

Local News Online: Aggregators, Geo-Targeting and the Market for Local News**

Lisa M. George[†]

Hunter College and the Graduate Center, CUNY

December 18, 2012

Abstract

This study examines the effect of adding geo-targeted local news links to the Google News web page on household visits to local news outlets. Using a sample of news visits by 24,139 US households from April through September 2010, results indicate that adding local news links to Google News increased both the level of local news consumption and the local share of news consumed online. However the magnitude of the effect is quite small, with the likelihood of a local news visit increasing by only 4-6% from a low baseline level. Adding geo-targeted links increased the variety of local sites visited per day, but not the number of unique sites visited per month, suggesting that increases in local news consumption arise from more frequent visits to familiar news outlets rather than visits to additional news providers. Results suggest that geo-targeting and other technology to reduce the costs of accessing local news are not likely to have an economically meaningful impact on local media outlets in the US.

1. Introduction

The role of news aggregators and other intermediaries such as Google News, Yahoo, and the Huffington Post have been a subject of fierce debate for close to a decade. Newspapers and other creative enterprises argue in the language of piracy, challenging the right of outsiders to collect, link and re-post news in ways that divert advertising impressions from original sources. Digital media giants such as Google but also a long tail of bloggers and internet advocates counter with arguments on the value of intermediaries in reducing search costs for consumers, the potential to raise news consumption overall and improve matches between consumers and content.

Economic research has only recently begun to work through the ways that intermediaries affect the market for news. Economic theory has advanced a few models exploring how both the substitution effect embedded in the piracy discussion and the complementarity effect embedded in digital media arguments can both operate. George and Hogendorn (2012) highlight the importance of transaction costs in news consumption, illustrating how aggregators can increase news consumption in ways different than search. Jeon and Esfahani (2012) show how

* This paper was written with financial support from the Time Warner Cable Research Program on Digital Communications. The views expressed are those of the author and not necessarily those of Time Warner Cable or the Time Warner Cable Research Program on Digital Communications.

[†] Lisa M. George, Associate Professor of Economics, Hunter College and the Graduate Center, CUNY, 695 Park Avenue, New York, NY 10065. Email: lisa.george@hunter.cuny.edu.

consumer preferences for quality can lead aggregators to increase or decrease demand for original content. Dellarocas et al. (2010) study incentives to link content within and across media.

Empirical research to date indicates that aggregators can indeed increase traffic for major producers. Chiou and Tucker (2011) show that a contract dispute between the associated press and Google which removed AP content from Google News for a short period in early 2010 reduced demand. On the consumer side, Athey and Mobius (2012) show that users who adopt a localization feature of Google News in France increase their consumption of local news in the short run, but over time most additional local news consumption derives from increased use of Google News.

An important element in understanding the role of aggregators is the extent to which intermediaries systematically *shift* consumption across media outlets rather than directly substitute for or complement original news. If aggregators reduce search costs uniformly, allowing consumers ready access to content previously too difficult to find, then observed shifts in readership are likely efficient. However if intermediaries reduce the costs of consuming some types of news relative to others, consumers may switch to less-preferred but more readily available material rather than search. While this outcome might be cost-minimizing overall, the re-allocation of attention changes relative demand for different media types and can have important effects on competition.

The effect of technology on relative demand for news has been documented in traditional media markets. George (2008) showed how the spread of the internet altered the composition of the audience for traditional local newspapers, pulling younger and more educated readers out of local newspaper markets. George and Waldfogel (2005) showed how national expansion of the *New York Times* made possible by satellite printing attracted highly educated readers away from local newspapers, with consequences for local media markets and also for local voting (George and Waldfogel 2008).

In the context of digital news, a particular concern in the US is that intermediaries have reduced the costs of locating and consuming national information much more than the cost of locating and consuming local content, facilitating readership shifts to national media that harm local outlets. Anecdotal evidence can be found in often cited figures that a small number of national news sites (*CNN*, *Fox News*, the *Wall Street Journal*, and the *New York Times*) attract the vast majority of internet news views. Until very recently, few aggregators or search engines offered tools to systematically and accurately identify local content.

The focus of this study is whether and how much technologies that reduce the cost of consuming local content raise consumption of local news. Using a July 2010 modification to Google News that automatically placed geo-targeted local news links on the page, I estimate whether access to local links increased visits to local news outlets in the US overall and as a share of total news visits. Results indicate that adding geo-targeted links to Google News did increase local news consumption and the local share of news consumed online. However the magnitude of the effect is quite small, with the likelihood of a local news visit increasing by only 4-6% from a low

baseline level. Adding geo-targeted links increased the variety of local sites visited per day, but not the number of unique sites visited per month, suggesting that increases in local news consumption arise from more frequent visits to familiar news outlets rather than visits to additional news providers. There is also evidence that the effect of adding geo-targeted local news links erodes over time. Overall, the research suggests that geo-targeting and other technology to reduce the costs of accessing local news are not likely to have an economically meaningful impact on local media outlets in the US.

In addition to interest by media firms, the analyses presented here are relevant to policy at several levels. Strong provisions have long been in place in the US to promote local media, and localism is one of the three principles (along with diversity and competition) guiding Federal Communication Commission policy. Technology that reduces barriers to local news consumption, even with small effects, are of great interest in reforming policy to reflect modern markets. This research is also related to debate on privacy standards on the internet. While older personalization features of Google News and other aggregators allowed consumers with an interest in local media to “opt in” by providing geographic information, the current version of Google News studied here automatically identifies local content based on IP addresses and other geo-targeting technologies. The low baseline of local news consumption among Google News users during the “opt in” period and the measurable increase with automatic targeting suggests that these technologies can have modest positive social effects that can offset privacy concerns that much of the debate.

The remainder of the report proceeds as follows. Section 2 describes the research design. Section 3 outlines construction of the data. Sections 4 and 5 describe the empirical specifications and present results. Section 6 concludes the study.

2. Research Design

On June 30, 2010, Google News introduced the first and only comprehensive redesign of the Google News page since the start of the service. The redesign altered the presentation of content in several ways and allowed users to prioritize subject areas. Most important for this study, the redesign added a strip of content on the right side of the page with a set of local headlines and local news links, with the location of the user identified automatically through the IP address. While some customization of topics and local content had been possible on the site since early 2008, these were opt-in features requiring registration, log-in and user input. After June 30, 2010, local content was reported automatically through geo-targeting technology and could not be removed through customization.

Figure 1 shows a screen shot of a Google News page on June 28, 2010 before the design change and Figure 2 shows a screen shot on July 2, 2010 just after the change. The June 28 page reveals an area where viewers can enter preferences for topics or local content. This screenshot, scraped from the Internet Archives “Wayback Machine” shows a non-customized page. The screenshot in Figure 2, taken two days later, reflects the new format. The page still includes personalization options, but local headlines and links are now fixed on the right side of the page, shown with a large arrow (added for clarity). The location of the scraping server is

identified as San Francisco, and the local content includes two headlines from the *San Francisco Chronicle* and one from the *San Jose Mercury News*.

The basic research strategy of this study is to measure the effect of adding local news links to Google News on overall visits to local news sites and the share of news visits to local sites by a sample of internet. In one set of tests, local news consumption patterns of a sample of heavy Google News users is compared to a sample of Yahoo users before and after the redesign. In a second set of tests, local news consumption patterns before and after the redesign are compared based on intensity of Google News use prior to the change.

Figure 1: Google News Page Before Redesign, June 28, 2010

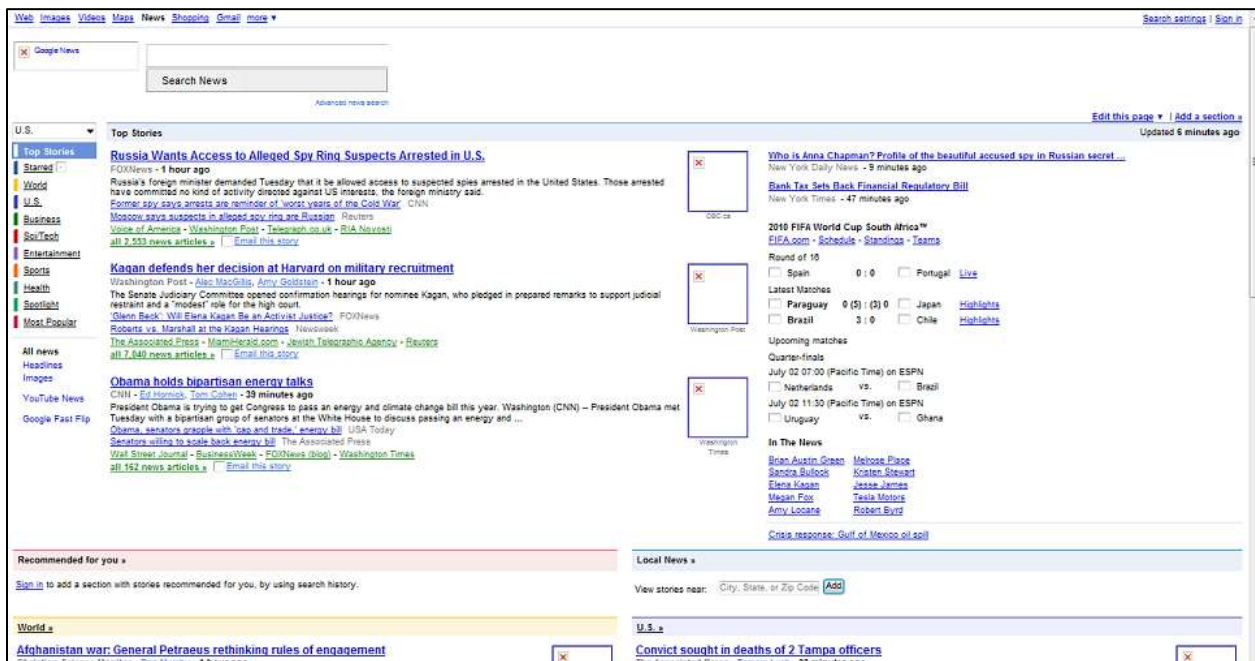
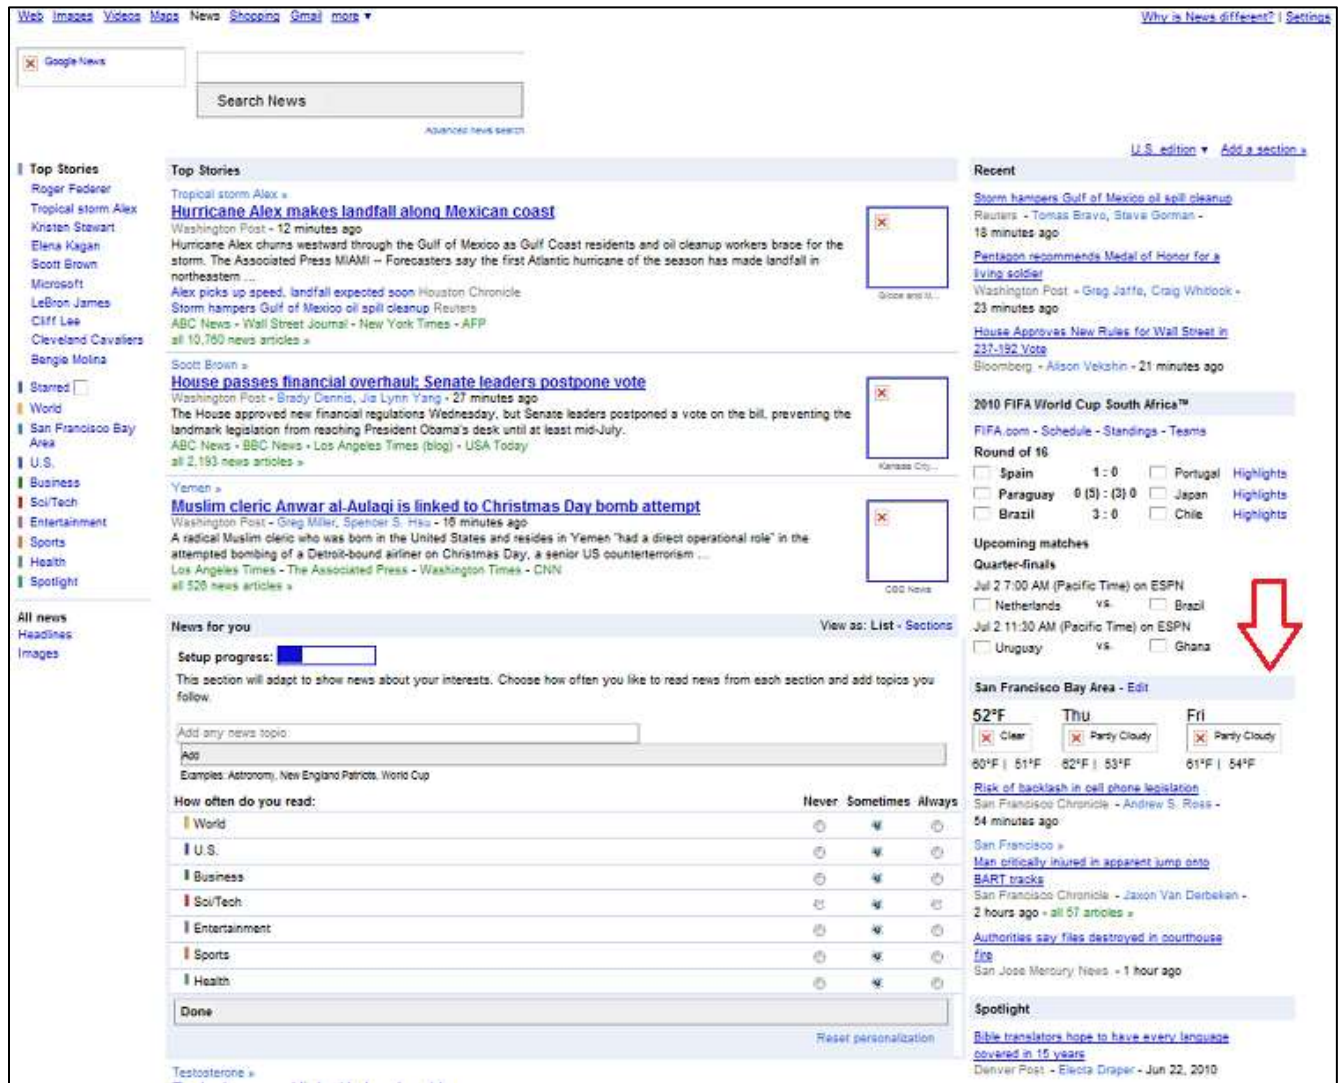


Figure 2: Google News Page After Redesign, July 2, 2010



The empirical strategy has features in common with Athey and Mobius (2012), who study the effect of a Google News design change on local news consumption in France. A key difference here is examination of a redesign that placed geo-targeted news links in front of users simultaneously rather than as an “opt-in” feature. In addition to allowing a simpler estimation strategy, the results offer a more complete measure of the effects of lower transaction costs on local news consumption and a closer link to privacy policies associated with geo-targeting. A second key difference is insight on the larger and more complex US news market, which requires demand-based measures of local news and local content.

3. Data

The basic working data is a panel of 24,859 household news visits each day from April 1 through September 30, 2010. For each household each day, the data record the total number of news visits and the number of news visits local to the household. The data also record the

number of news visits and local news visits referred by Google and Yahoo (news and search). For each household, the data include the Metropolitan Statistical Areas (MSA) of residence and basic demographic indicators for race, Hispanic origin, and income status.

The working data set is constructed from several underlying sources, discussed in turn.

3.1. Site Visit Data

Site visits logs come from ComScore Inc.'s MediaMetrics series extracted through the Wharton Research Data Service (WRDS) interface. The data include complete browsing history for a sample of approximately 50,000 households that have opted to allow tracking of internet use. Data is collected at the machine level, so may reflect behavior of more than one user and does not capture computer use on mobile devices or computers outside the home. The opt-in nature of the collection program means that the sample may not be representative of the population in terms of geography or use characteristics. These limitations are common to most sources of computer use microdata and largely affect the interpretation or generalization of results rather than estimation.

The raw data include basic user demographics and household zip code. Zipcodes are merged with US Census geography to identify the MSA for each household. Households living outside of MSA's are not included in the study.

3.2. News Outlets

News sites are identified from several proprietary databases and public sources. Newspaper, radio, magazine and television outlets are identified primarily from the *Burrelle's Media Directory* (2000 & 2005 edition), Bulldog Reporter's *MediaPro Directory* (2008) and the Newspaper Association of America web site (2010). A list of major blogs, aggregators and other digital intermediaries was scraped from *Technorati*, a blog reference site. With a few exceptions, all sources linked from Google News are classified as news sites.¹

News visits are identified in the raw session data by merging the news site URL's with the raw session data by domain name. Infrequent news readers, defined as users visiting fewer than 10 news sites in the first half of 2010, are excluded from the study. The final working data include visits to 3,184 domains.

The raw site visit data available for this study include only top-level domains. This limitation means that media outlets co-hosted on the same domain cannot be distinguished. (For example local television and radio stations sometimes share a web site, as do some broadsheet and tabloid daily newspapers.) Since co-hosted sites are local to the same market and household visits are aggregated over domains each day for the analysis, this limitation does not affect research results. A more difficult problem is distinguishing referrals from Google News from

¹ Google News occasionally posts links to other aggregators or portal sites. These sites are not classified as news sources.

Google search. The procedure for indirectly measuring Google News use is described in the subsection on intermediation, below.

3.3. Local Visits

In US markets, there is no consensus measure of what constitutes “local” media, especially on the internet. For this research, local news visits are identified based on demand. For each domain, the number of visits from each MSA over the entire year are counted. The MSA with the highest number of visits is defined as the home MSA for each media outlet. The share of visits to each domain from the home MSA is recorded as the home share for each domain. To reduce measurement error, outlets with fewer than 10 visits over the year are excluded from the sample. The final working sample includes 3,184 news domains. Once a home market has been assigned for each outlet, a household visit is recorded as local if the household MSA matches the home MSA for the domain.

With this approach, every media outlet is assigned a home MSA. Some threshold is needed to identify sites with a national focus, or for which the appeal in any particular MSA is not much different than the appeal in other MSA’s. In the analyses, outlets with less than 10% of visits from the home MSA are classified as non-local and visits to these sites are never considered local to any user. The 621 non-local outlets constitute about 20% of the news domains and encompass most national news outlets, topical news sources, and international media. The cut-off, though somewhat arbitrary, was chosen to preserve the *New York Times* (11% New York share) and the *New York Post* (10% New York share) as local sites for New York residents but at the same time exclude the majority of national news outlets such as *CNN* (3.9% in Atlanta), *USA Today* (4.2% in Las Vegas) and *Fox News* (5.4% in San Diego). With this criteria, the *Wall Street Journal* (7.1% New York) is not considered local to New York. Regression results are robust to alternate specifications, but for this study the inclusive sample is most useful for capturing changes in behavior.

For purposes of this study, the revealed preference measure of local interest has the feature of capturing implicit localism in different types of outlets which would not be identified with an exogenously defined media list characterized by the place of publication. For future research, the measure offers a useful way of characterizing competition between broadcast and print media, and also for understanding the extent to which demand for different types of information by different types of users is satisfied locally. For example, minority-targeted news outlets vary considerably in the share of viewing each receives in different markets. Relating demand for non-local media to individual and population demographics offers a systematic way of studying how groups with distinct tastes satisfy demand for information. For this study the relevant geography is set at the MSA level, but the data allow for finer definitions of local media for future research on community information needs.²

² There is considerable policy interest in this topic. For example, in 2012 the FCC commissioned a literature review examining how communities meet “critical information needs.” Though the authors concluded that digital media were not likely to satisfy those needs, they highlighted the subjective nature of current research and need for generalizable measures of local supply and demand.

The share of visits from households in the home MSA provides a measure of localism for each news outlet. The 20 domains with the largest number of local visits in the sample are shown in table 1 along with visit counts and the local share. In most markets the media outlet with the most local visits is the major urban daily newspaper in the market, but in some markets the top outlet is a local radio or television station website. For example, WXPI News in Pittsburgh receives more visits from inside the MSA than the Pittsburgh Post-Gazette, although the newspaper receives more visits overall (4,985). In general, the local visit share for radio and television stations is considerably higher for broadcast than for print media, with many broadcast sites reaching shares of over 90%. In table 1, the only sites with a local share exceeding 90% are broadcast sites.

Table 1: Top 20 Local News Sites, Visits & Local Share

News Site	URL	Local Visits	Total Visits	Local Share	Home MSA
Atlanta Journal Constitution	ajc.com	8,227	12,087	68%	Atlanta, GA
New York Times	nytimes.com	8,206	72,419	11%	New York, NY
KSL News Radio	ksl.com	7,739	10,845	71%	Salt Lake City, UT
News Channel 9	9wsyr.com	6,751	6,920	98%	Syracuse, NY
Gannett Arizona Newspapers	azcentral.com	5,577	9,052	62%	Phoenix, AZ
Washington Post	washingtonpost.com	5,549	21,963	25%	Washington, DC
NY Daily News	nydailynews.com	5,315	23,473	23%	New York, NY
WRAL News	wral.com	5,223	7,275	72%	Raleigh-Durham, NC
Boston Globe Portal	boston.com	5,131	14,448	36%	Boston, MA
LA Times	latimes.com	5,179	23,470	22%	Los Angeles, CA
Houston Chronicle	chron.com	5,041	8,247	61%	Houston, TX
The Post-Standard	syracuse.com	4,493	5,937	76%	Syracuse, NY
Philadelphia Inquirer, Daily News	philly.com	4,077	7,033	58%	Philadelphia, PA
Cincinnati Enquirer, Community Press	cincinnati.com	4,021	4,696	86%	Cincinnati, OH
Chicago Tribune	chicagotribune.com	3,895	8,680	45%	Chicago, IL
Cleveland Plain Dealer/Sun News	cleveland.com	3,705	7,389	50%	Cleveland, OH
WPXI News	wpxi.com	3,511	3,834	92%	Pittsburgh, PA
Minneapolis Star Tribune	startribune.com	3,362	4,642	72%	Minneapolis, MN
Newark Star Ledger	nj.com	3,199	3,340	96%	Newark, NJ
WYFF4 TV	wyff4.com	3,408	15,509	22%	Greenville, SC
New Orleans Times Picayune	nola.com	3,183	6,758	47%	New Orleans, LA

3.4. Intermediation

News visits directed by intermediaries are identified from a referral field in the raw session data that lists the referring domain. Most relevant for this study are referrals from Google and

Yahoo.³ The raw session data identifies only top-level domains for referrals as well as visits. Because of this, referrals by Google News cannot be directly distinguished in the raw data from Google search referrals. The basic identification strategy, which relies on changes to Google News that do not affect search, does not require distinguishing referrals in the data. However, construction of a treatment group of users most affected by the redesign of the Google News page does require identifying the most active Google News users. I develop an indirect measure of Google News usage by linking referrals in the visit data to outlets appearing on Google News. Specifically, Google News headlines are scraped from the archival site *Archive.org* that operates a program called the “Wayback Machine.” The scraped data identify for each domain-day whether or not an outlet appeared on Google News. In the working data, news visits referred by Google on days the domain was listed on the Google News page are classified as Google News referrals. For example, a visit to the *Atlanta Journal Constitution* referred by Google on a day the Constitution appeared on Google News would be coded as a Google News referral. A visit to the *Atlanta Journal Constitution* referred by Google on a day the newspaper did not appear on Google News would not be recorded as a Google News referral. A visit to the *Constitution* on a day it appeared on Google News but was not referred by Google not be identified as a Google News referral. The share of all news visits referred by Google News in the first six months of 2010 offers a measure of intensity in Google News use that distinguishes users most likely to be affected by the redesign.

The indirect measure of Google News use is limited in several ways. Most important is that because the measure is based on referrals, it only captures behavior of users who “click through” to media outlets. It is not possible to fully separate users who never visit Google News from those who visit but do not follow links. I return to this point in interpreting results.

There is also some noise in measuring referrals. During the sample period, 1-3 snapshots are captured each day. If the Google News page is updated more often, I will miscount Google News referrals in the visit data. Second, the web scrape does not pick up all user customization available before the redesign, which can lead to under-counting referrals. Finally, the procedure attributes all Google referrals to domains listed on Google News as a Google News referral, though it might be the case that users visited different news stories than those posted on Google News. With domain-level attribution, the procedure would over count Google News referrals. In general, since the primary purpose of counting Google News referrals is to identify more active and less active Google News users, the noise introduced by the indirect measure does not undermine the basic empirical approach. However because the redesign changed the number of personalized links in ways not captured by the scrape, the number of Google News referrals cannot be used directly as an independent variable in the empirical work.

³ The referral field can be supplemented by recording the domain visited prior to each news visit and classifying news visits following a Yahoo or Google visit as a referred news visit, but for this final report only the referral field is used.

Table 2: Sample Statistics (Media Outlets)

	Variable	N	Mean	SD	Min	Max
<i>Year Totals</i>						
All News Visits	Visits	4,885	738	4,572	10	205,590
Local News Visits	Lvisits	4,885	170	467	0	8,227
Local Visit Share	Lshare	4,885	0.434	0.31	0	1
Google News Visit Share (Google News Referrals/Visits)*	GNVshare	4,885	0.044	0.13	0	1
Conditional Local Google News Referral Share (Local Google News Referrals/All Google News Referrals)	CGNLshare	892	0.234	0.28	0	1
All Google Referral Share (Google Referrals/Visits)	RGshare	4,885	0.317	0.23	0	1
Conditional Local Google Referral Share (Local Google Referrals/All Google Referrals)	CRGLshare	892	0.294	0.32	0	1
Likelihood of a Google News Link	Gmax	4,885	0.210	0.41	0	1
<i>Daily</i>						
All News Visits	Visits	1,311,510	2.749	15.33	0	974
Local News Visits	Lvisits	1,311,510	0.633	1.79	0	64
Local Visit Share	Lshare	602,222	0.429	0.44	0	1
Google News Visit Share (Google News Referrals/Visits)*	GNVshare	262,346	0.082	0.22	0	1
Conditional Local Google News Referral Share (Local Google News Referrals/All Google News Referrals)	CGNLshare	46,660	0.210	0.37	0	1
All Google Referral Share (Google Referrals/Visits)	RGshare	602,222	0.270	0.38	0	1
Conditional Local Google Referral Share (Local Google Referrals/All Google Referrals)	CRGLshare	272,133	0.236	0.40	0	1
Likelihood of a Google News Link	Gmax	1,311,510	0.101	0.30	0	1

* Calculated before redesign, January- June 2010

3.5. Sample Statistics

Tables 2 and 3 reports summary statistics for domains and households. The top portion of table 2 presents yearly totals, the lower half summarizes daily data. The average number of visits to each news outlet is 738 over the year, with a range from 10 to 205,590. (Recall that outlets with fewer than 10 visits per day are dropped from the sample.) Total local visits average 170, with an average local share of 43%. The average share of visits referred by Google News in the first half of the year prior to the redesign is 4.4%. The share of Google News referrals from local users is 23.4%. The share of news outlets that appear at least once on Google News is 21%, or 1027 out of the 4,885 sites. The share of referrals from Google (news and search) is 31.7%, with 29.4% local.

The daily data show similar patterns, with a local share visit share averaging 42% and Google News referral share of 8.2%. Local visitors comprise 21% of Google News referrals and 23.6% of all Google referrals. The average number of visits to each domain each day is low, 2.7 visitors. The average likelihood a site appears on Google News each day is 0.10.

Table 3 reports analogous measures for the sample of households. Recall that households with fewer than 10 visits in the first half of the year are excluded from the sample as infrequent news readers. Households outside of MSA's are also excluded. For the remaining sample of 24,139 households, the average total number of news visits over the year is 140, with a substantial range. Households make an average of 33 local visits over the year, for a local share of 18%. The local share of household visits is much lower than the share of visits to outlets from local users, 43.4%, reflecting that many households frequently visit a small number of national outlets such as CNN. Households visit an average of 40.8 different news outlets per month, but only three different local domains. The estimated share of news visits in the first half of the year referred by Google News is 19.5%, with 9.8% of these to local outlets. The total share of visits referred by Google (news and search) is 30%, with a 12.5 % local share. About 86% of households have at least one Google News referral over the year.

Again, the daily data in the lower half of the table reflect the annual totals. The daily local visit share is 22%. The daily share of news visits referred by Google News is 19%, 11% of which are local. The total share of news visits intermediated by Google is 23% with a 14% local share. About 3% of households are estimated to have at least one Google News referral each day.

Before turning to the empirical analysis, it is useful to plot the trends in local news consumption in the raw data. Figure 3 shows a weekly average of the number of local news visits each day for households with at least one news visit. The bottom line shows local visits for households with more intense Google News use (referral share exceeding 20% in the first half of the year) and the upper line shows visits for households with no estimated Google News referrals prior to the redesign. (Recall that Google News intensity is estimated from referrals, so Google News users that never follow links to sites are considered low-intensity users.) Figure 4 repeats the graph for the share of news visits to local outlets. Weekly average shares are plotted over time, with the June 30 redesign date marked with a vertical line. For intensive Google

News users, local visits and visit shares increase at the point of the redesign and remain elevated through the study period. For less intense users, local consumption is more variable and increases slightly after the redesign before dropping off. There is no clear shift in the local share for low intensity users.

The next section outlines the estimation strategy for testing the patterns shown in the graphs and other measures of local consumption.

Table 3: Sample Statistics (Households)

	Variable	N	Mean	SD	Min	Max
<i>Year Totals</i>						
All News Visits	Visits	24,139	140	289	10	7,098
Local News Visits	Lvisits	24,139	33	121	0	2,848
Local Visit Share	Lshare	24,139	0.180	0.22	0	1
Google News Visit Share (Google News Referrals/Visits)*	GNVshare	24,139	0.195	0.18	0	1
Conditional Local Google News Referral Share (Local Google News Referrals/All Google News Referrals)*	CGNLshare	19,894	0.098	0.20	0	1
All Google Referral Share (Google Referrals/Visits)	RGshare	24,139	0.308	0.26	0	1
Conditional Local Google Referral Share (Local Google Referrals/All Google Referrals)	CRGLshare	21,541	0.125	0.19	0	1
Likelihood of Google News Link	Gmax	24,139	0.863	0.34	0	1
Unique News Outlets (Month)	udomains_m	24,139	40.817	61.81	1	1,747
Unique Local News Outlets (Month)	uldomains_m	24,139	2.921	5.67	0	170
<i>Daily Averages April 1-September 30, 2010</i>						
All News Visits	Visits	3,682,827	0.454	1.493	0	164
Local News Visits	Lvisits	3,682,827	0.106	0.577	0	48
Likelihood of Local News Visit	lvisit	3,682,827	0.057	0.232	0	1
Local Visit Share	Lshare	708,308	0.222	0.379	0	1
Google News Visit Share (Google News Referrals/Visits)*	GNVshare	1,876,831	0.190	0.181	0	1
Conditional Local Google News Referral Share (Local Google News Referrals/All Google News Referrals)*	CGNLshare	79,526	0.106	0.296	0	1
All Google Referral Share (Google Referrals/Visits)	RGshare	708,308	0.232	0.396	0	1
Conditional Local Google Referral Share (Local Google Referrals/All Google Referrals)	CRGLshare	204,398	0.144	0.332	0	1
Likelihood of Google News Link	Gmax	3,682,827	0.031	0.173	0	1
Unique News Outlets (Day)	udomains_day	3,682,827	0.350	1.009	0	127
Unique Local News Outlets (Day)	uldomains_day	3,682,827	0.073	0.334	0	13

* Calculated before June 30 redesign.

Figure 3: Local News Visits, April 1-September 30, 2010

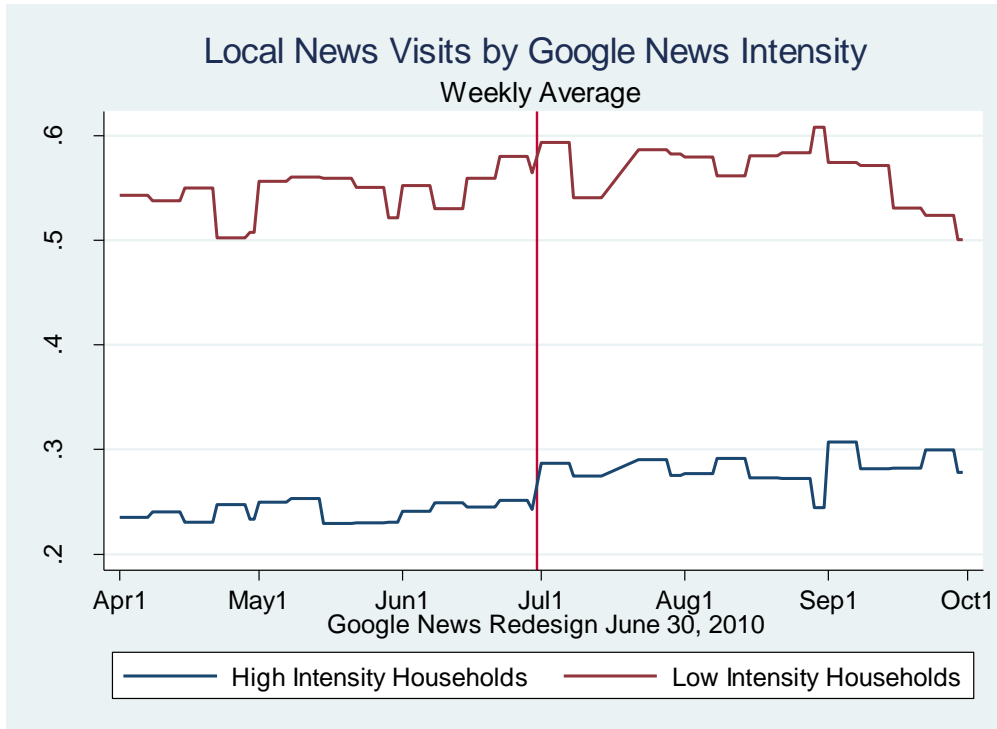
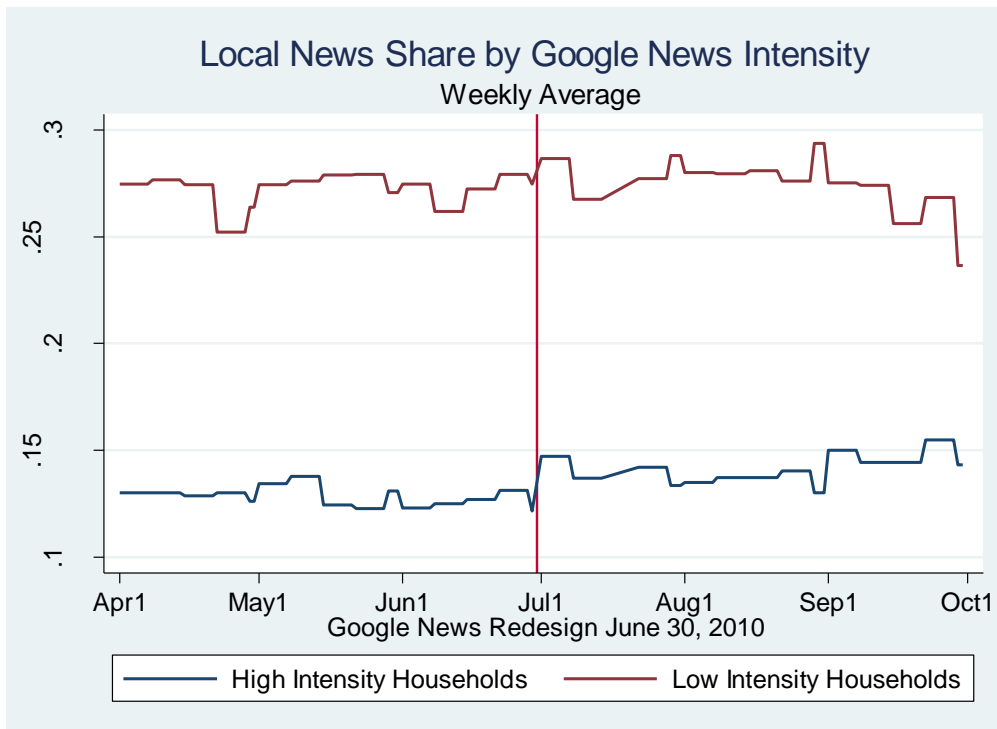


Figure 4: Local News Visit Share, April 1-September 30, 2010



4. Empirical Methodology

The goal of the analysis is to identify the effect of adding local news links to the Google News page on local news consumption. The baseline specification for estimation is a standard fixed effects formulation:

$$(1) \quad Y_{it} = \beta_0 + \beta_1 Post + \beta_2 PostX + \tau + \gamma_i + \varepsilon_{it}$$

where the independent variable Y captures a local consumption measure for each household i each day t , $Post$ is a treatment dummy set to 1 after local news was added to the Google News site, and X is a treatment measure for internet users more strongly affected by changes to the Google News site. The error terms ε_{it} are assumed to be independent across users. A time trend τ and household fixed effect γ are included in all specifications, as are dummy variables for days of the week (not shown). Equation (1) is estimated with four measures of local news consumption (Y) and two treatment specifications (X). The four independent variables are the number of local news visits (with log transformation, $lnLvisits$), the probability of a local news visit ($ILvisit$), the share of news visits to local sites ($Lshare$), and the share of all news visits referred by Google that are local ($CRGLshare$). The first two measures address the amount local news consumption, the second two measures concern the share of attention to local relative to non-local news.

The choice of modeling approach for count data is not straightforward, especially when much of the variation is captured in the difference between zero (no local news visit) and one. In the tables that follow, I use a semi-log specification and a linear probability model to capture changes in local news consumption. Because of the large number of users who make no local news visits on any given day, local news visits are transformed as $lnLvisits = \ln(Lvisits + 1)$. With this specification, the coefficients can be interpreted as the percent change in the local visits after the Google redesign. A second and more interpretable approach for capturing changes in the quantity of local news visits is to use a linear probability model, where the coefficient estimates show the effect of the redesign on the likelihood a treatment or control household makes a local news visit.⁴ In evaluating results, I focus on this measure.

The third and fourth metrics consider the share of attention to local relative to non-local media. The third specification measures the share of all news visits households make to local sites. The fourth and final measure considers the share of all Google referrals that are local. Because the set of links on the Google News page varies by household after the redesign in ways that are not captured by the scraped archival pages, the local share is calculated from the universe of Google referrals not from the estimated Google News referrals. All variables are measured at the household-day level.

Two identification strategies are adopted for each set of local news consumption measures. The first uses a treatment and group and control group. The treatment group is comprised of

⁴ I also estimate negative binomial fixed effects models on visit counts, which offer a better fit for data with the overdispersion exhibited here (Cameron, A. C. and Pravin K. Trivedi (1998). See Section 5 for other robustness checks.

frequent Google News users in the first half of 2010. The control group consists of households with no estimated Google News referrals that are frequent Yahoo users in the first half of 2010. Adding local links to Google News should affect the treatment group more than the control group.⁵ In terms of equation (1), X is set equal to one for the Google News users and to zero for the control group of Yahoo users. A positive coefficient β_2 supports the hypothesis that the effect of the redesign is greater for the Google News users.

The ideal treatment group would be chosen to resemble the control group in every way but for the use of Google News. As in Chiou Tucker (2011), Yahoo users are chosen as a comparison group because they are expected to be more similar to Google News users than individuals who do not use intermediaries. However in practice, the number of regular Yahoo users that do not also use Google is small, only about 1,000 households. The presence of top-level domains also means that Yahoo users are identified from both referrals from both news and search. To make better use of available data, I adopt a second strategy that identifies the effect of adding local links to the Google News site from the intensity of Google News use prior to the redesign. Intensity is measured as the share of all news visits referred from Google News ($GNVshare$) prior to the redesign, and in these specifications replaces the treatment dummy as X in equation (1). A positive coefficient β_2 then indicates the effect of adding local links to Google News for more avid Google News users.

The treatment group is comprised of households with more than 20% of news visits referred by Google News in the first half of the year and no referrals from Yahoo. The control group is comprised of users with more than 20% of news visits referred by Yahoo prior to the redesign and no Google News referrals. Results are not highly sensitive to the cutoff, but the cutoffs do affect the sample size and the magnitude of measured effects.

Sample statistics for the treatment group and control group are shown in table 4. Looking first at visit counts, total news consumption by the two groups are similar, with an average of about 0.25 news visits per day. The local share is considerably higher for Yahoo users, (29% vs 13% of all news visits). The share of Hispanic and black users in the two groups is similar, with slightly higher minority shares in the Yahoo group. The share of high income users is higher for the Google News group (21% vs. 10%). Google users visit a greater number of unique outlets per month (40 vs. 30), but fewer local outlets (4 vs. 5.6).

With regard to intermediation, by construction the share of visits referred by Google News is zero for the control group. The average share of news visits referred by Google News in the treatment group is 39%. The share of all Google referrals for the control group is very small but non-zero, as some of the Yahoo users also use Google search. (Constructing the control group with households that never use Google search would cut the sample size by one third, significantly reducing the power of the analysis.)

⁵ Because Google News intensity is measured from referrals, individuals who use Google News but do not click through to individual web pages are recorded as infrequent users and can be included in the control group. For this reason the effect of the redesign on the control group may not be zero but should be less than effects on the control.

Table 4: Sample Statistics (Treatment & Control Households)

	Variable	HH	N	Mean	SD	Min	Max
<i>Google News Users (Treatment)</i>							
All News Visits	Visits	5,663	865,449	0.25	1.01	0.00	65.00
Local News Visits	Lvisits	5,663	865,449	0.03	0.28	0.00	20.00
Likelihood of Local News Visit	ILvisit	5,663	865,449	0.02	0.15	0.00	1.00
Local Visit Share	Lshare	5,663	115,537	0.13	0.32	0.00	1.00
Google News Visit Share*	GNVshare	5,663	865,449	0.39	0.14	0.20	0.97
Conditional Local Google News Ref. Share*	CGNLshare	5,663	40,210	0.12	0.31	0.00	1.00
All Google Referral Share	RGshare	5,663	115,537	0.53	0.47	0.00	1.00
Conditional Local Google Referral Share*	CRGLshare	5,656	69,140	0.12	0.31	0.00	1.00
Likelihood of Google News Link	Gmax	5,663	865,449	0.05	0.21	0.00	1.00
Unique News Outlets (Day)	udomains_d	5,663	865,449	0.22	0.78	0.00	47.00
Unique Local News Outlets (Day)	udomains_d	5,663	865,449	0.03	0.21	0.00	10.00
Unique News Outlets (Month)	udomains_m	5,663	717,591	39.88	55.93	1.00	940.00
Unique Local News Outlets (Month)	ulddomains_m	5,663	252,072	4.04	7.83	1.00	108.00
Hispanic %	hhispanic	5,663	865,449	0.08	0.27	0.00	1.00
Black %	hblack	5,663	865,449	0.12	0.33	0.00	1.00
High Income %	hinc	5,663	865,449	0.21	0.40	0.00	1.00
<i>Yahoo Users (Control)</i>							
All News Visits	Visits	1,131	175,437	0.24	0.77	0.00	43.00
Local News Visits	Lvisits	1,131	175,437	0.07	0.36	0.00	16.00
Likelihood of Local News Visit	ILvisit	1,131	175,437	0.05	0.22	0.00	1.00
Local Visit Share	Lshare	1,131	24,943	0.29	0.43	0.00	1.00
Google News Visit Share*	GNVshare	1,131	175,437	0.00	0.00	0.00	-
Conditional Local Google News Ref. Share*	CGNLshare	-	-	0	0	0.00	1.00
All Google Referral Share	RGshare	1,131	24,943	0.03	0.16	0.00	1.00
Conditional Local Google Referral Share*	CRGLshare	539	865	0.19	0.38	0.00	1.00
Likelihood of Google News Link	Gmax	1,131	175,437	0.00	0.03	0.00	1.00
Unique News Outlets (Day)	udomains_d	1,131	175,437	0.21	0.62	0.00	12.00
Unique Local News Outlets (Day)	udomains_d	1,131	175,437	0.06	0.28	0.00	7.00
Unique News Outlets (Month)	udomains_m	1,131	147,053	29.51	31.30	1.00	203.00
Unique Local News Outlets (Month)	ulddomains_m	1,131	73,079	5.63	7.77	1.00	53.00
Hispanic %	hhispanic	1,131	175,437	0.07	0.25	0.00	1.00
Black %	hblack	1,131	175,437	0.13	0.34	0.00	1.00
High Income %	hinc	1,131	175,437	0.10	0.30	0.00	1.00

* Calculated before redesign, January - June 2010

The time period of analysis is the three months before and after the June 30th redesign when local links were added to the Google News site, from April 1, 2010 through September 30, 2010. I eliminate from the sample June 30, when local links first appeared on the Google News page but had not yet been announced, and also July 21, 2010 when Google News was unavailable for part of the day. While it is possible to estimate equation (1) over a longer time period, changes in local and non-local news consumption patterns into the fall season introduce considerable noise into the data and complicate causal links to the Google News redesign. Incremental change to Google search algorithms in early 2010 and in the fall further limit the ability to study

longer time periods with specifications akin to equation (1). However the clear and uniform date for the redesign does allow for a somewhat more extended study than was possible in Athey and Mobius (2012), where the opt-in nature of the French redesign limited inference beyond a two week study window.

A key assumption in the error structure of (1) is that the timing of the Google News redesign is uncorrelated with unobserved trends in local news consumption. The effect of Google News is also restricted to impact the mean of local consumption at a constant level. At the end of section 5 I present a set of results with a more flexible time structure to consider the robustness of the basic results and whether the effects of the redesign persist over time.

Table 5: Do Geo-Targeted News Links Increase Local Visits Among Google News Users?

	Log Local Visits (1)	Prob. Local Visit (2)	Local Visit Share (3)	Local Referral Share (4)
Post Treatment	-.0011 (-1.43)	-.0013 (-1.46)	.0008 (.17)	-.1048** (-3.37)
Post x Google	.0021** (3.02)	.0020* (2.50)	.0048 (1.14)	.1096** (3.54)
Time Trend	-.0002 (-.41)	-.0003 (-.59)	.0014 (.51)	.0050 (1.36)
M	.0034** (7.14)	.0044** (7.97)	-.0001 (-.02)	-.0012 (-.30)
T	.0033** (6.92)	.0040** (7.19)	-.0044 (-1.54)	-.0047 (-1.18)
W	.0035** (7.19)	.0041** (7.27)	-.0045 (-1.54)	-.0018 (-.44)
Th	.0038** (7.88)	.0043** (7.88)	-.0055+ (-1.92)	-.0048 (-1.21)
F	.0028** (5.78)	.0035** (6.37)	.0028 (.96)	-.0025 (-.62)
Sat	.0006 (1.26)	.0007 (1.33)	.0026 (.87)	.0098* (2.41)
Constant	.0211** (49.45)	.0252** (51.46)	.1606** (61.95)	.1172** (33.09)
Adj. R-Squared	.20	.16	.35	.26
N	1040886	1040886	140480	70005

Dependent variable in column 1 is transformed log of local news visits. Dependent variable in column 2 is probability of a local visit. Dependent variable in column 3 is share of news visits to local outlets. Dependent variable in column 4 is share of google referrals to local outlets. All specifications include household fixed effects. See text for details. T-statistics in parentheses: + p 0.1 * p 0.05, ** p 0.01.

5. Results

5.1. Basic Specifications

Estimates of equation (1) for the treatment and control group are shown in table 5. In the first row, *Post* (β_1) reports the effect of adding local links to Google News on the control group while the interaction *Post x Google* (β_2) reports the result for Google News users. The first two

columns measure the volume of local news visits. After the Google redesign, Yahoo user make fewer local visits and Google users make more. Significance tests show the difference in effects across the two groups to be positive and statistically significant, indicating that local visits increased after the redesign among Google News user relative to Yahoo users. The magnitude of the coefficient estimates is modest, indicating an increase in local news visits by only 0.1% for the treatment group. The probability of a local news visit for the treatment group increased somewhat more, by 0.0007 on an average of 0.023 or 3%.

Results in columns 3 and 4 show a similar pattern, with the redesign showing a negative effect on the Yahoo users and a positive effect for Google News users. A test of the difference between the effect of the redesign on Google relative to Yahoo users shows the effect to be positive but below standard significance levels. The magnitudes imply that the redesign increased the local share for Google News users by 4.2%. For the share of local referrals in column 4, the difference is statistically significant. The magnitudes of the coefficients indicate that adding local links to the Google News page increased the local share of Google referrals by .0048 on an average of 0.12, or 4%, which is consistent with the local share results.

Table 6: Does the Effect of Geo-Targeted News Links Increase with Google News Use?

	Log Local Visits (1)	Prob. Local Visit (2)	Local Visit Share (3)	Local Referral Share (4)
Post Treatment	-.0040** (-8.93)	-.0043** (-9.33)	.0037** (2.59)	.0017 (.49)
Post x Google News Share	.0149** (13.52)	.0142** (12.45)	.0136** (3.48)	.0126+ (1.71)
Time Trend	-.0017** (-4.97)	-.0016** (-4.55)	.0010 (.92)	.0014 (.63)
M	.0090** (24.84)	.0091** (24.09)	-.0004 (-.30)	-.0001 (-.05)
T	.0100** (27.59)	.0100** (26.66)	-.0018 (-1.51)	-.0024 (-1.01)
W	.0105** (28.37)	.0101** (26.42)	-.0034** (-2.72)	-.0007 (-.30)
Th	.0104** (28.78)	.0101** (27.16)	-.0026* (-2.16)	.0011 (.46)
F	.0074** (20.31)	.0077** (20.35)	.0008 (.67)	-.0012 (-.48)
Sat	.0007+ (1.81)	.0007+ (1.78)	.0029* (2.31)	.0077** (3.13)
Constant	.0500** (155.32)	.0527** (158.59)	.2193** (199.84)	.1401** (64.84)
Adj. R-Squared	.40	.31	.52	.33
N	3,646,360	3,646,360	699,963	201,669

Dependent variable in column 1 is transformed log of local news visits. Dependent variable in column 2 is probability of a local visit. Dependent variable in column 3 is share of news visits to local outlets. Dependent variable in column 4 is share of Google referrals to local outlets. All specifications include household fixed effects. See text for details. T-statistics in parentheses: + p 0.1 * p 0.05, ** p 0.01.

Turning to the second set of estimates, table 6 reports results for the full sample of households where X is the share of visits referred from Google News prior to the redesign. The post treatment indicator β_1 shows the baseline effect of the redesign for households with no Google News referrals. The interaction term *Post x Google News Share* (β_2) shows the incremental effect for more intensive use. The daily average Google News share before the redesign from table 3 is 0.19 with a standard deviation of 0.18.

As in table 5, the top row shows effects of the redesign on households with no Google News referrals through June 2010. The second row shows the interaction of the post period and the share of all news visits referred by Google News prior to the change. The first two columns consider the effect of adding local links on the level of local news consumption. Both transformed local visits and the likelihood of a visit are estimated at about -0.004 with small standard errors, indicating baseline effects close to zero. For a household with mean (0.19) Google News referral share, the effect of the redesign is slightly negative, reducing local visits by 0.1%. For a household one standard deviation above the mean, the redesign increases local news visits by 0.2%. For a user at the 95th percentile (53% Google referral share), the redesign increases local visits by 0.4%. These are small changes from a low baseline of 0.11 local visits per day.

Effects with the linear probability model are larger. At the mean, the effect of the redesign is negative, $(0.0142 \times 0.19 - 0.0043 = 0.0068)$ with a 2.8% drop in the likelihood of a local visit. For households one standard deviation above the mean in Google News referrals, the likelihood of a local news visit increases by 0.001 after the treatment, about 1.7%. For users at the 95th percentile in Google News intensity, the likelihood of a local visit increases by 0.003 or about 5.7% after the redesign.

Results for shares in columns 3 and 4 are larger still. The baseline effect is positive as well as the intensity interaction. Households with mean Google News referrals see the local share of visits rise by 0.006, or 2.9%. The share of local visits increases by about 4% for users one standard deviation above the mean, and by 5% for households at the 95th percentile of Google News intensity.

A few comments on the results are warranted. Visit shares are defined only on days with positive news visits (or positive Google referrals for column (4)), so the share samples are not directly comparable to the visit or linear probability samples. Restricting the visit and linear probability models to households making at least one news visit increases the magnitude of the results considerably. For the visit specification in column (1), local visits increase for users with mean Google News referrals by 0.6%, and by 0.9% and 1.1% at one standard deviation above the mean and at the 95th percentile, respectively. For the linear probability model, the likelihood of a local news visit increases by 4.4% after the redesign for households with mean Google News referrals, and by 6.2% and 8.4% for households one standard deviation above the mean and at the 95th percentile. Full specifications are included in appendix table A1.

Separately, it is not surprising that the measured effect of the redesign is higher for the linear probability models, since much of the variation in consumption is captured in the difference

between zero and one local visit. However the fact that the local shares change more than local visits suggests that the redesign might induce households to shift news consumption from outside to local outlets. An estimate of equation (1) for non-local visits and the probability of a non-local visit suggest that this is indeed the case, with non-local visits declining after the redesign more strongly for more intense Google News users. Full specifications are reported in the appendix table A1.

Taken together, results in tables 5 and 6 indicate that the Google News redesign did increase local news consumption overall and also shift attention from non-local to local news sources, with the largest effects among most intense users. The magnitude of the increase is small, with the likelihood of a local link increasing by about 4-6% and total local visits by 0.2-0.9% for a household one standard deviation above the mean in Google News referrals prior to the redesign. Given the low level of local visits by Google News users, adding geo-targeted links to the Google News page is not likely to have an economically meaningful impact on local news outlets.

5.2. Aggregators and Consumption Variety

In addition to net increases in local news consumption, another important avenue by which aggregators impact the market is by diffusing or concentrating attention across outlets. Theoretically, lower transaction costs can trigger either effect, but under some plausible specifications of demand aggregation increases consumption variety more than improvements in search.⁶ To test this, equation (1) can be estimated with the dependent variable defined as the unique number of local news outlets visited per day or per month.

Table 7 reports the relationship between the Google redesign and the number of unique local news sites visited per day and per month. Columns (1) and (3) report results for the treatment and control sample, while columns (2) and (4) report results for the full sample and Google News use intensity. Both specifications indicate an underlying decrease in the number of different local outlets visited per day for households that do not use Google News prior to the redesign and an increase for active Google News users. An increase in Google News use intensity of one standard deviation (0.18) increases the number of unique sources visited per day by 0.3%, a small but positive effect. An increase from the 5th to the 95th percentile increases the number of unique sources visited per day by 0.6%. None of the results on unique sites visited per month are statistically different than zero, suggesting that increases in variety come about by more frequent visits to multiple sites rather than overall introduction of new outlets into the household set.

The lack of overall impact of geo-targeted links on local consumption variety suggest that the process of discovery and matching often attributed to aggregators does not play an important role in local news markets, where households are already familiar with the media landscape. The results also suggest that geo-targeted links on Google News are taken largely from

⁶ George and Hogendorn 2012 consider the differential impact of aggregation and search on demand, with implications for targeted advertising.

traditional outlets rather than niche sources. The relationship between aggregators and consumption variety along different news dimensions is a topic worth of further study.

Table 7: Do Local Links Increase Consumption Variety?

	Local Outlets per Day		Local Outlets per Month	
	(1)	(2)	(3)	(4)
Post Treatment	--.0012+	--.0030**	.0307	--.0147
	(--1.69)	(--8.19)	(.60)	(--.73)
Post x Google	.0018**		.0058	
	(2.88)		(.13)	
Post x Google News Share		.0108**		.0379
		(11.73)		(.69)
Time Trend	--.0002	--.0014**	--.0003	--.0018
	(--.52)	(--5.00)	(--.01)	(--.12)
M	.0032**	.0074**		
	(7.39)	(24.44)		
T	.0031**	.0082**		
	(7.12)	(27.18)		
W	.0031**	.0084**		
	(6.92)	(27.25)		
Th	.0035**	.0085**		
	(8.10)	(28.16)		
F	.0027**	.0062**		
	(6.13)	(20.53)		
Sat	.0006	.0006+		
	(1.27)	(1.95)		
Constant	.0194**	.0416**	.6963**	.9323**
	(50.41)	(155.29)	(42.05)	(115.69)
Adj. R-Squared	.18	.34	.41	.50
N	1040886	3646360	10848	46859

Dependent variable in columns 1-2 is unique local news outlets visited per day (log transform). Dependent variable in columns 3-4 is unique local news outlets visited per month (log transform). All specifications include household fixed effects. See text for details. T-statistics in parentheses: + p 0.1 * p 0.05, ** p 0.01.

5.3. Robustness

To investigate more fully the likely persistence of effects from local news links and to evaluate the robustness of results, equation (1) can be modified to include separate time trends for the pre- and post-redesign period, and to interact each of the trends with the Google News intensity measure. Equation (1) becomes:

$$(2) \quad Y_{it} = \beta_0 + \beta_1 \text{Post} + \beta_2 \text{PostX} + \beta_3 \tau + \beta_4 \tau \text{Post} + \beta_5 \tau X + \beta_6 \tau \text{PostX} + \gamma_i + \varepsilon_{it}$$

With this specification, β_3 and β_4 show the baseline time trend before and after the redesign, and β_5 and β_6 show how the time trend before and after the redesign vary with Google News

intensity. As in equation (1), β_2 shows how the effect of the redesign varies with Google News intensity.

Table 8: Local News Consumption and Google News Intensity Over Time

	Log Local Visits	Prob. Local Visit	Local Visit Share	Local Referral Share
	(1)	(2)	(3)	(4)
Post Treatment	.0136** (3.27)	.0102* (2.40)	.0036 (1.05)	.0006 (.08)
Post x Google News Share	.0643** (6.53)	.0604** (5.99)	.0576** (6.97)	.0722** (4.76)
Time Trend	.0023 (1.02)	.0003 (.14)	.0008 (.39)	.0107* (2.03)
Post x T	-.0073* (-2.28)	-.0030 (-.91)	.0018 (.68)	.0079 (1.42)
T x Google News Share	.0134 (1.54)	.0170+ (1.90)	.0135+ (1.83)	-.0222 (-1.62)
Post x T x Google News Share	-.0545** (-27.61)	-.0578** (-28.56)	-.0573** (-34.53)	-.0527** (-22.92)
M	-.0002 (-.12)	-.0004 (-.23)	-.0008 (-.59)	-.0003 (-.11)
T	-.0005 (-.27)	-.0012 (-.66)	-.0013 (-.89)	-.0033 (-1.12)
W	-.0014 (-.81)	-.0041* (-2.31)	-.0045** (-3.12)	-.0008 (-.27)
Th	.0003 (.17)	-.0015 (-.88)	-.0030* (-2.19)	.0013 (.49)
F	-.0045** (-2.64)	-.0030+ (-1.75)	-.0013 (-.90)	-.0032 (-1.12)
Sat	-.0010 (-.56)	-.0008 (-.43)	.0006 (.38)	.0046 (1.55)
Constant	.2781** (174.55)	.2930** (179.33)	.2181** (162.95)	.1385** (52.34)
Adj. R-Squared	.59	.51	.52	.33
N	545959	545959	545959	157581

Dependent variable in column 1 is transformed log of local news visits. Dependent variable in column 2 is probability of a local visit. Dependent variable in column 3 is share of news visits to local outlets. Dependent variable in column 4 is share of Google referrals to local outlets. All specifications include household fixed effects. See text for details. T-statistics in parentheses: + p 0.1 * p 0.05, ** p 0.01.

Estimates of equation (2) are shown in table 8. As with earlier results, the first line shows the effect of the redesign at the baseline and the second shows how this effect changes with Google News intensity. The baseline measures are zero or positive, and the intensity interactions (β_2) are all positive and statistically significant, confirming the results in table 6 that the redesign increases local visits and local shares.

The time interactions show little evidence of a pre-existing baseline trend but a positive trend for more intensive Google users. This differential trend before the redesign may be due to

field testing of the redesign, which was reported periodically in technology blogs in months leading up to the launch. The baseline trend in visits drops off after the redesign in the levels specification, but overall the results do not offer much indication of a baseline trend before or after the redesign. In all specifications the trend for Google News intensity switches from positive to negative after the redesign, indicating that the redesign effect erodes over time. The time trend is measured in units of 90 days, which indicates that for a fixed user intensity the positive impact on local news consumption of the redesign would not persist much beyond that time.

6. Conclusion

This research studies the effect of adding geo-targeted local news links to Google News in US online news markets. Results indicate that adding local news links increases both the amount and share of local news consumption, and that the size of the effect is higher for more active Google News users. However the magnitude of the effects are small, with increases in the likelihood of a local news visit increasing by 4-6% and total local news visits by slightly less than 1% for households one standard deviation above the mean in Google News referrals prior to the redesign. Access to local links increases the variety of local sites visited per day, but not per month, suggesting that increases in local news consumption arise from more frequent visits to familiar news outlets rather than visits to new news providers. Given the low level of local visits by Google News users, adding geo-targeted links to the Google News page is not likely to have an economically meaningful impact on local news outlets.

References

- Athey, Susan and Markus Mobius. *The Impact of News Aggregators on Internet News Consumption: The Case of Localization*. Working paper, February 2012.
- Cameron, A. C. and Pravin K. Trivedi (1998), *Regression Analysis of Count Data*, Econometric Society Monograph No.30, Cambridge University Press, 1998.
- Chiou, Lesley and Catherine Tucker. *Copyright, Digitization, and Aggregation*. NET Institute Working Paper No. 11-18, December 2011.
- Chrysanthos Dellarocas, Zsolt Katona, and William Rand. *Media, Aggregators And The Link Economy: Strategic Hyperlink Formation In Content Networks*. Working Paper, February 2012.
- George, Lisa M. and Christiaan Hogendorn (2012). "Aggregators, Search and the Economics of New Media Institutions." *Information Economics and Policy*, Vol. 22: Issue 1 pp. 40-51.
- George, Lisa M. (2008) "The Internet and the Market for Daily Newspapers." *The B.E. Journal of Economic Analysis & Policy*, Vol. 8, No. 1 (Advances), Article 26.
- George, Lisa M. and Joel Waldfogel (2006). "The *New York Times* and the Market for Local Newspapers." *American Economic Review*, Vol. 96, No. 1, pp. 435-477.

George, Lisa M. and Joel Waldfogel (2008). “National Media and Local Political Participation: The Case of the *New York Times*” in Roumeen Islam, ed., *Information and Public Choice: From Media Markets to Policymaking*. Washington, DC: World Bank Publications, pp. 33-48.

Jeon, Doh-Shin and Nikrooz Nasr Esfahani. *News Aggregators and Competition Among Newspapers in the Internet*. Working paper, June 2012.

Appendix

Table A1: Robustness Checks: Conditional Visits and Non-Local Visits

	Log Local Visits (Visits>0) (1)	Prob. Local Visit (Visits>0) (2)	Log Non-Local Visits (3)	Prob. Non- Local Visit (4)
Post Treatment	.0033+ (1.93)	.0038* (2.17)	-.0027 (-1.18)	-.0046** (-2.99)
Post x Google News Share	.0247** (5.30)	.0196** (4.12)	-.0120+ (-1.90)	-.0044 (-1.05)
Time Trend	-.0014 (-1.08)	-.0004 (-.26)	-.0106** (-5.86)	-.0016 (-1.34)
M	.0006 (.39)	.0001 (.07)	-.0004 (-.21)	.0006 (.46)
T	-.0007 (-.47)	-.0014 (-.93)	.0033+ (1.67)	.0017 (1.28)
W	-.0002 (-.11)	-.0029+ (-1.93)	.0071** (3.53)	.0045** (3.37)
Th	.0004 (.28)	-.0013 (-.85)	.0040* (2.03)	.0036** (2.76)
F	-.0021 (-1.43)	-.0007 (-.44)	-.0087** (-4.38)	-.0024+ (-1.77)
Sat	.0019 (1.25)	.0018 (1.21)	-.0072** (-3.57)	-.0031* (-2.27)
Constant	.2825** (216.39)	.2956** (221.18)	.8644** (488.17)	.8422** (711.56)
Adj. R-Squared	.58	.51	.43	.40
N	699,963	699,963	699,963	699,963

Dependent variable in column 1 is transformed log of local news visits conditional on at least one household news visit. Dependent variable in column 2 is probability of a local visit conditional on at least one household news visit. Dependent variable in column 3 is transformed log of non-local news visits. Dependent variable in column 4 is probability of a non-local visit. All specifications include household fixed effects. T-statistics in parentheses: + p 0.1 * p 0.05, ** p 0.01.

Columns (1) and (2) in table A1, above, replicate the results in columns (1) and (2) in table 6 restricting the sample to households making at least one news visit that day. Columns 3 and 4 repeat the analysis in columns (1) and (2) in table 6 studying non-local visits.