was conducted for each of the 386 newspapers covered in the previous section, using the same weighting methods described above.

Table 3.6: Changes in Violent Crime & Sports Coverage

	(1) Sports Articles
Average change violent crime	-1908.7** (907.4878)
Constant	1768.7*** (442.65)
N	7215

Notes: In this table we present an alternative first-stage regression to the one specified in Table 3.3 (Panel B), this time using the number of sports-related stories published in each of the respondents' leading newspapers over the period December 26, 2004 to June 26, 2005 as our dependent variable. Data on the sports articles data which were collated via the NewsLibrary website, while the violent crime data are from the FBI Uniform Crime Reporting Statistics 2000-2005. Robust standard errors, clustered at the Beale code level, are shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The results are shown in Table 3.6. We can observe a clear, negative relationship between changes in violent crime and newspaper coverage of sports news, similar to the one observed for tsunami articles in our first stage IV results. Hence, this supports the argument that a sudden spike in violent crime and therefore crime-related stories in 2005 led to less coverage of the tsunami, and indeed less coverage of other major stories, due to crowding out resulting from the newspaper's limited column inches.

[16]In 2005 no major international sporting events were held, like for example the Olympic Games, therefore minimising the possibility that our sports-related coverage variable is in fact capturing international stories.

3.5 Concluding Remarks

This paper has sought to examine the causal link between media coverage of a news story and people's responsiveness to the informational content of the article. More specifically, we looked at the influence of newspaper coverage of the 2004 Indian Ocean tsunami on both the likelihood as well as the magnitude of individual private donations. As part of the analysis, we used county-level survey data on charitable gift giving in the U.S. for the year 2005 as well as data on the number of tsunami-related news stories published by the top-circulating newspapers pertaining to each county in the six-month period immediately after the disaster (December 26, 2004 to June 26, 2005). The first set of results showed a positive and statistically significant correlation between media coverage and the probability/size of charitable disbursements, although the reliability of these estimates was called into question due to endogeneity concerns. Therefore, it was decided to adopt an instrumental variables (IV) approach to estimation, with county-level changes in violent crime rates in 2005 employed as our main instrument.

Based on this approach, the results showed that newspaper coverage only had a modest (albeit, statistically significant) effect on the likelihood of donating money. By contrast, we observe that coverage of the tsunami had a non-trivial, positive and statistically significant impact on the amount of money donated. In fact, the results imply that an additional tsunami-news story per week published in the average American newspaper (over the six-month period under review) could have increased the mean value of donations by around 7%, which translates to approximately US$131 million - an amount which demonstrates the relative influence of the media in terms of encouraging higher levels of charitable donations from private individuals. The results are consistent under a number of different specifications, namely the type of weighting method used in order to calculate the newspaper coverage variable (by population vs. by circulation) as well as the type of change in violent crime rates employed as our instrument (i.e. changes in 2005 violent crime relative to violent crime in 2004 vs. the average of the preceding 5 years). Furthermore, the inclusion of various interaction terms constructed using the media coverage variable and several categorical controls showed that the media had a larger impact, in terms of encouraging higher donations, on people with prior volunteering experience as well as young adults in the 25-34 age bracket. A number of robustness checks were then performed in order to further test the reliability of the results; this involved the inclusion of additional control variables for the incidence and size of community/neighbourhood donations, the proportion of the county-level population of Asian descent and international aid donations as well as a test on the validity of our identifying relationship.

Therefore, the evidence strongly suggests that the media has a direct and significant impact on people's behaviour and decision-making processes. This point is further highlighted by the fact that although this paper is solely concerned with newspaper coverage of the tsunami, the magnitude of the media's influence on the size of charitable donations was still relatively high. The newspaper is a relatively passive form of media consumption when compared to other channels

like television or online news portals which include audio-visual news reports and other immersive content, in contrast with the standard articles found in the print media. In fact, McRoberts and Terhanian (2008) find that in the UK, France and Germany online content is almost twice as influential as television and approximately ten times more influential than newspapers or magazines.[17] Hence the results shown in this study merely represent a lower bound in terms of the media's impact on people's behaviour, and may in aggregate be much higher than reported, which further underlines the growing role of the media in society. The potential impact of biased media reportage on individual behaviour is well-documented and has already been discussed; however, what this paper suggests is that even the amount of coverage given to a particular news story or topic can influence people's perception of the subject matter and ultimately their behaviour. In this paper the topic in question was the 2004 Indian Ocean tsunami, a relatively uncontroversial and decidedly non-partisan issue which was characterised by the global nature of the subsequent relief effort. Although this goes beyond the current scope of the paper, it would be interesting to consider the implications of such findings on the impact that media coverage could have on various other issues ranging from health (for example the incidence of obesity) to public perceptions regarding global warming (and all that entails in terms of buying environmentally-friendly products).

Another key result from this paper is the fact that young adults in the 25-34 age bracket were more responsive to media coverage of the tsunami, in terms of raising the amount donated, relative to other age-groups, even after controlling for various factors like income, wealth, education and prior donating tendencies. As mentioned earlier, this complements numerous studies from psychology which show the extent to which young adults are influenced by media reportage in relation to various issues like concern for body image (Harrison and Cantor, 2006). What this paper suggests is that the impact of the media on young adults' individual behaviour also extents to their economic choices, which has important implications from an industrial organisation point of view in terms of the formulation of media advertising strategies. This is because although the results indicate that 25-34 year old individuals are more responsive to media content, this may easily be extended to include advertising messages. In fact, other studies have also shown the responsiveness of young adults to advertising messages in the media - for example Lamont, Hing, and Gainsbury (2011) find that young people in general are more likely to engage in sports-related gambling activities following exposure to advertisements promoting gambling in Australia. Hence, media outlets with a large proportion of its audience in the 25-34 age bracket could use this information in order to generate greater advertising revenue, which is of vital importance given the decline in traditional sources of income like paid subscriptions (particularly for newspapers) and the predominance of advertising-driven media firms like free-to-air TV channels and online news portals. The 18-34 age group has long been considered as the key advertising demographic (Stabile and

[17]Whereby the authors define the level of 'influence' as the total amount of time that consumers spend on each medium, combined with the importance that respondents attribute to each medium in terms of influencing their daily purchasing decisions via advertising.

Harrison, 2003) across various media due to the fact that this cohort typically has more disposable income and less brand loyalty than other age groups, and thus more receptive to advertising. What this paper suggests is that even after accounting for income and other individual characteristics, the young adult segment is still more likely to respond to media messages relative to others, which implies that this tendency may be driven by non-monetary factors like cognitive or psychological considerations.

In addition, the results also highlight the potential role of the media as an effective fund-raising tool for charitable causes. The relative merits of various solicitation techniques for charitable donations have been studied at length by various authors - for example Diamond and Iyer (2007) look at the numerous considerations related to the use of direct mail in order to solicit funds from the public, with response rates often as low as 0.5%. This paper suggests that while increased media coverage of a particular issue would only have a minor impact in terms of convincing people to donate money to the corresponding charitable cause/organisation, it may nonetheless prove to be a viable and relatively cost-effective means of increasing the monetary value of donations, particularly given today's technological landscape where several social media and online tools can be mobilised almost instantly in order to raise funds.

3.6 Appendices

Appendix A1: Description of Variables

Variable Name	Description of Variable	Source
Donations ¿ 0	Dummy Variable denoting whether respondent has donated money to the tsunami relief effort	CPPS 2005[18]
Amount Donated	Amount (in US$) of money donated towards the tsunami relief effort	CPPS 2005
More than $25	Dummy denoting whether respondent donated more than $25 to charitable causes in 2004	CPPS 2005
Family income	Income (in US$) of household in 2005	CPPS 2005
Wealth	Value (in US$) of household's net asset holdings, including property	CPPS 2005
High school	Dummy denoting whether respondent ceased formal education at High School level	CPPS 2005
Undergrad	Dummy denoting whether respondent ceased formal education at undergraduate level	CPPS 2005
Postgrad	Dummy denoting whether respondent ceased formal education at postgraduate level[19]	CPPS 2005
Catholic	Dummy denoting whether respondent is a Catholic	CPPS 2005
Jewish	Dummy denoting whether respondent is Jewish	CPPS 2005
Protestant	Dummy denoting whether respondent is a Protestant	CPPS 2005
Other	Dummy denoting whether respondent belongs to another formal religious organisation (including Hinduism, Buddhism, etc.)[20]	CPPS 2005
Between 25-34	Dummy denoting whether respondent's age is between 25 and 34 years	CPPS 2005
Between 35-44	Dummy denoting whether respondent's age is between 35 and 44 years	CPPS 2005
Between 45-54	Dummy denoting whether respondent's age is between 45 and 54 years	CPPS 2005
Between 55-64	Dummy denoting whether respondent's age is between 55 and 64 years[21]	CPPS 2005
Gender head	Dummy denoting whether the respondent is male	CPPS 2005
Working status head	Dummy denoting whether respondent is gainfully occupied	CPPS 2005

[18]Wilhelm (2006)

[19]Omitted category in this case is for respondents with no recognised educational qualifications.

[20]Omitted category in this case is for respondents with no religious beliefs (i.e. Atheists).

[21]In this case due to lack of data we have two omitted categories, namely respondents below the age of 25 and those above the age of 65.

Variable Name	Description of Variable	Source
Retired status head	Dummy denoting whether respondent is retired from work	CPPS 2005
Working wife	Dummy denoting whether respondent's wife is gainfully occupied	CPPS 2005
African American	Dummy denoting whether respondent is African American	CPPS 2005
Hispanic	Dummy denoting whether respondent is Hispanic[22]	CPPS 2005
Marital status	Dummy denoting whether respondent is married	CPPS 2005
Northeast	Dummy denoting whether respondent's county is located in the North East of the U.S.	US Census Bureau
South	Dummy denoting whether respondent's county is located in the South of the U.S.	US Census Bureau
West	Dummy denoting whether respondent's county is located in the West of the U.S.[23]	US Census Bureau
Voluntary work	Dummy denoting whether respondent undertook any voluntary work in 2004	CPPS 2005
Dummy int aid	Dummy denoting whether respondent donated any money towards international aid relief in 2005	US Census Bureau
Community donations	Amount (in US$) of money donated towards neighbourhood or community improvement causes in 2005	CPPS 2005
Asian population	Proportion of each county's population defined as being of Asian descent	US Census Bureau

[22]Omitted category in this case is all other ethnic backgrounds (the CPPS only indicates whether respondent is African American, Hispanic or Other in the survey).

[23]Omitted category in this case is counties located in the Midwest.

Appendix A2: Descriptive Statistics of Explanatory Variables

Variable	Mean	Std. Dev.	Min.	Max.	N
Donations ¿ 0	0.239	0.427	0	1	7215
Amount Donated to Tsunami	29.078	142.264	0	6000	7215
More than $25	0.617	0.486	0	1	7215
Total family income	63486.141	107568.903	-39000	5500000	7215
Wealth	233800.901	958876.278	-363000	42000000	7215
Voluntary work	0.279	0.474	0	8	7215
High school	0.235	0.424	0	1	7215
Undergrad	0.224	0.417	0	1	7215
Postgrad	0.048	0.214	0	1	7215
Catholic	0.196	0.397	0	1	7215
Jewish	0.02	0.14	0	1	7215
Protestant	0.609	0.488	0	1	7215
Other	0.055	0.229	0	1	7215
Between 25-34	0.234	0.423	0	1	7215
Between 35-44	0.215	0.411	0	1	7215
Between 45-54	0.227	0.419	0	1	7215
Between 55-64	0.119	0.324	0	1	7215
Gender head	0.302	0.459	0	1	7215
Marital status	0.501	0.5	0	1	7215
Working status head	0.751	0.432	0	1	7210
Working wife	0.38	0.485	0	1	7213
Retired status head	0.112	0.315	0	1	7210
African American	0.346	0.476	0	1	7215
Hispanic	0.075	0.264	0	1	7215
Northeast	0.146	0.353	0	1	7215
South	0.407	0.491	0	1	7215
West	0.192	0.394	0	1	7215
Community donations	12.307	166.419	0	10000	7215
Asian population	3.281	3.287	0.171	49.529	7215
Dummy int aid	0.027	0.162	0	1	7215

Appendix B: IV Regression Output - Yearly Change in Violent Crime

	(1) Donations ¿ 0		(2) Amount donated to tsunami	
Tsunami Articles	0.00445***	(0.00161)	1.117**	(0.479)
More than $25	0.693***	(0.0390)	203.7***	(13.30)
Total family income	0.000000427***	(0.000000139)	0.000277***	(0.0000906)
Wealth	1.27e-08	(1.86e-08)	0.0000204***	(0.00000501)
Voluntary work	0.310***	(0.0407)	90.13***	(6.330)
High School	0.0958	(0.0712)	15.91	(22.29)
Undergrad	0.170**	(0.0842)	42.67*	(23.61)
Postgrad	0.223***	(0.0722)	70.76***	(24.58)
Catholic	0.0125	(0.0456)	-1.635	(7.514)
Jewish	0.280***	(0.0974)	85.75***	(25.75)
Protestant	-0.00273	(0.0628)	11.43	(13.04)
Other	0.144*	(0.0758)	98.35***	(33.84)
Between 25-34	-0.0847	(0.0627)	-30.32	(20.01)
Between 35-44	0.0959**	(0.0476)	28.37**	(11.39)
Between 45-54	0.0252	(0.0569)	12.41	(19.92)
Between 55-64	0.156***	(0.0400)	38.09***	(9.973)
Gender head	0.188***	(0.0515)	42.32**	(17.49)
Marital Status	0.144***	(0.0487)	30.92	(19.13)
Working status head	0.208***	(0.0555)	60.33***	(14.95)
Working Wife	0.0349	(0.0330)	18.62*	(10.60)
Retired status head	0.169*	(0.0938)	52.73**	(25.72)
African American	-0.149***	(0.0518)	-26.42	(16.69)
Hispanic	0.00630	(0.0913)	16.39	(22.45)
Northeast	0.297***	(0.0532)	82.19***	(19.25)
South	0.239***	(0.0493)	55.03***	(13.98)
West	0.288***	(0.0617)	100.0***	(22.61)
Constant	-2.206***	(0.118)	-710.1***	(47.73)
Observations	7208		7208	

Notes: The table shows the second-stage results of the instrumental variables regression analysis. Column 1 estimates an IV Probit model where the dependent variable is a dummy that takes a value of 1 if the survey respondent donated money towards the 2004 Indian Ocean tsunami relief effort. Column 2 shows an IV Tobit model where the dependent variable denotes the value in US$ of donations. In both cases the explanatory variable of interest (which is 'instrumented out') is the number of tsunami articles published in the leading newspapers for each survey respondent's home county. For both models we use growth rates in 2005 county-level violent crime compared to 2004 as our instrument. All data is from the 2005 Center of Philanthropy Panel Study, with the exception of the tsunami articles which were collated from the NewsLibrary website as well as individual newspaper archives, while violent crime data are from the FBI Uniform Crime Reporting Statistics 2000-2005. Robust standard errors, clustered at the Beale code level, are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Appendix C: Marginal Effects for IV Regression Results

	(1) Donations ¿ 0		(2) Amount donated to tsunami	
Tsunami Articles	0.00119***	(0.000322)	0.268***	(0.0970)
More than $25	0.187***	(0.00666)	40.80***	(1.567)
Total family income	0.000000126***	(4.28e-08)	0.0000570***	(0.0000176)
Wealth	3.77e-09	(5.30e-09)	0.00000426***	(0.00000109)
Voluntary work	0.0897***	(0.0105)	18.90***	(1.602)
High School	0.0278	(0.0208)	3.540	(4.752)
Undergrad	0.0515**	(0.0259)	8.995*	(5.269)
Postgrad	0.0704***	(0.0240)	15.35***	(5.450)
Catholic	0.00402	(0.0134)	-0.349	(1.583)
Jewish	0.0942***	(0.0365)	18.01***	(6.052)
Protestant	-0.00122	(0.0177)	2.647	(2.799)
Other	0.0446*	(0.0249)	22.22***	(8.213)
Between 25-34	-0.0241	(0.0180)	-6.223	(4.108)
Between 35-44	0.0282**	(0.0141)	5.994**	(2.409)
Between 45-54	0.00738	(0.0165)	2.511	(4.221)
Between 55-64	0.0474***	(0.0119)	8.131***	(2.116)
Gender head	0.0568***	(0.0157)	8.968**	(3.774)
Marital Status	0.0416***	(0.0148)	6.541*	(3.916)
Working status head	0.0578***	(0.0140)	12.22***	(3.193)
Working Wife	0.00993	(0.00932)	4.061*	(2.172)
Retired status head	0.0515*	(0.0300)	11.37*	(6.164)
African American	-0.0403***	(0.0103)	-6.282*	(3.549)
Hispanic	0.00241	(0.0268)	3.134	(4.873)
Northeast	0.0927***	(0.0160)	18.02***	(4.692)
South	0.0674***	(0.0108)	12.57***	(2.276)
West	0.0849***	(0.0167)	23.33***	(4.912)
Observations	7208		7208	

Notes: The table above shows the computed marginal effects, evaluated at the mean, for the second-stage results of the instrumental variables (IV) regression analysis. Column 1 estimates an IV Probit model where the dependent variable is a dummy that takes a value of 1 if the survey respondent donated money towards the 2004 Indian Ocean tsunami relief effort. Column 2 shows an IV Tobit model where the dependent variable denotes the value in US$ of individual donations towards the relief effort. In both cases the explanatory variable of interest (which is 'instrumented out') is the number of tsunami related articles published in the leading newspapers for each survey respondent's home county, and the instrument of choice is the growth rates in 2005 county-level violent crime relative to the average from the previous five years. All data are from the 2005 Center of Philanthropy Panel Study, with the exception of the tsunami articles which were collated from the NewsLibrary website as well as individual newspaper archives, while data for the violent crime instrument are from the FBI Uniform Crime Reporting Statistics 2000-2005. Robust standard errors are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Chapter 4

Lewinsky and Gore: The Impact of Soft Media Reportage on Voting Behaviour

4.1 Introduction

The debate surrounding the media's reporting practices has gathered considerable pace in recent years, fueled by the perceived influence of such reportage on voting behaviour as well as government policy choices. Lately there has been growing concern with regards to the increased coverage of 'soft' stories over more substantive news items in the mainstream media, with the existence of such reporting tendencies highlighted in various contexts, including U.S. TV network evening news (Hamilton, 2004 and Bagdikian, 2000).

This paper represents the first attempt at assessing the impact of media reports on individual voting patterns, and more specifically the role played by soft or sensationalist news in influencing electoral choices. We use novel data on the amount of coverage afforded to the infamous Monica Lewinsky scandal over the period January 17, 1998 to August 31, 2000 by a sample of 257 leading daily newspapers in the U.S., and link this to survey data from the American National Election Studies (ANES) on individual voting behaviour during the 2000 Presidential election. To account for any potential endogeneity issues we adopt an instrumental variables approach, where the identification strategy exploits differences in the number of deaths caused by extreme weather events for each of the newspapers' main counties of circulation.

The results show that Lewinsky coverage had a positive and statistically significant impact on the likelihood of voting for George W. Bush in the 2000 Presidential election and, in contrast, a negative influence on the probability of voting for Al Gore. This pattern is visible among both Democrats and Republicans, with the effect strongest among the former subgroup. In fact, a 5% increase in newspaper coverage of the Lewinsky scandal could have increased the likelihood that a non-Republican voted for Bush in the 2000 Presidential election by around 1.9%, with a similar decline reported in the probability of voting for Gore. In

terms of the other control variables, prior political choices, race, religion and individual perceptions regarding Bill Clinton also yield a statistically significant impact on voting behaviour. On the other hand, we find that media coverage of the scandal had no impact on voter turnout across all subgroups.

A crucial feature of the Lewinsky scandal is that, as further shown in the results below, newspaper reportage was not driven by any of the demand or supply side determinants of political media bias proposed in the literature (Gentzkow and Shapiro, 2010) like readers' political beliefs, newspaper editorial bias and ownership. This enables us to abstract away from considerations related to partisan-driven media bias and its effect on voting, which has been dealt with by the likes of DellaVigna and Kaplan (2007), and focus solely on the electoral influence of soft or sensationalistic reportage. This independence from political bias also helps in allaying any doubts that an observed relationship between Lewinsky coverage and voting is merely a reflection of a general anti-Democrat or pro-Republican newspaper reporting agenda.

The Lewinsky scandal is the ideal case to analyse since, within the context of the 2000 Presidential election, it fits in perfectly with the idea of 'soft' news reportage. This is because whereas 'hard' news yields positive externalities to society (in terms of improved policy outcomes and greater accountability) through the provision of contextually-relevant information, soft news only provides private benefits in terms of entertainment or personal interest (Ellman, 2014), with limited information regarding the suitability of electoral candidates or policy efficacy. Thus in terms of the 2000 election, the Lewinsky scandal falls squarely within this latter definition since apart from dealing with the extra-marital activities of the President Clinton, it provided little or no information regarding either Al Gore or George W. Bush, the two candidates in the election, in terms of their presidential credentials. The Lewinsky scandal has long been held up within the media and journalism literature as the classic example of soft or sensationalistic reportage (Lawrence and Bennett, 2001), whereby sensationalism primarily refers to stories that have no relevance to people's lives and to society at large, typically related to celebrities, crime, sex and disasters (Stephens, 2007).

A possible contention might be that the Lewinsky scandal was informative in terms of Clinton's moral standards and thus reflected the ethical beliefs or trustworthiness of the entire Democrat-led administration (including Al Gore, who was Vice-President). However, this claim can be countered on several fronts. Firstly, it is unclear whether scandals related to politicians' personal lives has any impact on their electoral performance (Bhatti et al., 2013), let alone the political fortunes of their colleagues or successors. In fact, the Lewinsky scandal seemed to have little impact on Clinton's own reputation, since his approval ratings at the end of his Presidency were the highest recorded in 50 years at almost 70% (higher than the ratings recorded at any point before the scandal erupted in January 1998). Secondly, we explicitly account for such concerns by exploiting the richness of our survey data in order to control for respondent approval of Clinton's presidency as well as individual perceptions regarding his moral beliefs. Finally, we explicitly control for people's moral qualms with Clinton (and, by association, Gore) by including control variables denoting respondents' religious

beliefs as well as their perceptions regarding Clinton's impact on the country's moral climate.

The paper also considers a number of alternative hypotheses which may explain the observed relationship between Lewinsky coverage and voting. Firstly, we address concerns regarding the potential biases involved in self-identification by respondents as Republicans or Democrats by using an alternative sub-grouping classification system based on prior voting patterns in the 1996 Presidential election. The results are largely unchanged, and in particular the swing from Democrat to Republican is also evident among those who had previously voted for Bill Clinton. Next, we revisit the idea that the results are being driven by newspapers' political leanings by explicitly including an explanatory variable denoting editorial slant; the results are once again unaffected and merely serve to confirm that Lewinsky coverage was not motivated by partisan beliefs.

In addition, we also consider the possibility that our Lewinsky coverage variable is merely a reflection of media scrutiny afforded to several other scandals that rocked the Clinton administration, and in particular the 1996 Whitewater investigation on various controversies which, within the context of this paper, may have been informative in terms of Al Gore's presidential capabilities. Given that the investigation was opened prior to the 1996 Presidential election, we run a placebo treatment in order to assess whether newspaper coverage of the Lewinsky scandal had any impact on voting in 1996 election. We find no evidence of any relationship between Lewinsky coverage and 1996 voting patterns. Finally, the results are robust to a number of additional checks, including including the use of an alternative weighting system for our media coverage variable, the use of a different instrument and testing for attrition bias.

The results indicate that soft news stories which provide little or no information regarding the capabilities or suitability of electoral candidates may still have an impact on people's voting behaviour. This underscores the significance of media reportage as a determinant of political outcomes and individual choices, particularly in light of the recent emergence of more immersive forms of media consumption like social media. A greater focus on soft news stories is a major issue given the positive externalities associated with informative news, particularly in terms of the selection of high-quality politicians and ensuring government accountability (Prat and Stromberg, 2005). The results thus raise several questions regarding the media's evolving role within the democratic process, since although the media has often been regarded as a conduit of information between the political class and the people, thus enhancing government accountability (Besley and Burgess, 2001), these results also underscore how even soft news reportage may influence political outcomes.

This paper contributes to the growing literature dealing with the impact of the media on voting, with existing work firmly focussed on the role played by partisan media bias. DellaVigna and Kaplan (2007) analyse how the introduction of Fox News in the U.S. cable market influenced voting patterns in the 2000 Presidential and Senate elections, given Fox's well-documented conservative leanings. The results show a 0.4-0.7 percentage point swing in favour of the Republican Party

in the Presidential elections, as well as a 3-28% gain for Republican candidates in the Senate elections, in those towns where Fox News was launched between October 1996 and November 2000. Similarly, Chiang and Knight (2011) look at how newspaper endorsements of political candidates in the U.S. affect voting behaviour. The results show that the extent to which endorsements are influential in terms of affecting voting patterns largely depends on the perceived credibility of the message; hence endorsements for left-wing candidates are more likely to be influential if the emanate from neutral or right-wing candidates. Conversely, in a field experiment conducted in Virginia prior to the 2005 gubernatorial elections, Gerber, Karlan, and Bergan (2009) find no significant impact of newspaper exposure on individual political behaviour and opinions. As highlighted previously, this is the first paper to explicitly look at the influence of media reports on voting, and more specifically the impact of soft news reportage, abstracting away from all considerations related to partisan bias.

This paper also adds to the political science literature on the impact of scandals on political outcomes, with the results yielding mixed evidence thus far. On the one hand Lau, Sigelman, and Rovner (2007) find that negative or detrimental news has no significant impact on voting patterns, although the results suggest that it may help stimulate interest in the overall campaign. On the other hand, Blackwell (2012) finds that negative political advertising campaigns may prove to be effective if the candidate initiating such a campaign is a challenger, but may backfire if such a campaign is run by an incumbent. Fridkin and Kenney (2011) encapsulate the current impasse on the subject within the literature, stating that the impact of negative stories depends on individual voter characteristics - for example those with a strong political affiliation or interest in the campaign are less likely to be influenced by negativity.

4.1.1 The Lewinsky Scandal

On January 17, 1998 the Drudge Report published the first news report detailing allegations concerning a sexual relationship between Bill Clinton and Monica Lewinsky. The report alleged that the affair was conducted over the period November 1995 to March 1997, and that in 1996 attempts had been made by White House officials to relocate Lewinsky to the Pentagon amidst growing concerns regarding her relationship with the then-President. At first the President forcefully denied any allegations of impropriety, claiming that they were part of a right-wing conspiracy to tarnish his reputation. Despite initial hesitancy, in July 1998 Lewinsky decided to break her silence and testified in front of a grand jury. Faced with mounting evidence, President Clinton recanted his original stance and finally admitted to the affair. He was fined US$90,000 for giving false testimony under oath and had his license to practice law suspended for five years by his home state of Arkansas as well as the Supreme Court in 1999. Impeachment proceedings were launched in December 1998 by the House of Representatives against President Clinton resulting in a 21-day Senate trial, although he was subsequently acquitted of all charges on February 12, 1999.

Public interest in the affair was immense, fueled in part by the salacious nature of the details that emerged from the grand jury investigation as well as the media furore surrounding the trial. As a result, several academic studies have analysed the scandal from various perspectives, including the impact of the scandal on President Clinton's approval ratings (Fischle, 2000) and the media's role in the entire affair (Williams and Carpini, 2000). The sheer scale of the scandal is part of the reason why it was selected as part of this study. In addition, despite its infamy, it is unclear whether the scandal had any tangible impact on the 2000 Presidential elections. In fact as reported in Fischle (2000), during the height of the Lewinsky scandal in 1999 President Clinton's public approval ratings peaked at 73%, and he eventually bowed out of office in 2000 with approval ratings of 68%, the highest recorded for a U.S. President in almost fifty years. Several explanations have been forwarded in order to explain this apparent anomaly, including the fact that the scandal reflected the President's personality traits rather than his performance in Office (Kiousis, 2003). Others (Yioutas and Segvic, 2003) maintain that the media's incessant coverage of the scandal, and in particular its focus on the more salacious sexual details of the affair, only served to alienate the public and discard the information as irrelevant since, according to the authors, people tend to associate sexual or adulterous stories with gossip columns or entertainment rather than political news.

4.1.2 Conceptual Framework

Here we discuss the channels through which media reporting of soft news stories may influence individual voting behaviour. The first mechanism is related to the private benefits accruable from consuming 'hard' news (e.g. details regarding Al Gore's proposed environmental policies) and 'soft' news (e.g. the Lewinsky scandal). More specifically, the defining feature of hard news as described by Prat and Stromberg (2005) is that its consumption generates external benefits for the rest of society since it directly informs people's voting choices, leading to the election of more effective or suitable political candidates. And even though it is also plausible to assume that the consumption of hard news has some private benefit attached to it (e.g. satisfying the reader's desire to be informed or keep up-to-date with current affairs), a significant part of the benefits generated by such news is not internalised by readers. In contrast, soft news only generates private benefits in terms of entertainment, which are fully internalised.

Therefore, apart from the inevitable under-consumption of hard news from a social welfare standpoint (as pointed out by Ellman, 2014), it is also straightforward to see that consumers, faced with a choice between the two news types, will opt for soft news unless the private 'informativeness' benefit from hard news outweighs the entertainment benefit of soft news. This is particularly relevant in the Lewinsky scandal's case, since although it is 'soft' in nature it would still fall under the category of political news given that it dealt with the private life of the then-President, meaning that these stories would have been published in the same section of the newspaper as other hard political articles rather than in the entertainment or gossip section. Therefore, within this context the choice

between hard and soft political news will largely hinge on individual private benefits from either type; in turn, this choice will have a bearing on individual voting decisions.

A second, complementary mechanism considers the variation in people's cognitive capabilities to process different kinds of political news stories. This is based on the idea that, as stated by various authors (e.g. Bagdikian, 2000), on average hard news requires more mental resources in terms of processing and attention time than soft news. This difference in effort/time will further raise the attractiveness of soft political news relative to hard news. Consider an economic agent's choice between consuming 'hard' news stories and 'soft' news. We can express the utility derived from both types of stories as $U_M = u(s)+u(h)$, where $u(s)$ and $u(h)$ denote the private benefits obtained from soft and hard news respectively, and both are concave in s and h. Consumers have a finite amount of time/effort to allocate between hard (h) and soft (s) news, T. We can express this time constraint as $T = s + (1 + r)h$, where r denotes the additional time/processing effort required when reading hard news stories ($r \in [0, 1]$). It is easy to derive the following condition denoting consumption of soft news relative to hard news

$$u'(s) = \tfrac{u'(h)}{1+r} \ .$$

Briefly, this Euler-like equation states that, faced with the above problem, readers will consume soft news until the marginal utility from this type of news is equal to the marginal utility derived from hard news, discounted by the effort parameter r. This means that as r increases, so does the amount of time allocated towards soft news due to the additional effort required to process hard news. Therefore when the cognitive requirements of processing hard political news is high, people will opt for soft political news like Lewinsky scandal coverage, which will in turn inform their voting choices.

This effect would be even more pronounced if we consider that the utility derived from a newspaper U_M depends on consuming a bundle of news topics N that includes political news, sports news, international affairs, etc., with readers having to allocate time T optimally across each topic. In this case, consumption of soft political news makes even more sense since it frees up time to consume other news topics, particularly if the reader's utility varies across each topic in the bundle (for example s/he may prefer sports news to political news). Clearly, the size of r, and hence the reader's allocation of time towards soft news versus hard news, depends on several factors including the reader's level of political engagement/interest, cognitive capabilities (proxied by educational attainment, wages, etc.) and, in the case of the Lewinsky scandal, the importance that s/he attributes to morality or moral behaviour (proxied by religiosity, church attendance, etc.). This latter factor is important since people who attribute a strong importance to moral issues may have been more inclined to read Lewinsky-related stories since they were related to the moral behaviour of President Clinton (in line with the psychology literature on confirmatory information - see Lord, Ross, and Lepper, 1979), and in turn may have been more influenced in terms of their voting choices.

Apart from the externality/cognitive processing mechanism, there are other potential channels through which soft news coverage may influence voting. One alternative is that such stories simply 'crowd out' hard political news, given that newspapers have limited column inches within which to print stories. Thus in this situation it is not the consumer who is actively choosing the soft story, but rather it is the newspaper editor who for various reasons (Hamilton, 2004) is consciously opting to publish such stories at the expense of hard news. Thus, in the absence of any hard news voters would rely on soft political news in order to inform their voting decisions.

This argument can be countered on several fronts. Firstly, although there has been a well-documented rise in the number of soft or sensationalist news stories over the last few years, evidence suggests that this has not necessarily been at the expense of hard news. For example Prior (2003) finds that the number of hard news stories far outweighs that for soft news in the U.S. media (particularly with regards to newspapers), and that the rise in the latter is mainly the result of 24-hour news cycles with space to fill up, while Barnett, Ramsay, and Gaber (2012) find no evidence of a growing trend towards soft reportage in the leading UK news programs on television. Furthermore, Baum (2003) reports evidence to show that reading soft news may actually *increase* the factual political knowledge of politically-disengaged people, postulating that such stories may encourage individuals to read harder, more fact-heavy news stories.

4.2 Data

The data on voting behaviour pertaining to the 2000 U.S. Presidential elections were obtained from the American National Election Survey's (ANES) Time Series Study for the year 2000.[1] The ANES is a collaborative effort between Stanford University and the University of Michigan (with funding from the National Science Foundation), with the expressed aim of providing high-quality data on voting, public opinion and political participation. The ANES conducts several surveys as part of its remit, although the majority of its work is focussed on the Time Series studies which take place during years of national elections in the U.S. (i.e. both Presidential and Congressional elections). The Time Series studies are wide-ranging surveys which seek to gauge individual voting patterns and behaviour for the election in question, together with extensive data on various socio-economic and demographic characteristics, as well as other variables

[1]Another possible option would have been the 2000 National Annenberg Election Survey (NAES) conducted by the Annenberg Public Policy Center at the University of Pennsylvania, which is a large-scale public opinion survey related to U.S. presidential elections. Nonetheless, it is important to point out that the NAES is largely aimed at studying the dynamics of public opinion during the course of an electoral campaign, most of the survey is focussed on the pre-election months, with only around 7,000 interviews conducted after the election date. Taking into account data omissions and other factors, like availability of newspaper archives, it was decided to opt for the ANES Time Series Study due to its larger eventual sample size and improved reliability in terms of its national representativeness, given that this paper is primarily concerned with actual voting behaviour rather than voting intentions.

pertaining to respondents' values, predispositions, social ties and religious beliefs. The 2000 ANES Time Series study was conducted between September and December 2000, utilising a statistically representative sample of 1,807 households from across the country. The study consisted of two phases - the pre-election survey, which was carried out between September 5 and November 6 (prior to the election on November 7) and the post-election survey, which was held between November 8 and December 18. For the purposes of this study, we are mainly interested in the latter phase, since the focus is on actual voting behaviour rather than voting intentions.[2]

Hence in total our final sample, after accounting for data omissions and other shortcomings, consists of 1,555 households who participated in the post-election survey. One interesting feature of the 2000 Time Series study is that, apart from the standard variables found in each edition of the survey, it also includes specific questions related to respondents' perceptions regarding Bill Clinton and his legacy, including job approval as well as his impact on the country's moral climate, although the survey avoids any direct reference to Monica Lewinsky or the scandal. Figure 4.1 shows the voting patterns across our sample. As seen below, it appears that our sample is skewed slightly in favour of Al Gore given that although the Democratic Party did win the popular vote, the gap between the two parties was of around 0.5 percentage points, in contrast to the difference of around 4 percentage points reported in the dataset. This may partially be due to the higher reported turnout in our post-election dataset, where 75% of respondents claimed to have voted in the Presidential election, which is somewhat higher than the actual 67.5% turnout estimate.

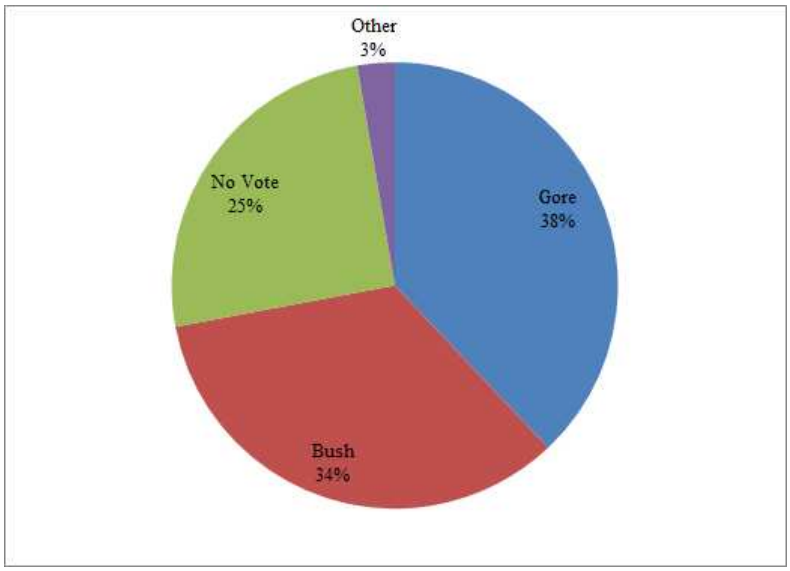

Figure 4.1: ANES 2000 Voting Patterns in U.S. Presidential Election (Source: ANES, 2000)

Data on the number of published newspaper articles related to the Monica

[2]As a robustness check we include the pre-election observations in our dataset, yielding markedly-similar results to those reported in the subsequent sections.

Lewinsky scandal over the period January 17, 1998 to August 31, 2000 were collated using three main resources - NewsBank's NewsLibrary website, which is a searchable online repository of newspaper articles in the U.S. covering over 4,000 newspapers across the country, the ProQuest database, as well as the individual newspapers' print archives.[3] This approach was employed in order to maximise the reliability and accuracy of the data; in case of a discrepancy, the source with the highest number of verifiable articles was selected. The time period was selected since as mentioned earlier the first mainstream media news story pertaining to the scandal broke on January 17, 1998, while the end-date coincides with the commencement of the pre-election ANES 2000 survey period. For the purposes of searching for relevant Lewinsky scandal news stories in each newspaper, articles containing the words "Monica Lewinsky" or "Lewinsky" or "Lewinskygate" or "Monicagate" or "tailgate" or "sexgate" or "zippergate" or "Clinton" in the headline, and "Monica Lewinsky" or "Lewinsky" or "Lewinskygate" or "Monicagate" or "tailgate" or "sexgate" or "zippergate" in the main body of text were sought out, excluding those articles containing the words "letter" or "letters", over the period January 17, 1998 to August 31, 2000.[4]

This entire process was conducted for 257 daily newspapers from across the country, including national, state-wide and local newspapers. Figure 4.2 shows the top 50 newspapers on the basis of the number of Lewinsky scandal news stories published over the period under review. As expected the large newspapers lead the way, with the New York Times topping the list with 867 news stories, followed by the Boston Globe with 768 articles and the Milwaukee Journal Sentinel with 709 articles.

The choice of newspapers was largely dictated by the demographic characteristics of the survey respondents in the 2000 ANES Times Series study, which records both the State and the congressional district associated with each respondent. For each congressional district, we identified (wherever possible) the top three daily newspapers by circulation, using data compiled by the Alliance for Audited Media (2000). As mentioned above, this yields a total of 257 newspapers which collectively account for over 70% of annual newspaper circulation in the whole country. To obtain the final figure for the number of news articles per congressional district (for each State) a weighted-average was taken based on the relative population of each newspaper's target counties.

An example will help to clarify matters. Consider the State of California, where the top three circulating newspapers in the 48th congressional district as defined in the year 2000 were the Orange County Register (circulation of 250,724 copies per week), the San Diego Union-Tribune (296,331 copies per week) and the Riverside Press-Enterprise (172,593 copies per week). In turn, these newspapers' main counties of circulation within this district were Orange County (population of 1,505,116 inhabitants), San Diego (413,493 inhabitants) and Riverside (100,097 inhabitants). The number of Lewinsky scandal news stories published in each of

[3]NewsLibrary and ProQuest have been used as the principal sources of news-related data in various other studies, like for example Gentzkow and Shapiro (2010).

[4]For a more detailed explanation of the data collection process undertaken for this paper, refer to Appendix C.

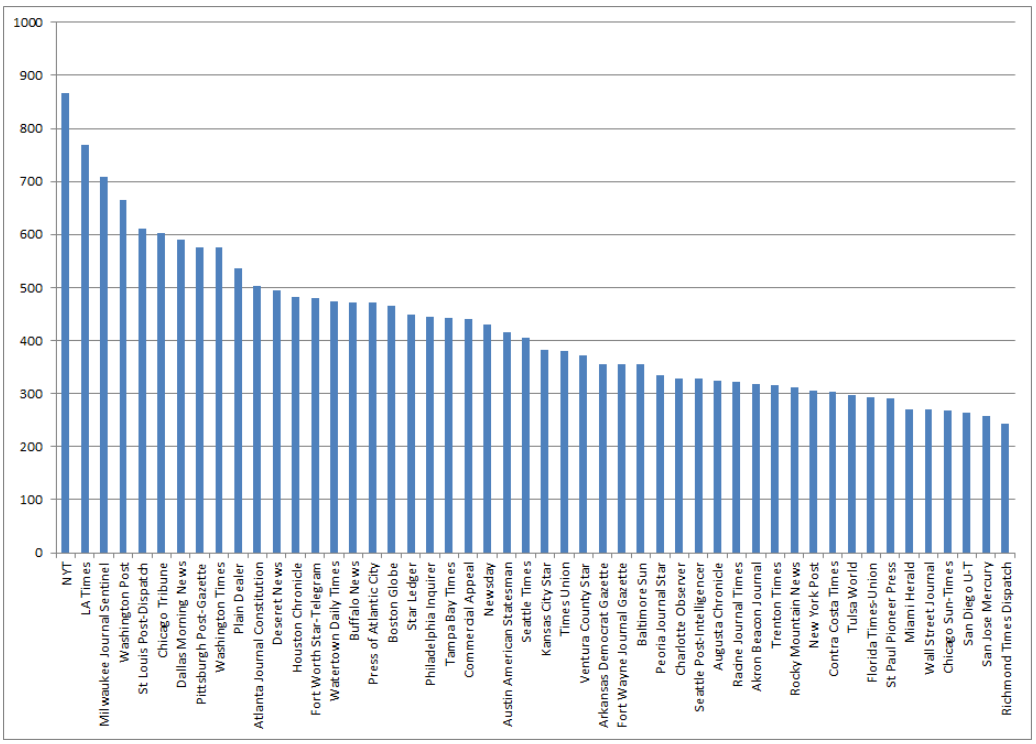

Figure 4.2: Top 50 US Daily Newspapers by Number of Lewinsky Scandal News Stories Published from January 17, 1998 to August 31, 2000 (Source: NewsLibrary, ProQuest, newspaper print archives, 2014)

these newspapers over the period January 17, 1998 to August 31, 2000 was 227, 264 and 110 respectively. Hence the weighted average number of Lewinsky news stories for California's 48th congressional district in this sample is 228.8 articles.[5]

Table 4.1 presents a brief summary of the main descriptive statistics for each of the key variables described above. In total 38% of respondents in the ANES 2000 Time Series study reported voting for Democratic candidate Al Gore, with 34% opting for eventual winner George W. Bush while just over 24% of respondents abstained from voting. On average around 226 news stories related to the Lewinsky scandal were published in each newspaper over the period January 17, 1998 to August 31, 2000, although this value varied considerably as seen from the quantitatively large standard deviation obtained.

4.3 The Determinants of Media Reporting Behaviour

In this section, we take a closer look at what influenced the level of coverage afforded by each newspaper to the Lewinsky scandal over the period under review.

[5]As part of the robustness tests we also consider an alternative weighting system based on each newspaper's relative circulation figures within each congressional district rather than population; the results obtained are quantitatively very similar.

Table 4.1: Summary statistics

Variable	Mean	Std. Dev.	Min.	Max.
% Votes for Al Gore	0.38	0.485	0	1
% Votes for George W. Bush	0.341	0.474	0	1
% Abstentions	0.242	0.429	0	1
No. of Lewinsky scandal articles	226.274	203.06	5.009	867
N		1552		

Notes: This table shows summary statistics for the two main variables of interest in this paper. Data on individual voting patterns are from the 2000 American National Election Survey (ANES), while the number of newspaper articles related to the Lewinsky scandal was collated using NewsLibrary, ProQuest and the individual newspapers' online news archives.

These results are crucial for the subsequent analysis since they enable us to eliminate any impact of partisan/political slant in the media (which is the subject of other related papers like DellaVigna and Kaplan, 2007) and focus exclusively on the effects of soft news reportage on voting behaviour. Moreover, by ensuring that Lewinsky scandal coverage is independent of partisan bias, we are also able to directly tackle concerns that any relationship between scandal reportage and voting patterns is being driven by a higher general level of anti-Democrat or pro-Republican news stories over the period under review. Given that such reporting behaviour is typical of right-wing newspapers as discussed by various authors (Gentzkow and Shapiro, 2010), it is therefore possible to uniquely identify the effect of media reportage (within the context of Lewinsky scandal coverage) on voting behaviour, irrespective of newspaper partisan slant.

Gentzkow and Shapiro (2008) identify two broad categories of determinants of media reporting behaviour and political bias, namely demand side factors and supply side factors. On the demand side, it is argued that newspapers simply tailor their reportage according to the opinions or beliefs of their target audience; thus political media bias is a direct consequence of biased readers (Mullainathan and Shleifer, 2005). On the other hand, supply side factors are related to the beliefs and/or preferences of journalists/editors (Baron, 2006), advertisers (Reuter and Zitzewitz, 2006) and owners (Germano and Meier, 2010). Within this context, it is possible to directly assess the main demand and supply side factors for each newspaper and their influence on the level of Lewinsky scandal coverage over the period January 17, 1998 to August 31, 2000. The motivation behind this analysis stems from the fact that although the Lewinsky scandal had a clear political element to it since Bill Clinton was a Democrat President, the somewhat tawdry nature of the story also lends itself to heightened media coverage across the board, irrespective of political slant (Sparks, Tulloch, et al., 2000). Therefore, the analysis may help to shed light on whether the newspapers' handling of the scandal was consistent with the demand and/or supply side factors associated with political media slant or else due to other factors like sensationalism.

The results are shown in Table 4.2. In column 1 we regress the amount of coverage afforded by each newspaper to the Lewinsky scandal [6] on the "canonical"

[6]Whereby coverage is denoted as the number of Lewinsky-related articles published over the period January 17, 1998 to August 31, 2000 by each newspaper, divided by the total number of articles published by each paper over that same period.

demand and supply side factors determining partisan media bias, namely readers' beliefs (in this case the proportion of people in each newspaper's home county who voted for Republican candidate Bob Dole in the 1996 Presidential election), ownership (where the variables Gannett, Knight Ridder and Advance refer to the top three media owners in our sample in the year 2000) and newspaper editorial slant (using data from Groseclose and Milyo, 2005). This latter variable is an index ranging from 0-100 that compares the number of times a newspaper refers to a particular think-tank to the number of times these think-tanks are cited by members of Congress with known political inclinations based on adjusted Americans for Democratic Action (ADA) scores for partisanship estimated in Groseclose, Levitt, and Snyder (1999) for the period 1993 to 1999. Low scores (below 50) denote that a newspaper has a conservative bias, while higher scores (over 50) denote a liberal slant.

As seen from the results, none of our explanatory variables yield statistically significant results. Column 2 introduces a number of additional control variables, including geographical location, circulation figures for each newspaper, a dummy denoting whether the newspaper is a broadsheet or not, and a proxy for readers' interest in political matters (as measured by the voter turnout rate in each newspaper's home county during the 2000 Presidential election). Once again, the impact of demand/supply side factors is still non-existent. In turn, Lewinsky coverage is positively-related to circulation and the East South-Central region of the U.S. (which includes Alabama, Kentucky, Mississippi and Tennessee), and negatively-related to the broadsheet dummy. These results suggest that coverage of the Lewinsky scandal may not have been due to some desire by newspapers to bias or slant their news reports in one direction or another, but rather fueled by other considerations. In fact the negative coefficient obtained for the broadsheet variable supports this latter argument, since it indicates that broadsheet newspapers' coverage of the scandal lagged behind that of tabloid newspapers, whose predilections for sensationalism are well-documented (Sparks, Tulloch, et al., 2000).

4.4 Media Reports and Voting Patterns

Having shown that the level of newspaper coverage afforded to the Lewinsky scandal was independent of political bias considerations, we now turn to analysing the extent to which reportage of the scandal influenced voting behaviour in the 2000 U.S. Presidential election. In this section, our primary outcome (dependent) variables of interest are the share of votes acquired by both leading candidates in the election. Let G_i refer to a binary variable indicating whether respondent i voted for Al Gore or not as observed in our data (where $G_i \in \{0, 1\}$), while G_i^* is the corresponding latent variable based on the unobserved propensity to vote for Gore. We can therefore express the individual's decision to vote for Al Gore or not in terms of the following latent variable formulation:

$$G_i^* = \beta A_i + \mathbf{X}_i'\alpha + \varepsilon_i$$

Table 4.2: Determinants of Media Reporting Behavior

	Lewinsky Coverage	
1996 Republican Votes	-0.0110	-0.00979
	(0.0115)	(0.0116)
Editorial Slant	0.00261	-0.0103
	(0.0144)	(0.0160)
Ownership: Gannett	-0.255	-0.318
	(0.240)	(0.261)
Ownership: Knight Ridder	0.401	0.401
	(0.252)	(0.272)
Ownership: Advance	-0.429	-0.425
	(0.339)	(0.331)
Circulation		0.000000601
		(0.000000368)
New England		-0.535
		(0.438)
Mid Atlantic		-0.157
		(0.329)
East North Central		-0.219
		(0.300)
West North Central		-0.182
		(0.381)
East South Central		0.865**
		(0.429)
West South Central		0.0678
		(0.295)
Mountain		0.110
		(0.534)
Pacific		-0.00305
		(0.305)
Broadsheet		-1.570*
		(0.826)
Voter Turnout		-0.0771
		(0.0517)
Constant	2.265*	4.500***
	(1.151)	(1.657)
Observations	200	200
R^2	0.032	0.125

Notes: In this table we regress newspaper coverage of the Lewinsky scandal on the key determinants of partisan media bias proposed in the literature, namely the political leanings of the newspapers' readership, the editorial slant of each newspaper, as well as the political bias of the owners (where Gannett, Knight Ridder and Advance Publications are the three largest newspaper owners in our sample). In Column 2 we also include a number of additional control variables like newspaper circulation, voter turnout and a dummy denoting whether a newspaper is a broadsheet (as opposed to being a tabloid). In both columns, 'coverage' is denoted as the number of Lewinsky-related articles published over the period January 17, 1998 to August 31, 2000 by each newspaper, divided by the total number of articles published by each paper over that same period.All data sources are listed in Appendix A1. Robust standard errors are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

$$G_i = \begin{cases} 1 & \text{if } G_i^* > 0 \\ 0 & \text{if } G_i^* \leq 0 \end{cases}$$

where A_i is the amount of coverage (per 1,000 articles) afforded to the Lewinsky scandal by the leading newspapers in respondent i's congressional district over the period January 17, 1998 to August 31, 2000 (which is calculated by expressing the number of Lewinsky articles published as a fraction total number of news stories published in each newspaper over the period under review), $\mathbf{X}_i$ is a vector of covariates[7] and ε_i is a random error term. We divide the number of Lewinsky articles by the total number of published news stories in each newspaper in order to properly gauge the amount of coverage given to the Lewinsky scandal by each newspaper relative to its size (in terms of column inches available). This is due to the fact that newspaper size varies considerably across the country, depending on the format of the newspaper, editorial/advertising decisions and circulation figures, meaning that the raw number of published Lewinsky articles may not be a fair reflection of the actual amount of coverage or importance afforded to the scandal by each paper.

Similarly, if we denote B_i as a binary variable indicating whether respondent i voted for George W. Bush or not as observed in our data, with B_i^* being the corresponding latent variable based on the unobserved propensity to vote for Bush, then we have the following specification:

$$B_i^* = \beta A_i + \mathbf{X}_i'\alpha + \upsilon_i$$
$$B_i = \begin{cases} 1 & \text{if } B_i^* > 0 \\ 0 & \text{if } B_i^* \leq 0 \end{cases}$$

where υ_i is a random error term. Under the assumption that both ε_i and υ_i are normally and independently distributed with zero mean and a constant variance, the above two expressions can be estimated consistently using a Probit model.[8]

The basic reduced-form equations estimated in this paper are thus specified as follows:

$$\left.\begin{matrix} G_i \\ B_i \end{matrix}\right\} = \mu + \beta A_i + \mathbf{X}_i'\alpha + \varepsilon_i \tag{4.1}$$

We estimate (4.1) separately for each of the two binary explanatory variables mentioned above using a Probit model.[9]

[7] A full description of the control variables used, including summary statistics, is provided in Appendices A1 and A2.

[8] In Section 4.4.3 we use a Probit IV model in order to account for endogeneity issues, particularly measurement error bias.

[9] Logit estimates were also produced for (4.1), yielding very similar results. A multinomial logit model was also estimated for (4.1) using the combined voting patterns of the survey respondents as our dependent variable. Once again, the results are quantitatively close to those obtained using separate Probit/Logit specifications. These results have been omitted due to space limitations.

4.4.1 Basic Regression Results

Table 4.3 reports the initial regression results for (4.1). Columns 1 and 2 report the regression results for the whole sample, with both G_i and B_i as dependent variables, while the rest of the table deals with different subsamples of the data in order to gauge the impact of Lewinsky-related media coverage on different respondent subgroups. Columns 3 and 4 estimate(4.1) for both dependent variables, focussing solely on those respondents who identified themselves as non-Republicans - i.e. Democrats, Independents and others. The idea behind this specification is to acquire an understanding of the media coverage-voting relationship for non-Republicans, since by definition this group would have been ex ante less likely to vote for George W. Bush. By the same token, columns 5 and 6 focus on respondents who identify themselves as non-Democrats, i.e. Republicans, Independents and others.[10] In each case, the explanatory variable of interest for the purposes of this paper is the one denoting the amount of coverage afforded to the Lewinsky scandal.

The regressions also contain various control variables related to each respondent, including geographical location, educational background, income levels, employment status, gender, marital status, age, race, religion, prior voting patterns in both the 1996 Presidential election as well as the 1998 House of Representatives election, perceptions of each Presidential candidate and opinions regarding Bill Clinton. We also include a variable denoting the average circulation figure of the top three newspapers within each survey respondent's congressional district of residence, in order to account for the fact that high-circulating newspapers afforded more coverage to the Lewinsky scandal relative to more local papers which tend to focus more on regional/micro-level news rather than national or international stories. Due to space limitations, in Table 4.3 (and indeed throughout the rest of this paper) we only present the coefficients of those control variables that yielded statistically significant results across the majority of specifications; for a complete list of all control variables used, consult Appendix A1.

As seen from the results, coverage of the Lewinsky scandal is positively and significantly correlated with the probability of voting for Bush. A closer look at the table suggests that this result is mainly being driven by non-Republicans, as shown in Column 3. This implies that newspaper coverage of the Lewinsky scandal may have influenced the voting behaviour of non-Republicans - who were a priori less inclined to vote for Bush - by increasing the probability of voting for the Republican candidate. By contrast, Lewinsky coverage is negatively correlated with the likelihood of a Gore vote, although the coefficients are imprecisely estimated. In terms of the other explanatory variables, it appears as though geographical location and age did not have any significant impact on voting patterns, although prior voting behaviour, Clinton perceptions and religion all yielded statistically significant results across most specifications.[11]

[10] Later on we also use a different subgroup classification systems based on respondents' voting patterns during the 1996 Presidential election, obtaining very similar results to those shown in the rest of this paper.

[11] In fact the results show that Protestants and Catholics were more likely to vote for Bush

Table 4.3: Basic Regression Results

	Full Sample		Non-Republicans		Non-Democrats	
	Bush	Gore	Bush	Gore	Bush	Gore
Lewinsky Coverage	0.0614*	-0.0168	0.121**	-0.0136	0.0569	-0.0492
	(0.0334)	(0.0406)	(0.0525)	(0.0448)	(0.0469)	(0.0602)
1998 House Republican	1.323***	-0.164	1.150***	-0.0743	1.342***	-0.112
	(0.110)	(0.112)	(0.136)	(0.134)	(0.123)	(0.144)
1996 Clinton	-0.271**	1.036***	-0.315**	1.040***	-0.0504	0.800***
	(0.132)	(0.0976)	(0.146)	(0.103)	(0.152)	(0.140)
1996 Dole	1.134***	-0.508***	0.992***	-0.440**	1.129***	-0.552***
	(0.136)	(0.159)	(0.185)	(0.197)	(0.142)	(0.199)
Like Gore	-0.709***	0.814***	-0.718***	0.721***	-0.659***	0.946***
	(0.113)	(0.0889)	(0.138)	(0.0980)	(0.128)	(0.119)
Like Bush	0.878***	-0.645***	0.808***	-0.551***	0.789***	-0.564***
	(0.106)	(0.0958)	(0.123)	(0.104)	(0.119)	(0.128)
Approve Clinton	-0.257**	0.447***	-0.199	0.423***	-0.222*	0.405***
	(0.113)	(0.111)	(0.136)	(0.122)	(0.120)	(0.136)
Clinton Moral Climate	0.297**	-0.363***	0.371**	-0.398***	0.202	-0.470***
	(0.120)	(0.117)	(0.151)	(0.131)	(0.132)	(0.168)
Protestant	0.342***	0.0973	0.293*	0.217*	0.338**	-0.0210
	(0.124)	(0.113)	(0.156)	(0.121)	(0.136)	(0.159)
Catholic	0.438***	0.141	0.430**	0.196	0.444***	0.0465
	(0.147)	(0.123)	(0.173)	(0.131)	(0.167)	(0.168)
Marital Status	0.207*	0.121	0.249*	0.146	0.288**	-0.00920
	(0.107)	(0.0972)	(0.130)	(0.106)	(0.123)	(0.139)
Advanced Education	-0.176	0.344***	-0.194	0.414***	-0.215*	0.312**
	(0.113)	(0.0993)	(0.136)	(0.109)	(0.128)	(0.134)
Dropout	-0.223	-0.497***	-0.178	-0.474***	-0.414**	-0.676**
	(0.178)	(0.151)	(0.212)	(0.157)	(0.191)	(0.263)
White	0.323**	-0.292**	0.391***	-0.311***	0.190	-0.137
	(0.128)	(0.116)	(0.150)	(0.120)	(0.149)	(0.182)
Constant	-1.279***	-1.140***	-1.288***	-1.009***	-1.071**	-1.532***
	(0.380)	(0.334)	(0.426)	(0.369)	(0.429)	(0.465)
Observations	1552	1552	1151	1151	1028	1028
Pseudo R^2	0.580	0.441	0.480	0.368	0.522	0.398

Notes: In this table we regress individual voting patterns during the 2000 U.S. Presidential election as reported in the 2000 ANES on the level of newspaper coverage afforded to the Lewinsky scandal by the leading 2-3 newspapers (by circulation) within each survey respondent's congressional district, as well as a host of control variables (in the table above we only include those controls which yield significant results; a full list of controls is provided in Appendix A1). We use two dependent variables - one is a dummy indicating whether the respondent voted for George W. Bush or otherwise, and the other denotes a vote for Al Gore (or otherwise). Columns 1 and 2 refer to the regression described above using the full sample; in columns 3 and 4 we focus only on the subsample of respondents whose political affiliation as recorded in the ANES is non-Republican (i.e. Democrat, Independent or other); while in columns 5 and 6 we analyze the subsample of respondents whose affiliation is non-Democrat (i.e. Republican, Independent or other). All data sources are listed in Appendix A1. Robust standard errors, clustered at the congressional district level, are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Nonetheless, we must also account for potential endogeneity in the regression results presented above. This may be a cause for concern since it may reasonably be argued that media coverage of the Lewinsky scandal is merely a reflection of a general, non-partisan rise in anti-Clinton/Democrat sentiment during the President's second term in office, particularly in light of the sizeable list of scandals associated with the Clinton administration. This may have led newspaper editors and journalists to increase their coverage of the Lewinsky scandal in order to account for this shift in public opinion, regardless of the individual newspaper's inherent political slant. Therefore, it could be the case that both Lewinsky coverage and individual voting patterns are being driven by this (omitted) general, cross-party antipathy towards the Democrats (an extreme form of the "Clinton Fatigue" idea described by Geer (2004)). The inclusion of various control variables related to respondents' approval of Bill Clinton and the two Presidential candidates was intended to mitigate against the potential presence of this problem.

Of greater concern is the likely presence of measurement errors which may bias our results downwards. This is because although we have implemented various measures to ensure the reliability of the media coverage data, it is not possible to identify the actual position of the news article within the newspaper. Plenty of research has been done in terms of analysing the likelihood of reading a news article based on its position in the newspaper, both in terms of page number (Bogart, 1989) as well its relative position on the actual page (Fico and Cote, 1999). For example, apart from the obvious front-page news, stories that are located in the first 5-6 pages of a newspaper are more likely to be read; the same thing can also be said for articles printed at the top of the page rather than towards the bottom. This means that our Lewinsky coverage variable contains several stories which have a statistically low probability of having been read due to their unfavourable position within the newspaper, which in turn would limit their likely impact on voting patterns.

4.4.2 Identification Strategy

The instrumental variables (IV) approach developed in this paper revolves around finding a suitable variable or instrument which can explain the variation in Lewinsky scandal media coverage (A_i) but is uncorrelated with the probability of voting for either Bush or Gore. With this in mind, the identification strategy employed is based on variations in the number of people who died, over the period January 17, 1998 to August 31, 2000, as a result of extreme weather incidences (at the county level).[12] The selection of this instrument stems from the fact that a typical newspaper has a limited amount of space in which to publish news stories, with several potential articles vying for publication. The battle for column inches is

than atheists or respondents from other religious denominations. By contrast, religion did not have a statistically significant effect in terms of the probability of voting for Al Gore.

[12]Extreme weather incidents, as defined by the U.S. National Oceanic and Atmospheric Administration (NOAA), include various events like hailstorms, excessive heat, hurricanes, tornadoes, thunderstorm winds, drought and flash floods.

perhaps even more critical in the presence of an incendiary news story like the Lewinsky scandal which dominated headlines for several months, as attested by the data presented in Figure 4.2.[13]

It would therefore take a relatively major local incident or news event within the media outlet's circulating area in order to usurp any potential reportage on the Lewinsky scandal, or at the very least relegate it to 'News in Brief' status. In this regard, extreme weather events appear to fit the description since they are typically localised incidents and have significant public interest due to their widespread ramifications in terms of deaths, injuries and/or damages for the local community. It therefore comes as no surprise that extreme weather events are one of the most widely-reported news topics in the U.S., with coverage levels on the rise. Ungar (1999) looks at media coverage of extreme weather incidents in the U.S. over the period 1968 and 1996, finding a significant positive trend over time (independent of the actual number of weather incidents recorded), postulating that this may be due to increased public awareness of climate change.

Therefore, based on the above discussion, we use the number of extreme weather-related deaths over the period January 17, 1998 to August 31, 2000 as our instrument of choice. By concentrating solely on deaths we automatically eliminate any minor weather incidents which may not have been particularly newsworthy and instead focus on major weather events. Ex-ante, the number of deaths from extreme weather events is expected to be negatively-related to newspaper coverage of the Lewinsky scandal, since as explained earlier extensive casualties as a result of extreme weather may have led newspapers to focus more on such stories at the expense of other articles, including the Lewinsky scandal. We also contend that the only way in which the selected instrument is related to the respondents' voting patterns is through the Lewinsky coverage variable and not via other channels. It may be argued that people living in areas which registered a high death toll from extreme weather events may have been more inclined to vote for Bush rather than Gore since the casualty rate may have been a reflection of poorly-managed or inadequate protection and relief efforts by the government - accusations which were, somewhat ironically, leveled at President Bush himself in the aftermath of Hurricane Katrina in late August/September 2005 (Hartman and Squires, 2006). However, this assertion can be countered on three fronts:

- It is entirely plausible that, in the event of a high death toll caused by a perceived lack of government preparation or a slow response to the crisis, people would apportion the majority of the blame to local/state governments rather than the Federal government. In fact, despite the criticism received by the Bush administration in the wake of Hurricane Katrina, a CNN/Gallup poll conducted on September 8, 2005 showed that 25% of people blamed state and local governments for the devastation wrought by the storm, compared to 13% who blamed President Bush;

[13] In fact, a Gallup poll (1998) reported that 72% of respondents believed that the U.S. media afforded too much coverage to the Lewinsky scandal.

- It is unclear whether a high casualty rate would favour either candidate, since in reality this may simply reflect the scale of the disaster rather than any incompetence from the government's end. Healy and Malhotra (2010) find no correlation between voting patterns and tornado-caused deaths in the U.S.; instead, they argue that people may either reward or punish incumbent governments on the basis of their perceived handling of the disaster.

Another key feature of our extreme weather deaths instrument is that its ubiquity in the news media helps us to directly address the problem of measurement error caused by a lack of information regarding the positioning of Lewinsky articles in newspapers. A Pew Research Center 2012a study shows that weather events are the most popular news topic among U.S. citizens, with 52% of people following such stories very closely, almost double the level of interest shown in the next most popular topics (including crime and sports stories). This popularity of extreme weather events has (as mentioned earlier) unsurprisingly resulted in greater media attention over time, both in terms of the actual coverage (Ungar, 1999) as well as its prime positioning in newspapers and TV news broadcasts (Sturken, 2001). Therefore, by using extreme weather deaths as our instrument of choice we are explicitly identifying those Lewinsky stories that feature prominently in the local newspapers, and which would be in competition for prime column inches with extreme weather stories, while the other non-prime Lewinsky stories would be eliminated from the reckoning since they would not exhibit any relationship with extreme weather stories. This would in turn result in higher coefficients for our Lewinsky coverage variable when using IV relative to those obtained using standard Probit/Logit in Table 4.3.

Table 4.4 Panel B reports the first-stage results of the linear regression of Lewinsky-related newspaper coverage on the number of weather-related deaths over the period January 17, 1998 to August 31, 2000. To construct our instrument, we first obtain data on casualties related to extreme weather instances from the NOAA for each newspaper's main county/ies of circulation within each congressional district in our ANES dataset. We then obtain a weighted average of the number of weather-related deaths per congressional district using the same weighting method utilised in the previous section, based on the relative population size of each county.[14] As expected, weather-related deaths are negatively and significantly related to the number of Lewinsky news articles, with a unitary increase in weather casualties associated with a 1% decline in newspaper coverage of the scandal.

[14]To further assess the validity of our instrument, we also use *changes* in extreme weather deaths in each newspaper's main circulating county/ies, relative to the average in the previous 20 years. The results obtained are almost identical to those presented below, and are omitted due to space limitations.

4.4.3 Instrumental Variables (IV) Results

Table 4.4 Panel A shows the Probit IV results obtained for (4.1), using the number of extreme-weather deaths as our instrument of choice.[15] We start with column 3, where once again we observe a positive and significant relationship between Lewinsky coverage and the likelihood of voting for George Bush among non-Republicans. The estimates are also economically important - the computed marginal effects (supplied in Appendix B) suggest that a 5% increase in newspaper coverage of the Lewinsky scandal could have increased the likelihood that a non-Republican voted for Bush in the 2000 Presidential election by around 1.9 percentage points. The magnitude of these coefficients is significantly higher than those obtained under standard Probit/Logit in Table 4.3, which confirms the presence of measurement error in the latter set of estimates which has now been dealt with through the use of our extreme weather deaths instrument.

Interestingly, the IV results corresponding to voting patterns for Al Gore among non-Republicans (column 4) this time also yield significant results: in fact, they indicate that a 5% increase in Lewinsky coverage over the period under review could have *decreased* the likelihood of non-Republicans voting for Al Gore by approximately 1.8 percentage points (based on the marginal effects presented in Appendix B), almost equivalent to the positive swing in Bush's favour resulting from an identical increase in news reportage. Similar results can also be observed across the entire dataset (columns 1 and 2), although it appears as though media coverage did not have any impact on the voting patterns of non-Democrats (columns 5 and 6), confirming what was observed in Table 4.3 previously.

Therefore the results obtained from the Probit IV analysis imply that Lewinsky scandal newspaper coverage had an important impact on individual voting behaviour during the 2000 U.S. Presidential elections by increasing the likelihood of a Bush vote at the expense of Al Gore, particularly among non-Republicans. When we consider the fact that, as mentioned earlier, the actual gap between the proportion of the popular vote obtained by both candidates was of 0.5 percentage points, then it becomes clear that the media could have played a crucial role in deciding the final outcome of the election, particularly given the controversy that erupted in the aftermath of such a close result with the Florida recount.

In Table 4.5 we further sub-divide our dataset according to whether respondents identify themselves as Republicans, Democrats or Independents, in order to further analyse the results shown in Table 4.4 which are broader in scope. We can clearly observe that Lewinsky scandal coverage had an important impact on voting patterns for both Democrats and Republicans (columns 1 to 4), where in either case the effect was to reduce the likelihood of voting for Gore while raising the probability of a Bush vote, as obtained above. By contrast, Independents (columns 5 and 6) were not influenced in their voting decisions by Lewinsky-related media reportage.

A potential cause for concern may be that our current subgroup classifications

[15] Once again, due to space constraints we only present the estimates for those control variables which yielded statistically significant results across most specifications.

Table 4.4: IV Regression Results

Panel A	Full Sample		Non-Republicans		Non-Democrats	
	Bush	Gore	Bush	Gore	Bush	Gore
Lewinsky Coverage	0.645***	-0.497**	0.784***	-0.463**	0.492	-0.474
	(0.184)	(0.195)	(0.142)	(0.204)	(0.349)	(0.307)
1998 House Republican	0.960***	-0.0975	0.743***	-0.0293	1.138***	-0.0619
	(0.266)	(0.112)	(0.236)	(0.131)	(0.340)	(0.144)
1996 Clinton	-0.256**	0.918***	-0.249*	0.938***	-0.136	0.790***
	(0.116)	(0.152)	(0.127)	(0.149)	(0.150)	(0.157)
1996 Dole	0.735***	-0.338*	0.514**	-0.284	0.888**	-0.398
	(0.265)	(0.174)	(0.249)	(0.199)	(0.351)	(0.245)
Like Gore	-0.561***	0.704***	-0.536***	0.657***	-0.577***	0.831***
	(0.151)	(0.130)	(0.154)	(0.117)	(0.182)	(0.206)
Like Bush	0.677***	-0.551***	0.579***	-0.495***	0.704***	-0.506***
	(0.182)	(0.121)	(0.167)	(0.112)	(0.200)	(0.152)
Approve Clinton	-0.171	0.360***	-0.130	0.370***	-0.142	0.305*
	(0.104)	(0.121)	(0.108)	(0.124)	(0.136)	(0.163)
Clinton Moral Climate	0.294***	-0.371***	0.316**	-0.401***	0.205*	-0.451***
	(0.105)	(0.109)	(0.128)	(0.122)	(0.123)	(0.167)
Protestant	0.355***	0.0207	0.268**	0.169	0.387***	-0.0965
	(0.111)	(0.111)	(0.130)	(0.120)	(0.128)	(0.162)
Catholic	0.371***	0.0931	0.249	0.209*	0.489***	-0.0580
	(0.139)	(0.115)	(0.159)	(0.121)	(0.153)	(0.173)
Marital Status	0.149	0.105	0.168	0.127	0.232	0.00950
	(0.0990)	(0.0880)	(0.112)	(0.0990)	(0.141)	(0.128)
Advanced Education	-0.0882	0.261**	-0.0820	0.335***	-0.150	0.243*
	(0.102)	(0.107)	(0.112)	(0.119)	(0.138)	(0.143)
Dropout	-0.139	-0.451***	-0.119	-0.421***	-0.272	-0.694***
	(0.159)	(0.147)	(0.164)	(0.157)	(0.256)	(0.254)
White	0.383***	-0.342***	0.391***	-0.335***	0.258*	-0.195
	(0.105)	(0.104)	(0.122)	(0.113)	(0.140)	(0.165)
Constant	-2.376***	0.142	-2.197***	-0.0719	-1.990***	-0.342
	(0.390)	(0.673)	(0.359)	(0.602)	(0.737)	(1.075)
Panel B						
Weather Deaths	-0.0179***		-0.0176***		-0.0175***	
	(0.00221)		(0.00243)		(0.00323)	
Constant	2.128***		2.115***		2.091***	
	(0.0317)		(0.0354)		(0.0392)	
F-statistic	65.72		52.18		29.27	
Observations	1552	1552	1151	1151	1028	1028

Notes: Panel B reports the first-stage results from our instrumental variables approach, whereby we regress newspaper coverage of the Lewinsky scandal on the number of deaths caused by extreme weather events in each of the ANES respondent's congressional districts. In Panel A we report the second-stage results from our Probit IV analysis, where we regress individual voting patterns during the 2000 U.S. Presidential election as reported in the 2000 ANES on Lewinsky coverage (instrumented by weather deaths), as well as a host of control variables (in the table above we only include those controls which yield significant results; a full list of controls is provided in Appendix A1). We use two dependent variables - one is a dummy indicating whether the respondent voted for George W. Bush or otherwise, and the other denotes a vote for Al Gore (or otherwise). Columns 1 and 2 refer to the regression described above using the full sample; in columns 3 and 4 we focus only on the subsample of respondents whose political affiliation as recorded in the ANES is non-Republican (i.e. Democrat, Independent or other); while in columns 5 and 6 we analyze the subsample of respondents whose affiliation is non-Democrat (i.e. Republican, Independent or other). All data sources are listed in Appendix A1. Robust standard errors, clustered at the congressional district level, are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4.5: IV Regression Results - Alternative Subgroups

	Democrats		Republicans		Independents	
	Bush	Gore	Bush	Gore	Bush	Gore
Lewinsky Coverage	0.956***	-0.621**	0.816***	-0.881***	0.421	-0.487
	(0.312)	(0.247)	(0.103)	(0.0634)	(0.503)	(0.369)
1998 House Republican	0.827	0.410*	0.572	-0.0774	0.910***	0.0671
	(1.436)	(0.234)	(0.457)	(0.117)	(0.294)	(0.205)
1996 Clinton	-0.0568	0.703**	-0.307	0.549**	-0.472**	0.882***
	(0.533)	(0.290)	(0.294)	(0.246)	(0.218)	(0.199)
1996 Dole	1.042	0.0966	0.154	0.295**	0.861**	-0.592
	(1.962)	(0.294)	(0.378)	(0.128)	(0.421)	(0.375)
Like Gore	-0.778	0.379**	-0.220	0.239	-0.540**	0.611**
	(1.093)	(0.176)	(0.244)	(0.282)	(0.240)	(0.257)
Like Bush	0.491	-0.284	0.309	-0.215	0.561***	-0.388**
	(0.967)	(0.188)	(0.237)	(0.205)	(0.180)	(0.173)
Approve Clinton	0.133	0.168	-0.0700	0.138	-0.212	0.448**
	(0.568)	(0.199)	(0.158)	(0.170)	(0.192)	(0.202)
Clinton Moral Climate	1.126	-0.276	0.135	-0.132	0.146	-0.535**
	(1.211)	(0.178)	(0.133)	(0.139)	(0.202)	(0.245)
Protestant	0.800	0.107	0.473**	-0.426**	0.309	0.229
	(1.121)	(0.183)	(0.195)	(0.174)	(0.219)	(0.210)
Catholic	0.437	0.283	0.646**	-0.454***	0.403*	0.132
	(1.176)	(0.180)	(0.266)	(0.168)	(0.243)	(0.224)
Marital Status	0.0188	0.265*	-0.00421	0.0301	0.253	-0.186
	(0.227)	(0.151)	(0.182)	(0.140)	(0.206)	(0.203)
Advanced Education	-0.201	0.443**	-0.126	-0.0194	-0.237	0.202
	(0.490)	(0.209)	(0.193)	(0.115)	(0.207)	(0.185)
Dropout	0.664	-0.188	0.0172	-0.294	-0.483	-0.708
	(1.218)	(0.204)	(0.277)	(0.224)	(0.317)	(0.436)
White	0.867	-0.399**	0.0989	-0.113	0.315	0.0143
	(1.008)	(0.167)	(0.194)	(0.177)	(0.214)	(0.227)
Constant	-5.473	1.378*	-2.809***	2.104***	-0.678	-1.251
	(4.405)	(0.776)	(0.484)	(0.702)	(0.887)	(0.984)
Observations	524	524	401	401	427	427

Notes: In the table above we report the (second-stage) results using instrumental variables (Probit IV), where we regress individual voting patterns during the 2000 U.S. Presidential election as reported in the 2000 ANES on Lewinsky coverage (instrumented by weather deaths), as well as a host of control variables (in the table above we only include those controls which yield significant results; a full list of controls is provided in Appendix A1). We use two dependent variables - one is a dummy indicating whether the respondent voted for George W. Bush or otherwise, and the other denotes a vote for Al Gore (or otherwise). Columns 1 and 2 refer to the regression described above using the subsample of respondents who identifed themselves as Democrats in the ANES; in columns 3 and 4 we focus only on the subsample of respondents whose political affiliation is Republican; while in columns 5 and 6 we analyze the subsample of respondents whose affiliation is Independent. All data sources are listed in Appendix A1. Robust standard errors, clustered at the congressional district level, are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

rely heavily on self-reported party affiliation rather than actual prior voting behaviour, which may be a better indication of true partisan beliefs.[16] Therefore, we also utilise prior voting behaviour in the 1996 U.S. Presidential elections as our new subgroups, in order to analyse whether the voting patterns observed in the results obtained above are still evident. The results are reported in Table 4.6, where the first two columns refer to the subgroup of voters who did not vote for Republican candidate Bob Dole, while columns 3 and 4 look at the subgroup of respondents who did not vote for Democrat Bill Clinton, with the final two columns highlighting the behaviour of those individuals who voted for Clinton. As seen below, the positive impact of Lewinsky coverage on the likelihood of a Bush vote is still evident and quantitatively very similar to that obtained previously across all specifications. In particular, it is interesting to highlight the final two columns and the clear swing in voting patterns from Gore to Bush among those who had previously voted for Bill Clinton.

4.4.4 Voter Turnout

It is also possible to analyse the impact of Lewinsky coverage on voter turnout. Voter turnout in the U.S. is relatively low - the 67.5% turnout estimate reported above is for registered voters only, with the actual figure plummeting to 51.3% if we consider the total voting-age population eligible to vote in the 2000 Presidential election. Several arguments have been put forward in order to explain this low voter turnout in the U.S., including disillusionment with the political establishment and the two-party system, trust in the government, demographic considerations and socio-economic factors (Timpone, 1998). Evidence regarding the influence of bad news on voter turnout appears to be mixed - Krupnikov (2011) finds that negative news may dampen turnout levels among those whose preferred candidate is the subject of such information, while Goldstein and Freedman (2002) argue that bad news actually induces a higher turnout across the board, irrespective of voters' individual partisan beliefs. Within the current context, it would therefore be of interest to analyse the impact of Lewinsky-related newspaper reportage on voter turnout in the 2000 U.S. Presidential election. The results are shown in Table 4.7, where the first 3 columns refer to the same subgroups specified in Table 4.4. It is clear that media coverage of the Lewinsky scandal had no impact on voter turnout across all subgroups. In addition, newspaper coverage did not affect the turnout rate when compared to the previous Presidential election in 1996, as shown in column 4 where we consider the subgroup of those respondents who voted in the previous election.

To summarise, the Probit IV estimates show that soft media reportage (in this case Lewinsky scandal coverage) had a significant impact on voting behaviour during the 2000 U.S. Presidential election. More specifically, increased media coverage of the scandal is consistent with an increase in the likelihood of voting for George W. Bush and a corresponding decline in the probability of an Al

[16]The rationale behind choosing self-reported partisan affiliation as our main subgroup classifier stems from the fact that the 1996 U.S. Presidential election was a comfortable victory for Bill Clinton, with a winning margin of 8.5%, one of the highest of the last century.

Table 4.6: IV Regression Results - Prior Voting Subgroups

	Dole = 0		Clinton = 0		Clinton = 1	
	Bush	Gore	Bush	Gore	Bush	Gore
Lewinsky Coverage	0.499*	-0.400	0.737***	-0.630**	0.888***	-0.500**
	(0.300)	(0.272)	(0.150)	(0.288)	(0.129)	(0.250)
1998 House Republican	1.102***	0.169	0.895***	0.209	0.486	-0.0560
	(0.253)	(0.123)	(0.299)	(0.148)	(0.305)	(0.161)
Republican	0.365**	-0.462**	0.474***	-0.439**	0.0565	-0.154
	(0.164)	(0.187)	(0.159)	(0.198)	(0.461)	(0.338)
Democrat	-0.724***	0.838***	-0.571***	0.797***	-0.396	0.624***
	(0.173)	(0.125)	(0.158)	(0.156)	(0.321)	(0.195)
Like Gore	-0.615***	0.689***	-0.267**	0.592***	-0.712	0.564***
	(0.158)	(0.123)	(0.132)	(0.217)	(0.460)	(0.176)
Like Bush	0.699***	-0.481***	0.394**	-0.263	0.893*	-0.665***
	(0.182)	(0.122)	(0.173)	(0.170)	(0.525)	(0.170)
Approve Clinton	-0.120	0.433***	-0.147	0.308*	0.0542	0.166
	(0.122)	(0.113)	(0.122)	(0.182)	(0.250)	(0.200)
Clinton Moral Climate	0.373***	-0.368***	0.294**	-0.476***	0.240	-0.260
	(0.139)	(0.124)	(0.114)	(0.169)	(0.212)	(0.170)
Protestant	0.363***	0.138	0.364***	0.0399	0.420**	-0.150
	(0.138)	(0.138)	(0.135)	(0.160)	(0.209)	(0.171)
Catholic	0.547***	0.156	0.525***	0.0442	0.145	-0.0372
	(0.203)	(0.119)	(0.155)	(0.178)	(0.283)	(0.182)
Marital Status	0.166	0.104	0.142	-0.102	0.192	0.335**
	(0.117)	(0.100)	(0.129)	(0.146)	(0.195)	(0.154)
Advanced Education	-0.100	0.427***	-0.0718	0.177	0.0624	0.325*
	(0.132)	(0.135)	(0.105)	(0.133)	(0.157)	(0.196)
Dropout	-0.222	-0.581***	-0.310	-0.498**	0.301	-0.500**
	(0.203)	(0.162)	(0.221)	(0.222)	(0.306)	(0.228)
White	0.389***	-0.263**	0.309***	-0.347***	0.399*	-0.158
	(0.148)	(0.113)	(0.120)	(0.134)	(0.205)	(0.162)
Constant	-2.009***	-0.0155	-2.169***	-0.204	-3.080***	1.397**
	(0.643)	(0.787)	(0.439)	(1.215)	(0.676)	(0.665)
Observations	1175	1175	974	974	578	578

Notes: In the table above we report the (second-stage) results using instrumental variables (Probit IV), where we regress individual voting patterns during the 2000 U.S. Presidential election as reported in the 2000 ANES on Lewinsky coverage (instrumented by weather deaths), as well as a host of control variables (in the table above we only include those controls which yield significant results; a full list of controls is provided in Appendix A1). We use two dependent variables - one is a dummy indicating whether the respondent voted for George W. Bush or otherwise, and the other denotes a vote for Al Gore (or otherwise). Columns 1 and 2 refer to the regression described above using the subsample of respondents whodid not vote for Republican Bob Dole in the 1996 U.S. Presidential election (i.e. voted for Bill Clinton, Ross Perot or others); in columns 3 and 4 we focus only on the subsample of respondents who did not vote for Democrat Bill Clinton in the 1996 election (i.e. voted for Bob Dole, Ross Perot or others); while in columns 5 and 6 we analyze the subsample of 1996 Clinton voters. All data sources are listed in Appendix A1. Robust standard errors, clustered at the congressional district level, are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4.7: IV Regression Results - Voter Turnout

	Full Sample	Non-Republicans	Non-Democrats	1996 Voters
Lewinsky Coverage	0.110	0.0739	0.198	0.306
	(0.294)	(0.278)	(0.432)	(0.315)
1998 House Republican	-2.055***	-1.976***	-2.099***	-1.323***
	(0.282)	(0.371)	(0.326)	(0.275)
1996 Clinton	-0.990***	-0.980***	-0.802***	-0.118
	(0.111)	(0.117)	(0.156)	(0.174)
1996 Dole	-1.200***	-1.372***	-1.127***	-0.278
	(0.166)	(0.258)	(0.170)	(0.211)
Like Gore	-0.393***	-0.412***	-0.376***	-0.0607
	(0.0970)	(0.104)	(0.127)	(0.138)
Like Bush	0.0252	0.0899	-0.125	0.0864
	(0.0997)	(0.113)	(0.119)	(0.136)
Approve Clinton	0.0822	0.0449	0.114	0.0207
	(0.109)	(0.124)	(0.129)	(0.163)
Clinton Moral Climate	0.170	0.212	0.292**	0.201
	(0.126)	(0.144)	(0.141)	(0.142)
Protestant	-0.226*	-0.280**	-0.0944	-0.0830
	(0.123)	(0.129)	(0.169)	(0.163)
Catholic	-0.262**	-0.213	-0.203	-0.129
	(0.130)	(0.140)	(0.200)	(0.167)
Marital Status	-0.275***	-0.293***	-0.286**	-0.408***
	(0.0966)	(0.108)	(0.122)	(0.133)
Advanced Education	-0.346***	-0.481***	-0.219	-0.340**
	(0.111)	(0.126)	(0.140)	(0.152)
Dropout	0.547***	0.452***	0.891***	0.427**
	(0.134)	(0.147)	(0.187)	(0.200)
White	0.303**	0.299**	0.194	0.253*
	(0.119)	(0.121)	(0.162)	(0.152)
Constant	-0.171	-0.134	-0.135	-1.700**
	(0.790)	(0.660)	(1.151)	(0.737)
Observations	1552	1151	1028	1123

Notes: In the table above we report the (second-stage) results using instrumental variables (Probit IV), where we regress a dummy variable denoting whether the 2000 ANES respondent voted or not during the 2000 U.S. Presidential election on Lewinsky coverage (instrumented by weather deaths), as well as a host of control variables (in the table above we only include those controls which yield significant results; a full list of controls is provided in Appendix A1). In columns 1 and 2 we focus on the subsample of respondents whose political affiliation as recorded in the ANES is non-Republican (i.e. Democrat, Independent or other); in columns 3 and 4 we analyze the subsample of respondents whose affiliation is non-Democrat (i.e. Republican, Independent or other); while in columns 5 and 6 we focus only on those respondents who also voted (for any party, and regardless of political affiliation) during the 1996 U.S. Presidential election. All data sources are listed in Appendix A1. Robust standard errors, clustered at the congressional district level, are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Gore vote; this swing is observed across both Republicans and Democrats. This pattern is also observed when considering an alternative subgroup classification based on prior voting behaviour during the 1996 Presidential election, where Lewinsky coverage led to a significant shift from Gore to Bush among people who had voted for Clinton. Conversely, Lewinsky coverage had no impact on voter turnout across all specifications and subgroups.

4.4.5 Interpretation of Results

We now look at the channels through which Lewinsky coverage may have influenced individual voting behaviour. In Section 4.1.2 we argued that people may opt to read soft news rather than more substantive hard stories due to the time/effort required in order to interpret the latter, whereas the former is more readily consumable as entertainment. This preference for soft political news may be related to several individual characteristics, most notably a person's level of political engagement or involvement, with several authors (e.g. Baum, 2003 and Prior, 2003) specifically citing this link as the key determinant in a voter's reading preferences. A priori, based on the earlier discussion we would expect that people with a low level of political engagement would be more likely to opt for Lewinsky-related stories over more substantive news, meaning that Lewinsky coverage would have had a greater impact on their voting choices relative to others.

We can therefore check whether our results are consistent with this idea by looking at different subsamples within our data on the basis of the ANES respondents' political knowledge/engagement (based on their survey responses). The results of this analysis are shown in Table 4.8, where we report the IV results obtained for different respondent subsamples in each column. In columns 1 and 2 we look at those respondents who were unable to mention at least one of the candidates involved in the 2000 U.S. House of Representatives election within their congressional district of residence. Columns 3 and 4 look at the subsample of respondents who, in the survey, reported that their knowledge of American politics and government was weaker than that of the average U.S. citizen. The results are in line with our prior expectations, since in both cases the coefficients on Lewinsky coverage are statistically significant and larger than those obtained in the previous section. [17] Note that these results also hold when looking at the non-Republican subgroup as in Table 4.4. Therefore, it appears as though Lewinsky coverage had a higher impact in terms of voting on those people who are politically disengaged or uninformed, in line with the discussion in Section 4.1.2.

We also postulated that respondents' religiosity may have played a key role in determining the impact of Lewinsky coverage on voting, since such stories would have been of greater interest for people whose concern for moral issues is high.

[17] Conversely, when looking at the subsample of people who *were* able to mention at least one House of Representatives candidate, or who reported an above-average knowledge of U.S. politics/government, the coefficients on Lewinsky coverage were not statistically significant. This suggests that the impact of Lewinsky articles on voting was mainly driven by people whose political knowledge or interest is weak.

Table 4.8: Interpretation of IV Results

	Candidate Knowledge		Political Knowledge		Religiosity	
	Bush	Gore	Bush	Gore	Bush	Gore
Lewinsky Coverage	0.760***	-0.512**	0.750***	-0.626***	0.735***	-0.666***
	(0.141)	(0.214)	(0.178)	(0.189)	(0.147)	(0.154)
1998 House Republican	0.905***	-0.124	0.911**	-0.0910	0.994***	-0.0942
	(0.254)	(0.135)	(0.379)	(0.146)	(0.311)	(0.133)
1996 Clinton	-0.241**	0.958***	-0.216	0.851***	-0.233*	0.794***
	(0.118)	(0.165)	(0.152)	(0.206)	(0.140)	(0.192)
1996 Dole	0.692***	-0.331*	0.692**	-0.335	0.553*	-0.209
	(0.218)	(0.179)	(0.327)	(0.216)	(0.292)	(0.209)
Like Gore	-0.460***	0.712***	-0.516***	0.684***	-0.429***	0.521***
	(0.129)	(0.129)	(0.194)	(0.149)	(0.160)	(0.149)
Like Bush	0.548***	-0.444***	0.480**	-0.414***	0.694***	-0.497***
	(0.153)	(0.122)	(0.213)	(0.143)	(0.194)	(0.143)
Approve Clinton	-0.190**	0.373***	-0.122	0.284**	-0.199*	0.262*
	(0.0954)	(0.127)	(0.119)	(0.144)	(0.120)	(0.137)
Clinton Moral Climate	0.315***	-0.358***	0.414***	-0.396***	0.356***	-0.377***
	(0.103)	(0.111)	(0.115)	(0.127)	(0.118)	(0.112)
Protestant	0.370***	-0.0567	0.316**	0.0722	0.439**	0.0309
	(0.127)	(0.139)	(0.128)	(0.159)	(0.173)	(0.194)
Catholic	0.330**	0.0814	0.405**	0.0777	0.318*	0.228
	(0.131)	(0.125)	(0.174)	(0.142)	(0.181)	(0.189)
Marital Status	0.123	0.134	0.103	0.207**	0.0419	0.139
	(0.0956)	(0.0971)	(0.107)	(0.102)	(0.112)	(0.106)
Advanced Education	-0.0933	0.297***	-0.0804	0.0736	0.146	0.0363
	(0.104)	(0.111)	(0.130)	(0.132)	(0.124)	(0.111)
Dropout	-0.104	-0.523***	-0.0318	-0.490***	-0.146	-0.257
	(0.178)	(0.161)	(0.192)	(0.172)	(0.212)	(0.168)
White	0.382***	-0.384***	0.519***	-0.315***	0.329***	-0.365***
	(0.109)	(0.107)	(0.161)	(0.116)	(0.127)	(0.118)
Constant	-2.584***	0.175	-2.504***	0.281	-2.567***	0.860
	(0.407)	(0.722)	(0.433)	(0.705)	(0.498)	(0.638)
Observations	1287	1287	1022	1022	933	933

Notes: In the table above we report the (second-stage) results using instrumental variables (Probit IV), where we regress individual voting patterns during the 2000 U.S. Presidential election as reported in the 2000 ANES on Lewinsky coverage (instrumented by weather deaths), as well as a host of control variables (in the table above we only include those controls which yield significant results; a full list of controls is provided in Appendix A1). We use two dependent variables - one is a dummy indicating whether the respondent voted for George W. Bush or otherwise, and the other denotes a vote for Al Gore (or otherwise). In columns 1 and 2 we look at those respondents who were unable to mention at least one of the candidates involved in the 2000 U.S. House of Representatives election within their congressional district of residence. Columns 3 and 4 look at the subsample of respondents who, in the survey, reported that their knowledge of American politics and government was weaker than that of the average U.S. citizen. Finally, columns 5 and 6 show the results for a subsample of respondents who stated that religion plays an important part in their everyday lives. All data sources are listed in Appendix A1. Robust standard errors, clustered at the congressional district level, are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Columns 5 and 6 show the results for a subsample of respondents who stated that religion plays an important part in their everyday lives.[18] Once again, the results are markedly similar to those obtained for the full sample in Table 4.4, confirming the idea that religious people may have been more acutely influenced in their voting decision by Lewinsky coverage.[19]

The evidence so far supports the idea that Lewinsky coverage mainly had an impact on voting behaviour for politically disengaged respondents, since they would be more inclined to focus on soft political stories like the Lewinsky scandal rather than more substantive material, meaning that the bulk of their political information (which forms the basis of their voting choices) would have been garnered via this channel. In Section 4.1.2 we also mentioned an alternative channel through which Lewinsky coverage may have influenced voting, namely by crowding out hard political news. However, the use of IV based on the crowding-out caused by extreme weather deaths significantly reduces the possibility that this mechanism is the main driver behind our results. This is because it is reasonable to expect that extreme weather stories would crowd out all kinds of news stories, not simply the Lewinsky scandal, including also hard political news (e.g. each candidate's proposed tax policies once in Office).

In other words, the instrument helps us to identify the variation in Lewinsky coverage caused by the presence (or absence) of major weather-related articles, which would also have a similar impact on other more substantive political stories. The fact that our IV results are still statistically significant therefore seems to preclude the possibility that the crowding out of hard stories by soft stories is the main mechanism through which Lewinsky coverage influenced the 2000 election.[20]

4.5 Robustness Tests

In this section we present a series of tests in order to examine the robustness of the estimates presented in this paper. In particular, we test whether the conclusions derived in the previous section are robust to the inclusion of an additional variable denoting newspaper editorial slant, an alternative weighting system for the Lewinsky scandal coverage variable based on circulation levels, potential attrition bias and the use of an alternative instrumental variable. We also test the reliability of our estimates via a placebo treatment as well as an overidentification test on our instrument.

[18] Various other measures of religiosity have been used, like for example the frequency with which respondents stated that they attended some form of religious service (e.g. weekly), all yielding very similar results to those reported above.

[19] Once again, the subsample of respondents who claimed that religion did not play an important part in their lives yielded results which are not statistically significant.

[20] In Section 4.5.4 we look at this crowding-out idea further by considering the possibility that the decision to publish extreme weather stories rather than Lewinsky articles may have been influenced by the newspaper's editorial slant.

4.5.1 Survey Response Validity

Firstly, we tackle one of the most commonly-cited criticisms of using election surveys for this kind of analysis - namely that people who participate in such surveys have a tendency to wilfully or unintentionally misrepresent their voting behaviour (Lax, 2014). In fact, Reuben (2014) finds that one third of UK citizens do not provide accurate information about themselves, or indeed their voting choices, when asked questions by pollsters. Clearly, this is a cause for concern in this paper since it may be the case that respondents have provided inaccurate information regarding their voting preferences, which in turn may be the driving force behind the results obtained above. For example, it is conceivable that non-Republicans may have lied about voting for Bush in order to conceal the fact that they had in fact voted for the losing candidate (i.e. Gore), due to various reasons like a distaste for losing or concerns regarding the way in which such data may be used or distributed.

We therefore use the results from a follow-up survey in order to test the validity of our current results. This is possible since in 2002 ANES conducted another survey among the same respondents from the 2000 edition, asking the same questions regarding voting patterns as before in order to assess the sensitivity of the original responses to misrepresentation. The results are shown in Table 4.9. As seen below, the coefficients on our Lewinsky coverage variable, in all cases, are very similar to those obtained previously in Table 4.4, where once again newspaper coverage of the scandal had a positive, significant impact on the likelihood of a Bush vote, and a negative effect on the probability of voting for Gore. As an additional robustness test, we also used the reported partisanship/ideology of respondents in the 2002 survey in our original (2000) specification, due to the potential pitfalls associated with using self-reported partisanship data. Once again the results obtained are almost identical to those reported in Table 4.4 using the 2000 self-reported partisanship data.[21]

4.5.2 Placebo Tests

Impact on 1996 election: In this section we address the possibility that our Lewinsky coverage variable is merely a reflection of the level of media attention afforded to various Clinton-related scandals over his tenure as President, and in particular the 1996 Whitewater investigation into numerous controversies surrounding the Presidency.[22] This is potentially problematic since many of these controversies portrayed the entire administration in a negative light, including then-Vice President Al Gore, and may thus be considered to be "informative" in terms of his presidential credentials.

[21]These results are omitted due to space limitations, but are readily available on request.

[22]The Whitewater controversy mainly centred on the Clintons' real estate investments in the failed Whitewater Development Corporation in the 1970s and 1980s, although the investigation also incorporated other Clinton-era scandals like the unprecedented dismissal of 7 White House Travel Office employees in 1993 and the improper access of confidential FBI files by White House officials in 1993-94.

Table 4.9: Robustness Tests - ANES 2002 Data

	Full Sample		Non-Republicans		Non-Democrats	
	Bush	Gore	Bush	Gore	Bush	Gore
Lewinsky Coverage	0.760***	-0.783***	0.900***	-0.852***	0.681***	-0.704***
	(0.102)	(0.0882)	(0.0731)	(0.114)	(0.184)	(0.191)
1998 Republican	0.418***	-0.158	0.282**	-0.178	0.449**	-0.220
	(0.148)	(0.109)	(0.141)	(0.134)	(0.191)	(0.158)
1996 Clinton	-0.201**	0.663***	-0.117	0.600***	-0.236*	0.885***
	(0.0977)	(0.158)	(0.0999)	(0.190)	(0.139)	(0.219)
1996 Dole	0.459**	-0.160	0.232	-0.243	0.477*	-0.0825
	(0.209)	(0.164)	(0.199)	(0.227)	(0.286)	(0.221)
Like Gore	-0.231**	0.360***	-0.136	0.304**	-0.264	0.658**
	(0.108)	(0.129)	(0.0997)	(0.139)	(0.163)	(0.296)
Like Bush	0.526***	-0.383***	0.426***	-0.399***	0.594**	-0.532**
	(0.178)	(0.144)	(0.156)	(0.155)	(0.247)	(0.242)
Approve Clinton	-0.198*	0.239**	-0.174	0.274*	-0.0893	0.150
	(0.104)	(0.120)	(0.113)	(0.141)	(0.120)	(0.162)
Clinton Moral Climate	0.218**	-0.296***	0.201*	-0.309**	0.115	-0.226
	(0.0982)	(0.107)	(0.107)	(0.121)	(0.117)	(0.160)
Protestant	0.236**	-0.168	0.0977	-0.0416	0.335***	-0.346**
	(0.106)	(0.103)	(0.115)	(0.114)	(0.129)	(0.151)
Catholic	0.301**	-0.173	0.0700	-0.0184	0.611***	-0.642***
	(0.124)	(0.112)	(0.127)	(0.128)	(0.164)	(0.186)
Marital Status	0.0528	0.0157	0.0310	0.0398	0.0805	0.0373
	(0.0834)	(0.0820)	(0.0900)	(0.0948)	(0.115)	(0.129)
Advanced Education	-0.0544	0.281**	-0.0966	0.354**	-0.0811	0.395**
	(0.0857)	(0.117)	(0.0959)	(0.152)	(0.111)	(0.193)
Dropout	-0.0748	-0.185	-0.161	-0.0918	-0.00694	-0.683**
	(0.144)	(0.134)	(0.154)	(0.157)	(0.200)	(0.273)
White	0.232**	-0.280***	0.149	-0.210*	0.220	-0.258
	(0.0993)	(0.101)	(0.104)	(0.113)	(0.142)	(0.177)
Constant	-2.147***	1.268***	-1.945***	1.132**	-1.776***	0.339
	(0.304)	(0.416)	(0.328)	(0.480)	(0.458)	(0.923)
Observations	1184	1184	864	864	781	781

Notes: In the above table we report the second-stage results from a Probit IV analysis, where we regress individual voting patterns during the 2000 U.S. Presidential election as reported in the 2002 ANES on Lewinsky coverage (instrumented by weather deaths), as well as a host of control variables (in the table above we only include those controls which yield significant results; a full list of controls is provided in Appendix A1). We use two dependent variables - one is a dummy indicating whether the respondent voted for George W. Bush or otherwise, and the other denotes a vote for Al Gore (or otherwise). Columns 1 and 2 refer to the regression described above using the full sample; in columns 3 and 4 we focus only on the subsample of respondents whose political affiliation as recorded in the 2002 ANES is non-Republican (i.e. Democrat, Independent or other); while in columns 5 and 6 we analyze the subsample of respondents whose affiliation is non-Democrat (i.e. Republican, Independent or other). All data sources are listed in Appendix A1. Robust standard errors, clustered at the congressional district level, are reported in parentheses.
Robust standard errors are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Since the investigation took place prior to the 1996 Presidential election, it is reasonable to expect that if coverage of such scandals had any impact on the 2000 election then it would also have had at least an equivalent impact on 1996 voting patterns. We can therefore test whether coverage of prior scandals is driving our results by looking at the relationship between Lewinsky articles and voting in the 1996 election. Given that the 2000 ANES Time Series study reports respondents' voting patterns for both the 2000 (in the post-election survey) and 1996 U.S. Presidential elections, it is possible to construct a placebo treatment by regressing 1996 voting patterns on our Lewinsky coverage variable and the other control variables specified in Section 4.2. In addition, the test also enables us to further analyse the validity of the reported causal relationship between media coverage and voting.

Table 4.10 shows the IV results, where the same explanatory variables, instrument and subgroup classifications are used as above, with the only difference being that now our dependent variable refers to the Presidential candidates that took part in the 1996 elections. Hence, in columns 1, 3 and 5 the dependent variable is a dummy denoting whether the respondent voted for Republican candidate Bob Dole or not, while in columns 2, 4 and 6 the dependent variable assumes a value of 1 if respondents voted for Bill Clinton. As seen from the results, coverage of the Lewinsky scandal had no impact on voting patterns across all columns and subgroups.

Impact of prior extreme weather incidents: We now assess the validity of our identification strategy by looking at the relationship between Lewinsky coverage and extreme weather deaths in the equivalent time period immediately *prior* to the actual scandal. As stated earlier, our identifying relationship rests on the premise that extreme weather deaths in a given county generate significant and prominent news stories, which in turn crowd-out coverage of the Lewinsky scandal from the county's top circulating local newspapers. This is backed up by the first-stage results reported in Table 4.4 Panel B as well as the use of the *change* in weather deaths in the period under review relative to the previous 20-year average, which yields similar results. Nonetheless, to further ensure that our identifying relationship is indeed valid, we regress our Lewinsky coverage variable on the number of extreme weather deaths (using NOAA data, as before) in the equivalent period before the Lewinsky story broke - namely January 17, 1995 to August 31, 1997. Any significant, negative relationship from such a regression would raise significant concerns regarding the validity of our identification strategy since it would indicate that our first-stage relationship is either spurious or driven by extraneous factors.

The results are shown in Table 4.11, where we regress Lewinsky coverage on both the number of extreme weather deaths in each respondent's congressional district over the period January 17, 1995 to August 31, 1997 (Column 1), as well as the change in extreme weather deaths over the same period relative to the average in the previous 20 years (Column 2). As seen below, neither measure yields statistically significant results.

Table 4.10: Robustness Tests - Impact on 1996 Election

	Full Sample		Non-Republicans		Non-Democrats	
	Dole	Clinton	Dole	Clinton	Dole	Clinton
Lewinsky Coverage	0.355	0.122	0.313	-0.0561	0.251	0.165
	(0.294)	(0.251)	(0.357)	(0.252)	(0.387)	(0.378)
1998 House Republican	0.570***	-0.285***	0.564***	-0.249**	0.567***	-0.126
	(0.138)	(0.102)	(0.162)	(0.125)	(0.140)	(0.130)
Like Gore	-0.473***	0.624***	-0.552***	0.600***	-0.363***	0.622***
	(0.109)	(0.0859)	(0.144)	(0.0924)	(0.117)	(0.114)
Like Bush	0.972***	-0.632***	0.928***	-0.492***	0.943***	-0.542***
	(0.153)	(0.0894)	(0.163)	(0.0998)	(0.142)	(0.121)
Approve Clinton	-0.536***	0.928***	-0.371***	0.848***	-0.494***	0.707***
	(0.127)	(0.102)	(0.143)	(0.115)	(0.144)	(0.120)
Clinton Moral Climate	0.395***	-0.136	0.483***	-0.123	0.384***	-0.157
	(0.0996)	(0.107)	(0.138)	(0.122)	(0.111)	(0.129)
Protestant	0.271**	0.272**	0.346**	0.259**	0.305**	0.212
	(0.115)	(0.106)	(0.170)	(0.117)	(0.134)	(0.146)
Catholic	0.219*	0.294***	0.223	0.338***	0.314**	0.250
	(0.123)	(0.107)	(0.197)	(0.119)	(0.149)	(0.156)
Marital Status	0.203*	-0.000617	0.0859	0.0782	0.222*	-0.0989
	(0.109)	(0.0892)	(0.158)	(0.0969)	(0.124)	(0.123)
Advanced Education	0.319***	0.243***	0.287**	0.275***	0.318***	0.237**
	(0.0973)	(0.0869)	(0.136)	(0.0980)	(0.107)	(0.116)
Dropout	-0.316	-0.183	-0.378	-0.194	-0.411	0.135
	(0.225)	(0.145)	(0.302)	(0.154)	(0.284)	(0.211)
White	0.397***	-0.0700	0.382**	-0.0550	0.282	-0.0738
	(0.138)	(0.115)	(0.176)	(0.113)	(0.173)	(0.170)
Constant	-2.110***	-1.205*	-2.095***	-0.872	-1.744*	-1.527
	(0.659)	(0.636)	(0.702)	(0.597)	(0.943)	(0.937)
Observations	1552	1552	1050	1151	1028	1028

Notes: In the table above we report the second-stage results from a Probit IV analysis, where we regress individual voting patterns during the 1996 U.S. Presidential election as reported in the 2000 ANES on Lewinsky coverage (instrumented by weather deaths), as well as a host of control variables (in the table above we only include those controls which yield significant results; a full list of controls is provided in Appendix A1). We use two dependent variables - one is a dummy indicating whether the respondent voted for Bob Dole or otherwise, and the other denotes a vote for Bill Clinton (or otherwise). Columns 1 and 2 refer to the regression described above using the full sample; in columns 3 and 4 we focus only on the subsample of respondents whose political affiliation as recorded in the ANES is non-Republican (i.e. Democrat, Independent or other); while in columns 5 and 6 we analyze the subsample of respondents whose affiliation is non-Democrat (i.e. Republican, Independent or other). All data sources are listed in Appendix A1. Robust standard errors, clustered at the congressional district level, are reported in parentheses.
Robust standard errors are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4.11: Robustness Tests - Validity of Identification Strategy

	(1)	(2)
Weather Deaths 1995-97	-0.00236	
	(0.00279)	
Change in Weather Deaths 1995-97		-0.00347
		(0.00305)
Constant	2.101***	2.102***
	(0.0314)	(0.0305)
Observations	1552	1552
Pseudo R^2	0.0003	0.0006
F	0.712	1.301

Notes: In this table we test the validity of our original identification strategy between newspaper Lewinsky coverage and extreme weather deaths, used throughout this paper as the backbone of our IV analysis. In column 1 we regress Lewinsky coverage on the number of extreme weather deaths in the period January 17, 1995 to August 31, 1997 (i.e. the period prior to the one analysed in this paper, and prior to the disclosure of the Lewinsky affair). In column 2 we regress Lewinsky coverage on the change in extreme weather deaths over the period January 17, 1995 to August 31, 1997, relative to the average from the previous 20 years. Robust standard errors are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

4.5.3 Attrition Bias

Here we address concerns that the results obtained are being driven by the non-random composition of respondents in the post-election survey. As mentioned earlier the ANES Time Series study consists of two phases - the pre-election interview, held prior to election day, which records respondents' voting intentions and the post-election interview which deals with actual voting patterns. The initial pre-election interview consisted of 1,807 survey participants; however, when it came to the post-election phase (which was used in this study) only 1,555 respondents participated in the interview. This drop in participation of 252 respondents (almost 14% of the original sample size) may be problematic since it may have significantly altered the composition of our sample, thus limiting the applicability of our subsequent conclusions. In particular, there may be concerns that these respondents refused to participate in the post-election phase since they abstained from voting in the 2000 Presidential election, or else were either unhappy at the outcome of the election since their favoured candidate was defeated (in the case of Democrat voters), or elated following electoral victory (in the case of Republicans).

We therefore account for any potential attrition bias in Table 4.12, where we combine the actual voting outcomes reported by the 1,555 post-survey respondents with the voting intentions stated by the remaining 252 dropouts, with the latter acting as a proxy for actual voting. This is possible since in the pre-election phase respondents were specifically asked to state who they would vote for in the upcoming Presidential election. The results obtained are broadly similar to those reported in Table 4.4 Panel A, particularly with regards to the positive and statistically significant impact of Lewinsky coverage on the likelihood of voting for

Bush. The only difference is that the relationship between coverage and voting for Al Gore is now statistically insignificant in all cases. This may be due to the fact that prior voting intentions do not perfectly match up to actual voting - in fact from the 1,555 post-election survey participants almost 30% of respondents did not vote according to their stated intentions in the pre-election phase. Furthermore, among these 1,555 respondents 668 said that they would vote for Gore in the pre-election survey, compared to the 589 who actually did as per the post-election responses. Therefore it is possible that the 252 dropouts may have overstated their intentions to vote for Gore prior to the election, and changed their vote subsequently as observed in the main sample.

4.5.4 Newspaper Editorial Slant

In Section 4.3 we showed that the level of Lewinsky coverage was not influenced by the newspapers' editorial slant. Nonetheless, a reasonable argument can be made that editorial slant may be correlated with the decision to publish Lewinsky stories versus news on extreme weather deaths, as postulated in our identification strategy. For example, it possible that a newspaper editor with strong Democratic tendencies would have been more predisposed to cease coverage of the Lewinsky scandal in order to focus on news related to extreme weather fatalities. Therefore, such behaviour could well be the driving force behind the strong negative relationship observed in Table 4.4 Panel B between the number of extreme weather deaths and Lewinsky coverage. To address this problem we include the editorial slant variable previously used in Section 4.3, once again based on data by Groseclose and Milyo (2005).[23]

Table 4.13 replicates the IV regressions estimated in Table 4.4 Panel A, this time with the inclusion of our editorial slant variable. The results are almost identical to those obtained in the previous section, with the only difference being that now Lewinsky coverage is negatively and significantly-related to the likelihood of voting for Al Gore within the non-Democrat subgroup (column 6). Notice that the editorial slant variable has no statistically significant relationship with voting patterns across all specifications set out in Table 4.13.

4.5.5 Alternative Weighting System

In this section we employ a different weighting method in order to derive our variable for Lewinsky-related news coverage, in order to address any doubts regarding the dependence of our results to the specific weighting method used. As described earlier, the current weighting method is based on the relative population size of each newspaper's target county/ies within each congressional district included in the survey. Here we utilise a similar method, although this time the basis of the calculation is the relative circulation level of each newspaper. The rationale behind this method is that a high-circulation newspaper within a given

[23] In total estimates of editorial slant are available for 200 out of the 257 newspapers used in this paper.

Table 4.12: Robustness Tests - Attrition Bias

	Full Sample		Non-Republicans		Non-Democrats	
	Bush	Gore	Bush	Gore	Bush	Gore
Lewinsky Coverage	0.599***	-0.291	0.750***	-0.282	0.401	-0.0920
	(0.182)	(0.245)	(0.146)	(0.243)	(0.326)	(0.387)
1998 House Republican	0.973***	-0.113	0.768***	-0.0425	1.173***	-0.0546
	(0.217)	(0.109)	(0.211)	(0.129)	(0.230)	(0.137)
1996 Clinton	-0.252**	0.935***	-0.231**	0.927***	-0.148	0.752***
	(0.104)	(0.107)	(0.110)	(0.108)	(0.136)	(0.136)
1996 Dole	0.784***	-0.469***	0.551**	-0.353*	0.947***	-0.591***
	(0.225)	(0.165)	(0.226)	(0.195)	(0.254)	(0.199)
Like Gore	-0.618***	0.827***	-0.564***	0.738***	-0.662***	0.985***
	(0.148)	(0.109)	(0.150)	(0.100)	(0.174)	(0.121)
Like Bush	0.749***	-0.653***	0.633***	-0.572***	0.790***	-0.579***
	(0.167)	(0.104)	(0.163)	(0.102)	(0.171)	(0.118)
Approve Clinton	-0.147	0.390***	-0.0957	0.370***	-0.128	0.317**
	(0.0916)	(0.107)	(0.0993)	(0.114)	(0.110)	(0.129)
Clinton Moral Climate	0.270***	-0.391***	0.313***	-0.402***	0.148	-0.512***
	(0.0973)	(0.105)	(0.118)	(0.116)	(0.116)	(0.156)
Protestant	0.346***	0.0806	0.291**	0.199*	0.391***	0.0121
	(0.107)	(0.105)	(0.123)	(0.111)	(0.123)	(0.152)
Catholic	0.417***	0.0928	0.311**	0.174	0.519***	0.0588
	(0.128)	(0.111)	(0.145)	(0.115)	(0.141)	(0.175)
Marital Status	0.185**	0.0768	0.173*	0.101	0.260**	-0.0392
	(0.0915)	(0.0855)	(0.101)	(0.0936)	(0.118)	(0.126)
Advanced Education	-0.107	0.276***	-0.125	0.340***	-0.140	0.295**
	(0.0924)	(0.0944)	(0.109)	(0.104)	(0.117)	(0.124)
Dropout	-0.0664	-0.385***	0.00581	-0.328**	-0.183	-0.398*
	(0.138)	(0.134)	(0.140)	(0.140)	(0.206)	(0.230)
White	0.208**	-0.320***	0.168	-0.308***	0.197	-0.202
	(0.101)	(0.0985)	(0.106)	(0.103)	(0.136)	(0.160)
Constant	-2.174***	-0.389	-2.092***	-0.380	-1.802***	-1.432
	(0.379)	(0.694)	(0.346)	(0.614)	(0.694)	(0.990)
Observations	1802	1802	1353	1353	1185	1185

Notes: In the table above we report the (second-stage) results using instrumental variables (Probit IV), where we regress individual voting patterns during the 2000 U.S. Presidential election as reported in the 2000 ANES on Lewinsky coverage (instrumented by weather deaths), as well as a host of control variables (in the table above we only include those controls which yield significant results; a full list of controls is provided in Appendix A1). In this case, we combine the actual voting outcomes reported by the 1,555 post-survey respondents with the voting intentions stated by the remaining 252 dropouts, with the latter acting as a proxy for actual voting. We use two dependent variables - one is a dummy indicating whether the respondent voted for George W. Bush or otherwise, and the other denotes a vote for Al Gore (or otherwise).Columns 1 and 2 refer to the regression described above using the full sample; in columns 3 and 4 we focus only on the subsample of respondents whose political affiliation as recorded in the ANES is non-Republican (i.e. Democrat, Independent or other); while in columns 5 and 6 we analyze the subsample of respondents whose affiliation is non-Democrat (i.e. Republican, Independent or other). All data sources are listed in Appendix A1. Robust standard errors, clustered at the congressional district level, are reported in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4.13: Robustness Tests - Newspaper Editorial Slant

	Full Sample		Non-Republicans		Non-Democrats	
	Bush	Gore	Bush	Gore	Bush	Gore
Lewinsky Coverage	0.546***	-0.479***	0.746***	-0.473***	0.353	-0.555***
	(0.193)	(0.160)	(0.147)	(0.166)	(0.330)	(0.209)
Editorial Slant	-0.00261	-0.00335	-0.00545	-0.00148	0.00106	-0.0154
	(0.00679)	(0.00618)	(0.00753)	(0.00691)	(0.00900)	(0.00963)
1998 House Republican	1.062***	-0.107	0.842***	-0.0861	1.196***	-0.0667
	(0.209)	(0.119)	(0.205)	(0.140)	(0.217)	(0.147)
1996 Clinton	-0.345***	0.935***	-0.340**	0.935***	-0.212	0.793***
	(0.131)	(0.131)	(0.142)	(0.135)	(0.162)	(0.162)
1996 Dole	0.868***	-0.428**	0.620***	-0.364*	0.984***	-0.413**
	(0.214)	(0.169)	(0.240)	(0.206)	(0.215)	(0.206)
Like Gore	-0.596***	0.714***	-0.527***	0.669***	-0.596***	0.737***
	(0.143)	(0.121)	(0.156)	(0.121)	(0.157)	(0.184)
Like Bush	0.758***	-0.533***	0.609***	-0.441***	0.769***	-0.470***
	(0.158)	(0.114)	(0.160)	(0.115)	(0.158)	(0.150)
Approve Clinton	-0.177	0.385***	-0.0989	0.418***	-0.173	0.267*
	(0.112)	(0.120)	(0.122)	(0.129)	(0.135)	(0.147)
Clinton Moral Climate	0.280**	-0.336***	0.329**	-0.380***	0.174	-0.367**
	(0.121)	(0.119)	(0.145)	(0.134)	(0.141)	(0.164)
Protestant	0.331***	0.0780	0.220	0.263**	0.339**	-0.141
	(0.124)	(0.116)	(0.146)	(0.126)	(0.145)	(0.147)
Catholic	0.342**	0.0590	0.237	0.172	0.408**	-0.0821
	(0.144)	(0.122)	(0.160)	(0.131)	(0.174)	(0.164)
Marital Status	0.0736	0.177*	0.0464	0.206*	0.140	0.155
	(0.112)	(0.0978)	(0.122)	(0.109)	(0.143)	(0.131)
Advanced Education	-0.162	0.280***	-0.208	0.374***	-0.186	0.229*
	(0.115)	(0.106)	(0.132)	(0.119)	(0.137)	(0.133)
Dropout	-0.371*	-0.519***	-0.265	-0.441**	-0.610**	-0.949***
	(0.199)	(0.171)	(0.192)	(0.186)	(0.279)	(0.298)
White	0.390***	-0.273**	0.431***	-0.282**	0.218	0.0119
	(0.124)	(0.114)	(0.138)	(0.123)	(0.158)	(0.173)
Constant	-1.903***	0.116	-1.734***	-0.163	-1.603**	0.610
	(0.510)	(0.521)	(0.559)	(0.537)	(0.651)	(0.597)
Observations	1314	1314	972	972	879	879

Notes: In the table above we report the (second-stage) results using instrumental variables (Probit IV), where we regress individual voting patterns during the 2000 U.S. Presidential election as reported in the 2000 ANES on Lewinsky coverage (instrumented by weather deaths), as well as a host of control variables (in the table above we only include those controls which yield significant results; a full list of controls is provided in Appendix A1). In this case, we also include newspaper editorial slant as an additional control variable. We use two dependent variables - one is a dummy indicating whether the respondent voted for George W. Bush or otherwise, and the other denotes a vote for Al Gore (or otherwise).Columns 1 and 2 refer to the regression described above using the full sample; in columns 3 and 4 we focus only on the subsample of respondents whose political affiliation as recorded in the ANES is non-Republican (i.e. Democrat, Independent or other); while in columns 5 and 6 we analyze the subsample of respondents whose affiliation is non-Democrat (i.e. Republican, Independent or other). All data sources are listed in Appendix A1. Robust standard errors, clustered at the congressional district level, are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

congressional district is more likely to be read by the average person than a lower-circulating paper, meaning that he/she is more likely to have been exposed to the former's coverage of the Lewinsky scandal over the period under review.

Recall the example presented in Section 4.2, where we considered California's 48th congressional district and its top circulating newspapers in 2000, namely the Orange County Register (circulation of 250,724 copies per week), the San Diego Union-Tribune (296,331) and the Riverside Press-Enterprise (172,593). As stated earlier, the amount of coverage (per 1,000 articles) afforded to the Lewinsky scandal in each of these newspapers over the period January 17, 1998 to August 31, 2000 was 2.19, 3.2 and 1.15 articles respectively. Hence the weighted average level of Lewinsky coverage for California's 48th congressional district based on this alternative weighting system was of 2.36 articles, which in reality is not far off the 2.34 obtained under the previous weighting system based on relative county population.

The results are shown in Table 4.14, where due to space limitations we only show the results obtained using IV for the new weighting method; they are very similar to those obtained in the previous section (Table 4.4 Panel A).

4.5.6 Alternative Instrument and Overidentification Tests

As a final robustness check, we now use the cost in US dollars of damages caused by extreme weather events in each congressional district as our instrument (once again using NOAA data). The rationale behind the use of this instrument is similar to the original identification strategy employed based on the number of weather-related deaths, although in this case its scope is somewhat more broad since it encompasses extreme weather events that may have not led to any casualties but which were nonetheless still large or significant enough to cause material damages. It is therefore plausible to check whether the use of this new instrument alters the results obtained above since it would capture a wider range of prominent stories that may crowd out Lewinsky coverage, thus helping to further alleviate the measurement error problem. Table 4.15 Panel B shows the first-stage regression results where our Lewinsky coverage variable is regressed on extreme weather-related damage costs, where as expected we observe a negative and statistically significant relationship, similar to the one obtained for the original instrument. The results of the IV regression using this new instrument are reported in Panel A. Once again, the results are almost identical to those derived earlier in Table 4.4 Panel A across all columns; the only difference is that Lewinsky coverage now has a statistically significant (negative) impact on the probability of voting for Al Gore among non-Democrats (column 6).

This alternative instrument also enables us to run a series of overidentification tests in order to check the validity of our identification strategy. To assess each instrument's exogeneity we use a modified overidentification test similar to the one employed in Acemoglu, Johnson, and Robinson (2001). These tests involve running the usual IV regressions using both extreme weather deaths and damage costs as our instruments and including each instrument separately as an explana-

Table 4.14: Robustness Tests - Alternative Weighting Method

	Full Sample		Non-Republicans		Non-Democrats	
	Bush	Gore	Bush	Gore	Bush	Gore
Lewinsky Coverage	0.675***	-0.530***	0.805***	-0.493**	0.523	-0.512*
	(0.176)	(0.192)	(0.135)	(0.201)	(0.337)	(0.300)
1998 House Republican	0.908***	-0.0827	0.699***	-0.0153	1.094***	-0.0475
	(0.274)	(0.112)	(0.236)	(0.131)	(0.360)	(0.144)
1996 Clinton	-0.244**	0.889***	-0.236*	0.915***	-0.126	0.766***
	(0.115)	(0.161)	(0.126)	(0.155)	(0.144)	(0.168)
1996 Dole	0.700***	-0.320*	0.496**	-0.270	0.860**	-0.382
	(0.266)	(0.173)	(0.245)	(0.197)	(0.357)	(0.244)
Like Gore	-0.552***	0.692***	-0.528***	0.647***	-0.573***	0.821***
	(0.151)	(0.134)	(0.151)	(0.119)	(0.182)	(0.212)
Like Bush	0.641***	-0.525***	0.549***	-0.477***	0.673***	-0.483***
	(0.187)	(0.126)	(0.167)	(0.116)	(0.215)	(0.160)
Approve Clinton	-0.171*	0.351***	-0.140	0.367***	-0.138	0.297*
	(0.102)	(0.121)	(0.106)	(0.123)	(0.134)	(0.164)
Clinton Moral Climate	0.296***	-0.371***	0.325***	-0.403***	0.205*	-0.445***
	(0.103)	(0.107)	(0.124)	(0.121)	(0.122)	(0.167)
Protestant	0.361***	0.00306	0.288**	0.146	0.385***	-0.103
	(0.109)	(0.111)	(0.127)	(0.122)	(0.126)	(0.158)
Catholic	0.363***	0.0828	0.249	0.200*	0.481***	-0.0634
	(0.138)	(0.113)	(0.155)	(0.120)	(0.153)	(0.168)
Marital Status	0.127	0.115	0.147	0.133	0.202	0.0351
	(0.0997)	(0.0865)	(0.111)	(0.0979)	(0.152)	(0.130)
Advanced Education	-0.0817	0.250**	-0.0878	0.331***	-0.125	0.215
	(0.0997)	(0.108)	(0.110)	(0.117)	(0.144)	(0.148)
Dropout	-0.144	-0.429***	-0.125	-0.406***	-0.284	-0.663***
	(0.153)	(0.147)	(0.159)	(0.156)	(0.242)	(0.256)
White	0.380***	-0.341***	0.394***	-0.337***	0.261*	-0.198
	(0.103)	(0.102)	(0.120)	(0.111)	(0.137)	(0.161)
Constant	-2.481***	0.306	-2.333***	0.0788	-2.075***	-0.188
	(0.381)	(0.694)	(0.358)	(0.627)	(0.716)	(1.095)
Observations	1552	1552	1151	1151	1028	1028

Notes: In the table above we report the (second-stage) results using instrumental variables (Probit IV), where we regress individual voting patterns during the 2000 U.S. Presidential election as reported in the 2000 ANES on Lewinsky coverage (instrumented by weather deaths), as well as a host of control variables (in the table above we only include those controls which yield significant results; a full list of controls is provided in Appendix A1). In this case, for every congressional district we weight the amount of Lewinsky coverage afforded by each of the top 2-3 newspapers according to their respective circulation levels. We use two dependent variables - one is a dummy indicating whether the respondent voted for George W. Bush or otherwise, and the other denotes a vote for Al Gore (or otherwise).Columns 1 and 2 refer to the regression described above using the full sample; in columns 3 and 4 we focus only on the subsample of respondents whose political affiliation as recorded in the ANES is non-Republican (i.e. Democrat, Independent or other); while in columns 5 and 6 we analyze the subsample of respondents whose affiliation is non-Democrat (i.e. Republican, Independent or other). All data sources are listed in Appendix A1. Robust standard errors, clustered at the congressional district level, are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4.15: Robustness Tests - Alternative Instrument

Panel A	Full Sample		Non-Republicans		Non-Democrats	
	Bush	Gore	Bush	Gore	Bush	Gore
Lewinsky Coverage	0.482*	-0.633***	0.758***	-0.756***	0.333	-0.611**
	(0.259)	(0.158)	(0.189)	(0.126)	(0.408)	(0.274)
1998 House Republican	1.139***	-0.0690	0.780***	0.0146	1.256***	-0.0424
	(0.247)	(0.106)	(0.282)	(0.117)	(0.263)	(0.141)
1996 Clinton	-0.279**	0.815***	-0.260*	0.688***	-0.106	0.743***
	(0.123)	(0.170)	(0.135)	(0.178)	(0.162)	(0.186)
1996 Dole	0.915***	-0.250	0.552*	-0.100	1.012***	-0.307
	(0.252)	(0.163)	(0.286)	(0.174)	(0.282)	(0.242)
Like Gore	-0.645***	0.618***	-0.561***	0.491***	-0.628***	0.732***
	(0.142)	(0.142)	(0.169)	(0.126)	(0.156)	(0.256)
Like Bush	0.782***	-0.477***	0.601***	-0.360***	0.759***	-0.450***
	(0.166)	(0.126)	(0.189)	(0.113)	(0.159)	(0.171)
Approve Clinton	-0.206*	0.309***	-0.128	0.267**	-0.177	0.252
	(0.111)	(0.118)	(0.114)	(0.116)	(0.135)	(0.160)
Clinton Moral Climate	0.314***	-0.343***	0.336**	-0.330***	0.211	-0.406**
	(0.112)	(0.104)	(0.131)	(0.109)	(0.129)	(0.177)
Protestant	0.369***	-0.0112	0.271**	0.0888	0.378***	-0.116
	(0.116)	(0.103)	(0.135)	(0.107)	(0.137)	(0.150)
Catholic	0.414***	0.0701	0.264	0.183*	0.485***	-0.0904
	(0.143)	(0.108)	(0.173)	(0.107)	(0.164)	(0.165)
Marital Status	0.175	0.0930	0.174	0.0931	0.260*	0.0130
	(0.107)	(0.0824)	(0.118)	(0.0881)	(0.138)	(0.119)
Advanced Education	-0.122	0.209**	-0.0880	0.207*	-0.179	0.197
	(0.112)	(0.104)	(0.119)	(0.111)	(0.137)	(0.145)
Dropout	-0.179	-0.402***	-0.129	-0.302**	-0.336	-0.660***
	(0.171)	(0.144)	(0.169)	(0.149)	(0.245)	(0.253)
White	0.375***	-0.341***	0.375***	-0.302***	0.235	-0.212
	(0.117)	(0.0971)	(0.132)	(0.102)	(0.151)	(0.149)
Constant	-2.133***	0.584	-2.172***	0.720	-1.683*	0.142
	(0.541)	(0.599)	(0.371)	(0.480)	(0.907)	(1.065)
Panel B						
Weather Damage Costs	-0.00175***		-0.00165***		-0.00147***	
	(0.000285)		(0.000327)		(0.000349)	
Constant	2.135***		2.121***		2.095***	
	(0.0322)		(0.0362)		(0.0396)	
F-statistic	37.78		25.60		17.75	
Observations	1552	1552	1151	1151	1028	1028

Notes: Panel B reports the first-stage results from our instrumental variables approach, whereby we regress newspaper coverage of the Lewinsky scandal on the cost in US$ of damages caused by extreme weather events in each of the ANES respondent's congressional districts. In Panel A we report the second-stage results from our Probit IV analysis, where we regress individual voting patterns during the 2000 U.S. Presidential election as reported in the 2000 ANES on Lewinsky coverage (instrumented by weather costs), as well as a host of control variables (in the table above we only include those controls which yield significant results; a full list of controls is provided in Appendix A1). We use two dependent variables - one is a dummy indicating whether the respondent voted for George W. Bush or otherwise, and the other denotes a vote for Al Gore (or otherwise). Columns 1 and 2 refer to the regression described above using the full sample; in columns 3 and 4 we focus only on the subsample of respondents whose political affiliation as recorded in the ANES is non-Republican (i.e. Democrat, Independent or other); while in columns 5 and 6 we analyze the subsample of respondents whose affiliation is non-Democrat (i.e. Republican, Independent or other). All data sources are listed in Appendix A1. Robust standard errors, clustered at the congressional district level, are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

tory variable. The rationale behind these tests is that if either instrument were somehow related to respondent voting patterns via some direct channel other than through Lewinsky coverage then the slope coefficients for the instruments when included as explanatory variables would be statistically significant (these are the coefficients reported in Table 4.16, where for brevity we only report the results for the full sample). This approach is preferable since standard overidentification tests typically have low power and only test for the joint-exogeneity of the instruments used assuming one of them is indeed exogenous. As seen in Table 4.16, the point estimates obtained are all statistically insignificant across all columns.

Table 4.16: Robustness Tests - Overidentification Test

	Full Sample			
	Bush	Gore	Bush	Gore
Weather Deaths	-0.0135	-0.0123		
	(0.0260)	(0.0113)		
Weather Damage Costs			0.000588	0.00110
			(0.000911)	(0.00138)
Observations	1552	1552	1552	1552

In this table we test the validity of our instruments by implementing a modified overidentification test. The test involves running the usual IV regressions using both extreme weather deaths and damage costs as our instruments and including each instrument separately as an explanatory variable (in columns 1 and 2 respectively). The results reported above focus solely on these added instruments as explanatory variables - we omit all the other coefficients obtained from the IV regressions, although the explanatory variables used above are the same as those used throughout this paper. We use two dependent variables - one is a dummy indicating whether the respondent voted for George W. Bush or otherwise, and the other denotes a vote for Al Gore (or otherwise). All data sources are listed in Appendix A1. Robust standard errors, clustered at the congressional district level, are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

4.6 Concluding Remarks

In this paper we investigate the extent to which soft media reportage influences voting patterns. For this purpose we look at the infamous Monica Lewinsky sex scandal which erupted in the U.S. during the Presidency of Bill Clinton, and how newspaper reports of the scandal impacted on the outcome of the 2000 Presidential election. We first sought to analyse the main determinants of media reportage, and particularly the main demand/supply side factors affecting media political slant as promulgated in the literature, including readers' beliefs/opinions, newspaper ownership and editorial slant. We find that Lewinsky coverage was not driven by partisan media slant: this independence thus enables us to focus solely on the electoral impact of soft news reportage, abstracting away from all considerations related to political bias, in contrast to the majority of the literature (e.g. DellaVigna and Kaplan, 2007).

To account for any endogeneity issues we adopt an instrumental variables

approach, where the identification strategy employed exploits variation in the number of deaths caused by extreme weather events over the period under review for each of the newspapers' main counties of circulation. The results show that overall, media coverage had a positive and statistically significant impact on the likelihood of voting for George W. Bush in the 2000 Presidential election, and conversely a negative influence on the probability of voting for Al Gore. If we look at different subgroups within our dataset, then this pattern is visible among both Democrats and Republicans, with the effect strongest among the former subgroup. In fact, a 5% increase in newspaper coverage of the Lewinsky scandal could have increased the likelihood that a non-Republican voted for Bush in the 2000 Presidential election by around 1.9%, with a similar decline reported in the probability of voting for Gore. Various tests are conducted in order to check the robustness and validity of the main results, including checking the validity of survey responses, controlling for newspaper editorial slant, the use of a different instrument, testing for attrition bias, as well as the use of a placebo treatment.

The main conclusion that emerges from this paper is that soft media reports can have a significant impact on individual voting behaviour. This raises several questions regarding the role of the media in the democratic process, particularly since the media has traditionally been considered as a vital informational link between the political classes and the general populace (Eisensee and Stromberg, 2007). The rise of soft news or sensationalism has already diluted this role, as can be observed from dwindling public trust in media organisations over the years (Pew, 2012c). This paper suggests that such reportage may also have an impact on people's voting behaviour, despite the fact that these stories provide little meaningful information regarding the candidates' electoral credentials.

Various explanations have been proposed in order to justify this increase in media coverage of soft news. Hamilton (2004) states that the increased focus on such stories is driven by numerous factors including ownership changes, deregulation of media markets and increased competition among cable TV channels. Bagdikian (2000) on the other hand maintains that such reportage is fuelled by a desire to appease advertisers, since readers would be more attentive to advertising signals if the news content on offer were "dumbed-down" as described by the author. It would be interesting to analyse these factors individually in more detail in order to understand what, if anything, can be done by policymakers to change the current situation.[24]

[24]For example, Ellman (2014) suggests that government subsidies to media outlets may be a reasonable policy prescription given the positive externalities emanating from the consumption of informative news.

4.7 Appendices

Appendix A1: Description of Variables

Variable Name	Description of Variable	Source
Bush	Dummy denoting whether respondent voted for George W. Bush in the 2000 U.S. Presidential election	ANES 2000
Gore	Dummy denoting whether respondent voted for Al Gore in the 2000 U.S. Presidential election	ANES 2000
No Vote	Dummy denoting whether respondent abstained from voting in the 2000 U.S. Presidential election	ANES 2000
Lewinsky Coverage	Number of published newspaper articles on the Monica Lewinsky scandal over the period January 17, 1998 to August 31, 2000, expressed as a fraction of the total number of articles published by each newspaper over the same period	NewsLibrary, ProQuest, archives, 2014
Circulation	Number of newspaper copies sold per week for each newspaper	Alliance for Audited Media, 2000
New England	Dummy denoting whether respondent's county is located in the New England region of the U.S.	ANES 2000
Mid-Atlantic	Dummy denoting whether respondent's county is located in the Mid-Atlantic region of the U.S.	ANES 2000
East North-Central	Dummy denoting whether respondent's county is located in the East North-Central region of the U.S.	ANES 2000
West North-Central	Dummy denoting whether respondent's county is located in the West North-Central region of the U.S.	ANES 2000
East South-Central	Dummy denoting whether respondent's county is located in the East South-Central region of the U.S.	ANES 2000
West South-Central	Dummy denoting whether respondent's county is located in the West South-Central region of the U.S.	ANES 2000
Mountain	Dummy denoting whether respondent's county is located in the Mountain region of the U.S.	ANES 2000
Pacific	Dummy denoting whether respondent's county is located in the Pacific region of the U.S.[25]	ANES 2000
1998 House Republican	Dummy denoting whether respondent voted for a Republican candidate in the 1998 House of Representatives election	ANES 2000
Gender	Dummy denoting whether the respondent is male	ANES 2000

[25]Omitted category is South Atlantic

Variable Name	Description of Variable	Source
Children	Dummy denoting whether the respondent has any children	ANES 2000
1996 Dole	Dummy denoting whether respondent voted for Bob Dole in the 1996 U.S. Presidential election	ANES 2000
1996 Clinton	Dummy denoting whether respondent voted for Bill Clinton in the 1996 U.S. Presidential election	ANES 2000
Like Bush	Dummy denoting whether respondent expressed a liking for George W. Bush	ANES 2000
Like Gore	Dummy denoting whether respondent expressed a liking for Al Gore	ANES 2000
Approve Clinton	Dummy denoting whether respondent approved of Bill Clinton's performance as U.S. President	ANES 2000
Clinton Moral Climate	Dummy denoting whether respondent believed that Bill Clinton's Presidency had harmed the country's moral climate	ANES 2000
Protestant	Dummy denoting whether respondent is a Protestant	ANES 2000
Catholic	Dummy denoting whether respondent is a Catholic	ANES 2000
Other Religion	Dummy denoting whether respondent belongs to another formal religious organisation (including Judaism, Islam, etc.)[26]	ANES 2000
Marital status	Dummy denoting whether respondent is married	ANEs 2000
Advanced Education	Dummy denoting whether respondent has an advanced (i.e. tertiary) level of education	ANES 2000
Dropout	Dummy denoting whether respondent ceased formal education prior to high school graduation[27]	ANES 2000
Employed	Dummy denoting whether respondent is gainfully occupied	ANES 2000
Income < $10000	Dummy denoting whether respondent's annual income is below $10,000	ANES 2000
Income $10000-$34999	Dummy denoting whether respondent's annual income is between $10,000 and $34,999	ANES 2000
Income $35000-$64999	Dummy denoting whether respondent's annual income is between $35,000 and $64,999	ANES 2000
Income $65000-$94999	Dummy denoting whether respondent's annual income is between $65,000 and $94,999	ANES 2000
Income $94999-$249999	Dummy denoting whether respondent's annual income is between $95,000 and $249,999[28]	ANES 2000

[26]Omitted category in this case is for respondents with no religious beliefs (i.e. Atheists).
[27]Omitted category in this case is for respondents with a high school diploma.
[28]Omitted category in this case is for respondents with annual income of $250,000 and over.

Variable Name	Description of Variable	Source
White	Dummy denoting whether respondent's race is Caucasian	ANES 2000
Age 18-24	Dummy denoting whether respondent's age is between 18 and 24 years	ANES 2000
Age 25-34	Dummy denoting whether respondent's age is between 25 and 34 years	ANES 2000
Age 35-44	Dummy denoting whether respondent's age is between 35 and 44 years	ANES 2000
Age 45-54	Dummy denoting whether respondent's age is between 45 and 54 years	ANES 2000
Age 55-64	Dummy denoting whether respondent's age is between 55 and 64 years[29]	ANES 2000
Editorial Slant	Index ranging from 0-100 denoting the ideological slant of each newspaper. The index compares the number of times a newspaper refers to a particular think-tank to the number of times they are cited by members of Congress. Low scores (below 50) denote conservatism, while higher scores (over 50) denote a liberal slant.	Groseclose and Milyo (2005)
Gannett	Dummy denoting whether newspaper was owned by Gannett in 2000	Individual newspaper websites
Knight Ridder	Dummy denoting whether newspaper was owned by Knight Ridder in 2000	Individual newspaper websites
Advance	Dummy denoting whether newspaper was owned by Advance Communications in 2000[30]	Individual newspaper websites
Broadsheet	Dummy denoting whether newspaper was a broadsheet in 2000	Individual newspaper websites

[29]In this case the omitted category is for respondents above the age of 65.

[30]Omitted category in this case is all other newspaper owners (like NewsCorp or Tribune).

Appendix A2: Descriptive Statistics of Explanatory Variables

Variable	Mean	Std. Dev.	Min.	Max.
Circulation	213967.921	321625.662	5235	2000000
New England	0.058	0.234	0	1
Mid-Atlantic	0.115	0.319	0	1
East North-Central	0.174	0.379	0	1
West North-Central	0.079	0.269	0	1
East South-Central	0.069	0.253	0	1
West South-Central	0.127	0.333	0	1
Mountain	0.063	0.242	0	1
Pacific	0.153	0.36	0	1
1998 House Republican	0.285	0.452	0	1
Gender	0.434	0.496	0	1
Children	0.738	0.44	0	1
1996 Clinton	0.372	0.484	0	1
1996 Dole	0.243	0.429	0	1
Like Gore	0.512	0.5	0	1
Like Bush	0.467	0.499	0	1
Approve Clinton	0.637	0.481	0	1
Clinton Moral Climate	0.214	0.41	0	1
Protestant	0.312	0.463	0	1
Catholic	0.2	0.4	0	1
Other Religion	0.171	0.377	0	1
Marital Status	0.523	0.5	0	1
Advanced Education	0.414	0.493	0	1
Dropout	0.095	0.294	0	1
Employed	0.631	0.483	0	1
Income < $10000	0.224	0.417	0	1
Income $10000-$34999	0.264	0.441	0	1
Income $35000-$64999	0.249	0.432	0	1
Income $65000-$94999	0.151	0.358	0	1
Income $95000-$249999	0.064	0.244	0	1
White	0.794	0.404	0	1
Age 18-24	0.077	0.266	0	1
Age 25-34	0.162	0.368	0	1
Age 35-44	0.233	0.423	0	1
Age 45-54	0.186	0.389	0	1
Age 55-64	0.151	0.358	0	1
N		1552		

Appendix B: Marginal Effects for IV Regression Results in Table 4.4 (Panel B)

	Full Sample		Non-Republicans		Non-Democrats	
	Bush	Gore	Bush	Gore	Bush	Gore
Lewinsky Coverage	0.2245***	-0.1754***	0.1948**	-0.1844**	0.1959	-0.0891
1998 House Republican	0.3520***	-0.0340	0.2237***	-0.0117	0.4245***	-0.0115
1996 Clinton	-0.0873**	0.3311***	-0.0613**	0.3604***	-0.0543	0.1924***
1996 Dole	0.2717***	-0.1134**	0.1517**	-0.1112	0.3382***	-0.0697**
Like Gore	-0.1940***	0.2438***	-0.1400***	0.2549***	-0.2269***	0.1774***
Like Bush	0.2351***	-0.1907***	0.1567***	-0.1933***	0.2750***	-0.1023***
Approve Clinton	-0.0601*	0.1232***	-0.0334	0.1448***	-0.0565	0.0575**
Clinton Moral Climate	0.1061***	-0.1232***	0.0865**	-0.1554***	0.0815*	-0.0742***
Protestant	0.1272***	0.0073	0.0706**	0.0672	0.1527***	-0.0178
Catholic	0.1352***	0.0333	0.0664*	0.0833*	0.1897***	-0.0107
Marital Status	0.0516	0.0368	0.0417	0.0505	0.0923*	0.0018
Advanced Education	-0.0306	0.0930**	-0.0202	0.1329***	-0.0597	0.0467*
Dropout	-0.0467	-0.1423***	-0.0282	-0.1621***	-0.1080	-0.0895**
White	0.1244***	-0.1256***	0.0879***	-0.1329***	0.1027*	-0.0397
Observations	1552	1552	1151	1151	1028	1028

Notes: The table above shows the computed marginal effects, evaluated at the mean, for the second-stage results of the instrumental variables (IV) regression analysis. We use two dependent variables - one is a dummy indicating whether the respondent voted for George W. Bush or otherwise, and the other denotes a vote for Al Gore (or otherwise). Columns 1 and 2 refer to the regression described above using the full sample; in columns 3 and 4 we focus only on the subsample of respondents whose political affiliation as recorded in the ANES is non-Republican (i.e. Democrat, Independent or other); while in columns 5 and 6 we analyze the subsample of respondents whose affiliation is non-Democrat (i.e. Republican, Independent or other). All data sources are listed in Appendix A1.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Appendix C: Additional Information on the Lewinsky Coverage Dataset

In this section we describe in greater detail the process involved in collating data on the number of newspaper articles on the Lewinsky scandal over the period January 17, 1998 to August 31, 2000. As mentioned in the paper, the choice of newspapers was largely determined by the geographical location of the survey respondents in the 2000 ANES Time Series Study, whereby the top 2-3 circulating newspapers in each of the respondents' congressional district of residence were identified using circulation data from the Alliance of Audited Media as well as (where necessary) advertiser information provided by the newspapers themselves.

Three main resources were used in order to acquire the actual Lewinsky coverage data, namely NewsBank's NewsLibrary database, ProQuest and the individual newspapers' online archives. In line with other papers within this literature (e.g. Gentzkow and Shapiro, 2010) we develop a suitable Boolean string-search protocol in order to identify relevant articles on the Lewinsky scandal. For the purposes of this paper, articles containing the words "Monica Lewinsky" or "Lewinsky" or "Lewinskygate" or "Monicagate" or "tailgate" or "sexgate" or "zippergate" or "Clinton" in the headline, and "Monica Lewinsky" or "Lewinsky" or "Lewinskygate" or "Monicagate" or "tailgate" or "sexgate" or "zippergate" in the main body of text were sought out, excluding those articles containing the words "letter" or "letters", over the period January 17, 1998 to August 31, 2000. The search protocol used can be justified on the basis of the following points:

- The somewhat exhaustive list of search terms reflects the relative 'richness' of the vocabulary used by the media in order to describe the scandal. By incorporating a large number of search terms, it is hoped that omissions of relevant articles would be kept at a minimum across all newspaper formats, since certain terms like "zippergate" were almost exclusively used by tabloid or tabloid-style newspapers as opposed to broadsheets;

- By directing the search towards articles containing the list of terms in both the headline and the main body of text, this enabled us to focus on stories that dealt specifically with the Lewinsky scandal, eliminating those articles which simply alluded to the scandal briefly as part of some turn-of-phrase within the scope of an altogether distinct story[31];

- The approach also minimised the number of Lewinsky scandal news stories appearing as part of a 'News in Brief'-type section, where the actual story would only form a small part of the overall content;

- The exclusion of articles containing the words "letter" or "letters" was done to eliminate all letters to the editor, since it was felt that despite their

[31]Notice that although the search term "Clinton" is included in the headline search, it was excluded from the main body search. This was purposely done in order to exclude any Clinton-related stories that were not related in any way to the Lewinsky scandal.

ubiquity they do not constitute news stories within the context of 'media coverage' as defined in this study.

As an additional precautionary measure, the search results obtained for each newspaper were then individually checked in order to eliminate any invalid articles (like for example news in brief or letters to the editor) which somehow evaded the search protocol. As a general rule, all articles below 200 words in length were eliminated in order to avoid minor stories or news in brief. This entire process was conducted for a total of 257 daily newspapers from across the U.S., including national, state-wide and local newspapers

Bibliography

ABC (2012): *Local butchers benefit from 'pink slime' controversy*. ABC-WHAM, Rochester, New York, March 2012.

Acemoglu, D., S. Johnson, and J. A. Robinson (2001): "The colonial origins of comparative development: An empirical investigation," *American Economic Review*, 91(5), 1369–1401.

Agliata, D., and S. Tantleff-Dunn (2004): "The impact of media exposure on males' body image," *Journal of Social and Clinical Psychology*, 23(1), 7–22.

Anderson, S. P., and J. McLaren (2012): "Media Mergers and Media Bias," *Journal of the European Economic Association*, 10(4), 831–859.

Andreoni, J. (1990): "Impure altruism and donations to public goods: A theory of warm-glow giving," *Economic Journal*, 100(401), 464–477.

Andreoni, J., A. Payne, et al. (2011): "Diversity and donations: The effect of religious and ethnic diversity on charitable gift giving," *NBER Working Paper*, 17618.

Bagdikian, B. (2000): *The media monopoly, 6th Edition*. Beacon Press: Boston, MA.

Bagwell, K. (2007): *The economic analysis of advertising*. Handbook of Industrial Organization 3, 1701-1844.

Barak, G. (1994): *Media, process and the social constriction of crime: Studies in newsmaking criminology*. New York and London: Garland Publishing.

Barnett, S., G. N. Ramsay, and I. Gaber (2012): *From Calaghan to Credit Crunch: Changing trends in British television news, 1975-2009*. University of Westminster: London.

Baron, D. (2006): "Persistent media bias," *Journal of Public Economics*, 90(1-2), 1–36.

Bartels, L. M. (1993): "Messages received: the political impact of media exposure.," *American Political Science Review*, 87(2), 267–285.

Baum, M. A. (2003): "Soft news and political knowledge: Evidence of absence, or absence of evidence?," *Political Communication*, 20, 173–190.

Becker, E., and S. Strom (2005): "Groups pledge to account for tsunami aid," *New York Times*.

Bennett, W. L., R. G. Lawrence, and S. Livingston (2008): *When the press fails: Political power and the news media from Iraq to Katrina*. University of Chicago Press.

Berelson, B., et al. (1954): *Voting*. University of Chicago Press.

Besley, T., and R. Burgess (2001): "Political agency, government responsiveness and the role of the media," *European Economic Review*, 45(4-6), 629–640.

Bhatti, Y., et al. (2013): "Political hypocrisy: The effect of political scandals on candidate evaluations," *Acta Politica*, 48, 408–428.

Blackwell, M. (2012): "A framework for dynamic causal inference in political science," *American Journal of Political Science*, 57(2), 504–520.

Blasco, A., P. Pin, and F. Sobbrio (2015): "Paying positive to go negative: Advertisers' competition and media reports," *Working Paper*.

Blasco, A., and F. Sobbrio (2012): "Competition and commercial media bias," *Telecommunications Policy*, 36(5), 434–447.

Bloch, F., and D. Manceau (1999): "Persuasive advertising in Hotelling's model of product differentation," *International Journal of Industrial Organization*, 17(4), 557–574.

Bogart, L. (1989): *Press and public, 2nd Edition*. Erlbaum: Hillsdale, NJ.

Boykoff, M. T., and J. M. Boykoff (2004): "Balance as bias: global warming and the US prestige press," *Global Environmental Change*, 14, 125–136.

Brown, S., M. N. Harris, and K. Taylor (2012): "Modeling charitable donations to an unexpected natural disaster: Evidence from the U.S. Panel Study of Income Dynamics," *Journal of Economic Behavior and Organization*, 84(1), 97–110.

Chiang, C.-F., and B. Knight (2011): "Media bias and influence: Evidence from newspaper endorsements," *Review of Economic Studies*, 78(3), 795–820.

Chomsky, N., and E. Herman (1988): *Manufacturing consent: The political economy of the mass media*. Pantheon Books; New York.

Cragg, J. (1971): "Some statistical models for limited dependent variables with application to the demand for durable goods," *Econometrica*, 39, 829–844.

Crampes, C., et al. (2009): "Advertising, competition and entry in media industries," *Journal of Industrial Economics*, 57(1), 7–31.

DellaVigna, S., and E. Kaplan (2007): "The Fox News Effect: Media bias and voting," *Quarterly Journal of Economics*, 122(3), 1187–1234.

Diamond, W. D., and E. S. Iyer (2007): "Creating effective direct mail charitable solicitations: The effects of enclosures and different appeals," *Journal of Nonprofit and Public Sector Marketing*, 18, 81–100.

Doms, M., and N. Morin (2004): "Consumer sentiment, the economy and the news media," *Federal Reserve: Finance and Economics Discussion Series Working Paper*, 51, 1–70.

Druckman, J. N., and M. Parkin (2005): "The impact of media bias: How editorial slant affects voters," *Journal of Politics*, 67(4), 1030–1049.

Dukes, A. (2004): "The advertising market in a product oligopoly," *Journal of Industrial Economics*, 52(3), 327–348.

Dyck, A., N. Volchova, and L. Zingales (2008): "The corporate governance role of the media: Evidence from Russia," *Journal of Finance*, 63(3), 1093–1135.

Eisensee, T., and D. Stromberg (2007): "News floods, news droughts, and U.S. disaster relief," *Quarterly Journal of Economics*, 122(2), 693–728.

Ellman, M. (2014): "Softness bias in the news: Optimal subsidies, price floors and competitive threats," *Unpublished manuscript*.

Ellman, M., and F. Germano (2009): "What do the papers sell? A model of advertising and media bias," *Economic Journal*, 119(537), 680–704.

Erbring, L., E. N. Goldenburg, and A. H. Miller (1980): "Front-page news and real-world cues: A new look at agenda setting by the media," *American Journal of Political Science*, 24(1), 16–49.

FBI (2005): *Uniform Crime Statistics, 2000-2005: Violent Crime (by County), Federal Bureau of Investigation*. Uniform Crime Reporting Program (UCR).

——— (2015): *Uniform Crime Statistics, 2005-2013: Violent Crime (by County), Federal Bureau of Investigation*. Uniform Crime Reporting Program (UCR).

Fico, F., and W. Cote (1999): "Fairness and balance in the structural characteristics of newspaper stories on the 1996 Presidential election," *Journalism and Mass Communication Quarterly*, 76(1), 124–137.

Fischle, M. (2000): "Mass response to the Lewinsky scandal: Motivated reasoning or Bayesian updating?," *Political Psychology*, 21(1), 135–159.

Fletcher, R. (2003): "Adverts in medical journals: caveat lector," *The Lancet*, 361, 10–11.

Fridkin, K. L., and P. Kenney (2011): "Variability in citizens' reaction to different types of negative campaigns," *American Journal of Political Science*, 55(2), 307–325.

Gal-Or, E., T. Geylani, and T. P. Yildirim (2012): "The impact of advertising on media bias," *Journal of Marketing Research*, 41(1), 92–99.

Gallup (2011): *Gallup 2011 Environment Poll*. Gallup Historical Trends.

Geer, J. G. (2004): *Public opinion and polling around the world: A historical encyclopedia*. ABC-CLIO.

Gentzkow, M., and J. M. Shapiro (2006): "Media bias and reputation," *Journal of Political Economy*, 114(2), 280–316.

——— (2008): "Competition and truth in the market for news," *Journal of Economic Perspectives*, 22(2), 133–154.

——— (2010): "What drives media slant? Evidence from U.S. daily newspapers," *Econometrica*, 78(1), 35–71.

Gerber, A. S., D. Karlan, and D. Bergan (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(2), 35–52.

Germano, F., and M. Meier (2010): "Concentration and self-censorship in commercial media," *Barcelona Economics Working Paper Series*, No. 527.

Goldstein, K., and P. Freedman (2002): "Campaign advertising and voter turnout: New evidence for a stimulation effect," *Journal of Politics*, 64(3), 721–740.

Groseclose, T., S. D. Levitt, and J. M. Snyder (1999): "Comparing interest group scores across time and chambers: Adjusted ADA scores for the U.S. Congress," *The American Political Science Review*, 93(1), 33–50.

Groseclose, T., and J. Milyo (2005): "A measure of media bias," *Quarterly Journal of Economics*, 120(4), 1191–1237.

Hamilton, J. T. (2004): *All the news that's fit to sell: How the market transforms information into news*. Princeton University Press.

Hansen, S., et al. (2014): "Transparency and deliberation within the FOMC: A computational linguistics approach," .

Harrison, K., and J. Cantor (2006): "The relationship between media consumption and eating disorders," *Journal of Communication*, 47(1), 40–67.

Hartman, C. W., and G. Squires (2006): *There is no such thing as a natural disaster: Race, class and Hurricane Katrina*. Routledge: New York.

Healy, A., and N. Malhotra (2010): "Random events, economic losses and retrospective voting: Implications for democratic competence," *Quarterly Journal of Political Science*, 5, 193–208.

Iyengar, S., and D. Kinder (2010): *News that matters: Television and American opinion*. University of Chicago Press: Chicago.

Jaramillo, E. (2001): "The impact of media-based health education on tuberculosis diagnosis in Cali, Colombia," *Health policy and planning*, 16(1), 68–73.

Johnson, J. P., and D. P. Myatt (2006): "On the simple economics of advertising, marketing and product design," *American Economic Review*, 96(3), 756–784.

Kantar (2014): *U.S. Advertising Expenditures 2014*. Kantar Media: New York.

Kerkhof, A., and J. Muenster (2015): "Quantity restrictions on advertising, commercial media bias, and welfare," *Unpublished Maunscript*.

Kiousis, S. (2003): "Job approval and favorability: The impact of media attention to the Monica Lewinsky scandal on public opinion of President Bill Clinton," *Mass Communication and Society*, 6(4), 435–451.

Kitchen, H. (1992): "Determinants of charitable donations in Canada: A comparison over time," *Applied Economics*, 24(7), 709–713.

Klapper, J. T. (1960): *The effects of mass communication*. Free Press.

Krupnikov, Y. (2011): "When does negativity demobilize? Tracing the conditional effect of negative campaigning on voter turnout," *American Journal of Political Science*, 55(4), 797–813.

Lamont, M., N. Hing, and S. Gainsbury (2011): "Gambling on sport sponsorship: A conceptual framework for research and regulatory review," *Sport Management Review*, 14, 246–257.

Lau, R. R., L. Sigelman, and I. B. Rovner (2007): "The effects of negative political campaigns: A meta-analytic reassessment," *Journal of Politics*, 69(4), 1176–1209.

Lawrence, R. G., and W. Bennett (2001): "Rethinking media politics and public opinion: Reactions to the Clinton-Lewinsky scandal," *Political Science Quarterly*, 116(3), 425–446.

Lazarsfeld, P., et al. (1944): *The People's Choice*. Columbia University Press.

Lobb, A., N. Mock, and P. Hutchinson (2012): "Traditional and social media coverage and charitable giving following the 2010 earthquake in Haiti," *Prehosp Disaster Med.*, 27(4), 319–324.

Lord, C. G., L. Ross, and M. R. Lepper (1979): "Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered efforts," *Journal of Personality and Social Psychology*, 37(11), 2098–2109.

MacKuen, M., and S. Combs (1981): *More than news*. Sage: Beverly Hills, California.

Martin, G. J., and A. Yurukoglu (2014): "Bias in cable news: Real effects and polarization," Discussion paper, National Bureau of Economic Research.

McRoberts, B., and G. H. Terhanian (2008): *Digital influence index study*. Fleishman-Hillard and Harris Interactive.

Mermin, J. (1997): "Television news and American intervention in Somalia: The myth of a media-driven foreign policy," *Political science quarterly*, pp. 385–403.

Milgrom, P. (1981): "Good news and bad news: Representation theorems and applications," *The Bell Journal of Economics*, 12(2), 380–391.

Milgrom, P., and J. Roberts (1986): "Price and advertising signals of product quality," *Journal of Political Economy*, 94(4), 796–821.

Mullainathan, S., and A. Shleifer (2005): "The market for news," *American Economic Review*, 95(4), 1031–1053.

NAA (2012): *Annual advertising and circulation expenditures, 1990-2009*. Newspaper Association of America.

Nelson, P. (1974): "Advertising as information," *Journal of Political Economy*, 82, 729–754.

Niven, D. (2003): "Objective evidence on media bias: Newspaper coverage of congressional party switchers," *Journalism & Mass Communication Quarterly*, 80(2), 311–326.

Oreskes, N., and E. M. Conway (2010): *Merchants of Doubt*. Bloomsbury Press: New York.

Park, S.-Y. (2005): "The influence of presumed media influence on women's desire to be thin," *Communication Research*, 32(5), 594–614.

Peitz, M., and T. M. Valletti (2008): "Content and advertising in the media: Pay-TV versus free-to-air," *International Journal of Industrial Organization*, 26(4), 949–965.

Petrova, M. (2011): "Newspapers and parties: How advertising revenues created an independent press," *American Political Science Review*, 105(4), 790–808.

Pew (2007): *2007 Public Knowledge Survey*. Pew Research Center for the People and the Press.

——— (2012a): *2012 Media Consumption Survey*. Pew Research Center for the People and the Press.

——— (2012b): *Political Typology Survey*. Pew Research Center.

——— (2012c): *The state of the news media: An annual report on American Journalism*. Pew Research Center's Project for Excellent in Journalism.

——— (2013): *July 2013 Political Survey*. Pew Research Center for the People and the Press.

Phillips, D. P. (1983): "The impact of mass media violence on US homicides," *American Sociological Review*, 48, 560–568.

Prat, A., and D. Stromberg (2005): "Commercial television and voter information," *CEPR Discussion Paper*, 4989.

Prat, A., and D. Stromberg (2011): "The political economy of mass media," *CEPR Discussion Paper*, DP8246.

Prior, M. (2003): "Any good news in soft news? The impact of soft news preference on political knowledge," *Political Communication*, 20, 149–171.

Randolph, W. C. (1995): "Dynamic income, progressive taxes and the timing of charitable contributions," *Journal of Political Economy*, 103(4), 709–738.

Reuter, J., and E. Zitzewitz (2006): "Do ads influence editors? Advertising bias in the financial media," *Quarterly Journal of Economics*, 121(1), 197–227.

Rochet, J.-C., and J. Tirole (2003): "Platform competition in two-sided markets," *Journal of the European Economic Association*, 1(4), 990–1029.

Rudzitis, G. (1999): "Amenities increasingly draw people to the rural west," *Rural development perspectives*, 14, 9–13.

Schmitt-Beck, R. (2003): "Mass communication, personal communication and vote choice: The filter hypothesis of media influence in comparative perspective," *British Journal of Political Science*, 33(2), 233–259.

Shields, T. G., R. K. Goidel, and B. Tadlock (1995): "The net impact of media exposure on individual voting decisions in US senate and house elections," *Legislative Studies Quarterly*, pp. 415–430.

Sparks, C., J. Tulloch, et al. (2000): *Tabloid tales: Global debates over media standards*. Rowman and Littlefield Publishers.

Stabile, C. A., and M. Harrison (2003): *Prime time animation: Television animation and American culture*. Psychology Press.

Stephens, M. (2007): *A History of News*. Oxford University Press: New York.

Stromberg, D. (2004): "Radio's impact on public spending," *Quarterly Journal of Economics*, 119(1), 189–221.

Sturken, M. (2001): "Desiring the weather: El Nino, the media and California identity," *Public Culture*, 13(2), 161–189.

Sutter, D. (2002): "Advertising and political bias in the media: The market for criticism of the market economy," *American Journal of Economics and Sociology*, 61(3), 725–745.

Tella, R. D., and I. Franceschelli (2011): "Government advertising and media coverage of political scandals," *American Economic Journal: Applied Economics 3*, 4, 119–51.

Telser, L. G. (1964): "Advertising and competition," *Journal of Political Economy*, 72, 537–562.

Timpone, R. J. (1998): "Structure, behavior and voter turnout in the United States," *The American Political Science Review*, 92(1), 145–158.

Trenaman, D., and J. McQuail (1961): *Television and the political image*. Methuen.

Ungar, S. (1999): "Is strange weather in the air? A Study of U.S. national network news coverage of extreme weather events," *Climatic Change*, 41(2), 133–150.

USDA (2003): *The Rural-Urban Continuum Codes*. Economic Research Services.

Villani, S. (2001): "Impact of media on children and adolescents: a 10-year review of the research," *Journal of the American Academy of Child & Adolescent Psychiatry*, 40(4), 392–401.

Warner, K., L. Goldenhar, and C. McLaughlin (1992): "Cigarette advertising and magazine coverage of the hazards of smoking: A statistical analysis," *New England Journal of Medicine*, 326, 305–309.

Wilhelm, M. O. (2006): "New data on charitable giving in the PSID," *Economics Letters*, 92, 26–31.

Williams, B. A., and M. X. D. Carpini (2000): "Unchained reaction: The collapse of media gatekeeping and the Clinton-Lewinsky scandal," *Journalism*, 1(1), 61–85.

Williams, P., and J. Dickinson (2003): "Fear of Crime: Read all about it? The relationship between newspaper crime reporting and fear of crime," *British Journal of Criminology*, 33, 35–56.

Yioutas, J., and I. Segvic (2003): "Revisiting the Clinton/Lewinsky scandal: The convergence of agenda-setting and framing," *Journalism and Mass Communication Quarterly*, 80(3), 567–582.