

Attention Allocation in Information-Rich Environments: The Case of News Aggregators

Chrysanthos Dellarocas, Questrom School of Business, Boston University, Boston, MA 02215,
dell@bu.edu

Juliana Sutanto, Management Science Department, Lancaster University, Bailrigg, Lancaster, Lancashire
LA1 4YX, United Kingdom, j.sutanto@lancaster.ac.uk

Mihai Calin, The Boston Consulting Group AG (Switzerland), Munstergasse 2, Zurich 8001, Switzerland,
cm.calin@gmail.com

Elia Palme, Newscron SA, Via Carlo Maderno 24, Lugano 6900, Switzerland, elia.palme@newscron.com

Abstract

News aggregators have emerged as an important component of digital content ecosystems, attracting traffic by hosting curated collections of links to third party content, but also inciting conflict with content producers. Aggregators provide titles and short summaries (snippets) of articles they link to. Content producers claim that their presence deprives them of traffic that would otherwise flow to their sites. In light of this controversy, we conduct a series of field experiments whose objective is to provide insight with respect to how readers allocate their attention between a news aggregator and the original articles it links to. Our experiments are based on manipulating elements of the user interface of a Swiss mobile news aggregator. We examine how key design parameters, such as the length of the text snippet that an aggregator displays about articles, the presence of associated images, and the number of related articles on the same story, affect a reader's propensity to visit the content producer's site and read the full article. Our findings suggest the presence of a substitution relationship between the amount of information that aggregators offer about articles and the probability that readers will opt to read the full articles at the content producer sites. Interestingly, however, when several related article outlines compete for user attention, a longer snippet and the inclusion of an image increase the probability that an article will be chosen over its competitors.

Keywords: digital content, media curation, media economics, news aggregator, click-through rate

1. Introduction

The massive amount of news content available online has increased the importance of curation and aggregation, that is, of interfaces and services that help readers filter and make sense of the subset of content that is important to them. Historically such functions used to be the realm of professional editors.

Editors not only commissioned the production of content but also decided what content would be included in a newspaper and how it would be organized.

Web technologies allow this important function to be unbundled from content production. Specifically, the web's ability to place hyperlinks across content has enabled new types of players, commonly referred to as content aggregators, to successfully enter professional content ecosystems, attracting traffic and revenue by hosting collections of links to the content of others (Dellarocas et al. 2013; Dewan et al. 2004). Content aggregators produce little or no original content; they usually provide titles and excerpts (hereafter called snippets) of the articles they link to (Figure 1). Examples of well-known aggregators include Google News, the Drudge Report, and the Huffington Post. Google News (news.google.com) is a search engine of many of the world's news sources; it algorithmically aggregates headlines and groups similar articles together. The Drudge Report (www.drudgereport.com) aggregates selected hyperlinks to news websites all over the world; each link carries a headline written by the site's editors. The Huffington Post (www.huffingtonpost.com) is a hybrid of news aggregator and original content creator.



Figure 1: Example of a news aggregator article entry

Facing severe financial pressures, some content creators have turned against content aggregators, accusing them of stealing their revenues by free riding on their content.¹ Media tycoon Rupert Murdoch has been particularly outspoken on this issue, referring to aggregators as “parasites” and selectively blocking some from indexing the content of media sites he owns.² Since 2012, several countries, primarily in Europe, experimented with the imposition of a tax on news aggregators³ or passed laws that mandated payment from aggregators to publishers for the reproduction of news snippets.⁴ Other market actors point out that, in today's link economy, links bring valuable additional traffic to their target nodes.

¹ The 2009 dispute between the Associated Press and News Corporation with Google is a representative example. See <http://www.forbes.com/2009/04/06/google-ap-newspapers-business-media-copyright.html>

² <http://www.mediaite.com/online/rupert-murdoch-begins-blocking-new-aggregators-search-engines/>

³ See “Newspapers versus Google: Taxing times”, The Economist magazine, November 10, 2012; “French Publishers Forge Deal With Google, Breaking Ranks With Europe”, The New York Times, February 17, 2013.

⁴ See “Google News says ‘adiós’ to Spain in row over publishing fees”, The Guardian, December 16, 2014.

Therefore content creators should be happy that aggregators exist and direct consumers to their sites (Jarvis 2008; Karp 2007). Key aggregator executives, such as Google’s Eric Schmidt, assert that it is to their interest to see content creators thrive, since the value of links (and aggregators) is directly related to the quality of content that these point to.⁵

A central aspect of the debate focuses on the complex economic implications of the process of placing (for the most part) free hyperlinks across content nodes. The main argument in favor of aggregators is that, if links are chosen well, then they point to good quality content; as a result, they reduce the search costs of the consumers, which may lead to more traffic for higher quality sites. The main argument against aggregators is that some consumers satisfy their curiosity by reading an aggregator’s short snippet of a linked-to article and never click through to the article itself. In fact, the question of whether aggregators are legally permitted to reproduce an article’s title and snippet without obtaining permission from (and possibly paying) the content producer, is still unresolved.⁶

The question of whether the current generation of news aggregators is beneficial or harmful to content ecosystems remains open (Athey and Mobius 2012; Chiou and Tucker 2011). Nevertheless, we believe that the ever-increasing volume of available content makes some form of aggregation an inevitable, and valuable, component of every content ecosystem. The key question, therefore, is not whether aggregators should exist, but rather how the, partly symbiotic and partly competitive, relationship between aggregators and content creators can be optimized for the benefit of both parties.

To provide insights to these questions, we examine the distribution of readers’ attention between a news aggregator and the original articles it links to. The focus of our interest is a user’s decision to follow the provided link towards the content producer’s site and read the full text of an aggregated article. Our objective is to understand how key aggregator design parameters, such as the length of the text snippet that an aggregator provides about an article, the presence of associated images, and the presence of other related articles on the same topic, affect a reader’s propensity to click on an article. To achieve our objective, we conduct a set of field experiments with smartphone and tablet versions of a Swiss news aggregator application.

⁵ “CEO Eric Schmidt wishes he could rescue newspapers”, *Fortune* January 7, 2009.

⁶ Aggregators claim that the reproduction of titles and short snippets of text falls under the “fair use” provisions of copyright law. However, as stated by Isbell (2010), “for all of the attention that news aggregators have received, no case in the United States has yet definitively addressed the question of whether their activities are legal.”

We find evidence for the presence of a substitution relationship between the amount of information that aggregators display about an article and the probability that readers will opt to read its full text at the content producer sites. Our results suggest that an article's headline provides all the information users need to decide if an article is close enough to their interests. Any additional information provided by aggregators, in the form of text snippets or images, apparently satiates the appetite of some readers and can only serve to decrease click-through rates. Interestingly, however, when several related articles compete for user attention, a longer snippet and the inclusion of an image increases the probability that an article will be chosen over its competitors.

Besides contributing to research, the findings of this study are valuable for aggregators seeking to optimize their traffic patterns, as well as in terms of informing the public discourse between aggregators and content creators on the need for equitable business agreements between the two parties.

2. Related Work

The relationship between news aggregators and content producers is the subject of a small, but growing, body of scholarly work. Dellarocas et al. (2013) model how the ability to place costless hyperlinks to third party content affects the behavior of competing content producers, who can now choose between spending effort to write an original article on a story and simply linking to an article that someone else has written. They view aggregators as a limiting case of content nodes who are inefficient in original content production and, therefore, can only attract readers by placing links to interesting third-party content. The paper shows that the impact of an aggregator on the content ecosystems is the sum of two opposite effects. On one hand, a search cost reduction effect arises from the fact that aggregators generally place links to well-chosen content and provide some information (snippet) about this content that helps users decide whether it matches their interests. This effect is positive; it increases the overall consumption of content in the entire ecosystem and primarily benefits high quality content producers. On the other hand, a free riding effect is due to readers who browse aggregator headlines and snippets, and never click through to the original articles. The free riding effect is at the core of the controversy between aggregators and original content producers. It reduces the content producers' profits and incentives to produce quality content.

Chiou and Tucker (2011) offer an empirical contribution to the discourse about the net impact of aggregators. They empirically examine the effect of the removal of all hosted articles by The Associated Press from Google News at the end of 2009 (due to a dispute in licensing negotiations) on what sites consumers visited. They find that the removal of The Associated Press's content was correlated with a

decline in subsequent visits to traditional news sites (immediately after visiting Google News) as compared to other news aggregators that continued to host The Associated Press content. The results suggest the presence of a complementary relationship between aggregators and content producers, whereby article summaries hosted by aggregators induce readers to seek more news on those stories after visiting the aggregator.

In another empirical paper with a similar objective, Athey and Mobius (2012) look at how the addition of a localization feature on Google News affects the consumption of local content. They find that the addition of this feature increases local news consumption, including the number of direct visits to such sites (that presumably users discover via Google News and then begin to visit directly). However, the effect diminishes over time.

Hong (2011) focuses on the potential for aggregators to induce information cascades that concentrate traffic to a few “popular” sites. The author provides evidence of an association between the number of visitors to a news aggregator site and the online traffic concentration of that site. The author suggests design interventions for alleviating the adverse impact of such phenomena.

Our work also relates to the broader discourse on how readers allocate their attention in content networks. For example, Wu and Huberman (2008) analyze the role that popularity and novelty play in attracting the attention of users to dynamic websites. Agarwal et al. (2009) propose novel spatial-temporal models to estimate click-through rates in the context of content recommendation. Roos et al. (2011) propose a model of browsing behavior in hyperlinked media that takes into consideration a user’s utility and beliefs about the quality of cross-linked content.

Compared to this broader literature, our aims are more focused, looking specifically on how consumers allocate attention between news aggregators and news articles and how design parameters of the aggregator affect this allocation.

3. Research Hypotheses

Aggregators have emerged as an increasingly important layer in hyperlinked content ecosystems. Dellarocas et al. (2013) model aggregators as content network nodes who attract readers by placing links to interesting third-party content. They distinguish between the impact of aggregators on individual pieces of content and on competition among the related pieces of content. Accordingly, we develop hypotheses on both questions.

3.1. Impact of Snippet Length and Accompanying Images on Click-Through Rates

Text snippets (headlines plus article text excerpts) act as ‘free samples’ or ‘previews’ of news articles. News articles belong to the class of goods that economists refer to as ‘experience goods.’ Consumers must, thus, experience them in order to value them (Shapiro and Varian 1998). Experiencing digital content (e.g. reading a news article) requires substantial time and cognitive effort, which is wasted if the content does not end up satisfying the user’s needs. Offering free samples or previews is a way for publishers to allow consumers to have actual experience with the good’s quality and fit before purchase (Greene et al. 2000). For example, in the case of books, retailers, such as Amazon.com provide free previews of the contents of some titles (typically one or two chapters) for customers to browse online and evaluate.

The marketing literature has paid some attention to the study of free samples in the context of conventional products, such as durables and packaged goods (for example, Bawa and Shoemaker 2004; Heiman et al. 2001). Researchers have identified a tension between two opposite effects of distributing free samples: (1) the immediate cannibalizing impact of free samples – some consumers receiving free samples would otherwise be purchasing the product and, (2) the market expansion impact due to increased awareness and (hopefully) positive experiences with the sample that lead to repeat purchases and word of mouth (Bawa and Shoemaker 2004). In such settings, the quality of the sample is assumed to be the same as that of the actual product sold in the market; marketers try to optimize revenue by controlling the quantity of free samples that are distributed to the consumer population (Heiman et al. 2001).

In the context of digital goods, repeat purchase is usually not an option, since every piece of content is unique. Therefore, instead of limiting quantity, digital sampling strategies usually make available a reduced quality, or excerpted, free version of the good to all prospective consumers. For example, Amazon provides only a chapter or two of the online preview of a title, limiting the utility that consumers can derive relative to the full book. Furthermore, the image quality of these chapters is much lower than the corresponding print or Kindle version (Kannan 2013). In the case of news articles, relevant parameters of samples include the length of the snippet and the inclusion of an image from the complete article.

Halbheer et al. (2014) develop an analytical model of the problem of selecting the optimal length (more broadly: optimal information content) of a digital sample. Their model is driven by the assumption that longer samples have two simultaneous effects, roughly corresponding to the ‘cannibalization’ and ‘market expansion’ effects of Bawa and Shoemaker (2004): (1) they decrease demand for the full content, because the value of the residual content (i.e. content not contained in the free sample) is lower, and (2) they may

increase demand because consumers have a more precise idea about content quality and fit with their tastes and thus feel less reluctant to consume the full content.

If we further assume the following two limiting cases: (1) in the absence of any information about a news article (i.e. no headline and no text excerpt), no consumer would access it (e.g. because the consumers' perceived probability of a random article's fit with their taste is very low), and (2) in the presence of a snippet that contains the full article, no consumer will click through to the original article (because the residual information content of the full article is zero), Halbheer et al. (2014)'s model implies the following:

Hypothesis 1a: *News aggregator article click-through rates exhibit an inverse U-shaped relationship with article snippet length*

Besides text snippets, news article previews frequently display selected images from the article. Images reinforce text information and make it more concrete (David 1998). In the context of conventional products (i.e., consumer durables and packaged goods products), the presence of product images has been consistently found to increase sales (Di et al. 2014).

Unlike conventional products, where the utility from consumption is not affected by the prior viewing of a product image (Bawa and Shoemaker 2004, Heiman et al. 2001), in the case of news article previews, images are a part of the article itself; displaying them at the aggregator level thus reduces the value of the residual content that users can access by clicking-through to the full article.

Inclusion of images in news article previews is, thus, expected to have similar tradeoffs to those of increasing the length of text snippets: on the one hand it might increase demand because it conveys what the article is about more concretely, on the other hand it decreases demand because it gives away part of the story. Given that, in our setting, images are displayed in addition to headlines and, possibly, text snippets, we formulate the following hypothesis:

Hypothesis 1b: *If the amount of information about an article provided by its headline and text snippet is small, inclusion of an image increases the article click-through rate; if the amount of information about an article provided by its headline and text snippet is large, inclusion of an image reduces the article click-through rate.*

3.2. Impact of the Presence of Multiple Snippets on Group Click-Through Rates

Popular stories typically have multiple newspaper articles written about them. News aggregators collect such articles together and display their snippets next to each other. This is an interesting aspect of

aggregator behavior, whose implications have not yet received sufficient attention. We denote group click-through rate the probability that a user will click at least one article from among a group of related articles.

An important question is whether the presence of multiple related article snippets increases or decreases a user's propensity to click on any of these articles. Both possibilities are plausible. On the one hand, one can argue that, the more article snippets are available on a story, the more likely that a user will find at least one of them appealing. Furthermore, the presence of multiple related articles tends to signal important stories that are worth reading about. On the other hand, the combined presence of multiple article snippets potentially provides more information about the story itself. This reduces the residual utility of the original articles (Halbheer et al. 2014) and may satisfy the curiosity of readers who will then decide it is not necessary to read the original articles.

Besides this possible information complementarity effect, multiple articles on the same topic might lead to a choice overload effect. According to the choice overload theory, an increase in the number of available options may lead to adverse consequences, such as a decrease in the motivation to choose (Iyengar and Lepper 2000; Kuksov and Villas-Boas 2010; Scheibehenne et al. 2010). If too many alternatives are offered, then users have to spend more time and exert more cognitive resources to make a choice (Kahn and Lehmann 1991), which may deter them from making a choice altogether (Kuksov and Villas-Boas 2010). If too few alternatives are offered, then users may not choose, fearing that an acceptable choice is unlikely (Kuksov and Villas-Boas 2010).

Both theories (information complementarity and choice overload) seem to suggest the existence of an optimal number of related articles that maximizes user's propensity to click on any of the articles. This discussion leads to the following hypothesis:

Hypothesis 2: *The number of available articles about the same story exhibits an inverse U-shaped relationship with group click-through rates*

3.3. Impact of Snippet Length and Accompanying Images on an Article's Choice Probability

For stories that have multiple articles competing for user attention, another important question is what factors make users choose among the competing articles. It is well documented that position matters a lot. The higher the article is on the list, the higher the probability that it will be chosen (see for example,

Ghose and Yang 2009). What has not been researched is the impact of an article's snippet length and presence of accompanying image on the choice probability (controlling for position). Taking a perspective of utility maximization, Xiang and Soberman (2011) argued that an article preview confers a positive externality on a competitor that displays a significantly shorter preview. The reason is that when there are competing article previews or snippets, rational readers will choose to click-through a snippet that maximizes the utility of the remaining content of the article. This argument assumes that readers pay equal attention to all competing snippets, process all available information and act as rational utility maximizers. Previous empirical studies suggest that this may not be the case.

Information foraging theory (Pirolli and Card 1999) posits that user behavior in an information-rich environment is guided by *information scent*, which in turn, is determined from the user's perception of the value and cost of distal information sources, as obtained from proximal cues (Chi et al. 2001; Pirolli and Card 1999). Article snippets act as proximal cues that are used by information foragers to assess the distal content, i.e. the content of full articles at the other end of the Web link. Cutrell and Guan (2007) examined the competing snippets of search engine results and found that the length of snippets may act as a *spotlight*, such that, as the snippet length gets longer, the relative gaze fixation or attention to the snippet increases. They also found that, as the snippet length gets longer, the accuracy and hence, satisfaction of the choice made by search engine users improved.

We hypothesize that news aggregator users behave similarly to search engine users. The above findings then imply that when there are several snippets competing to be chosen, users are likely to pay more attention to the longer snippets and perceive the distal content represented by the longer snippets to be more valuable. This leads to the following hypothesis:

Hypothesis 3a: *When several articles about the same story compete for user attention, controlling for position, readers are, on average, more likely to click on articles whose snippet lengths are longer*

In an information rich environment, images have been known to attract the attention of web users (Sargent 2007; Zillman et al. 2001) and inspire interest in related texts (Knobloch et al. 2003). We, thus, expect that if users behave as information foragers, the presence of accompanying images will act as a spotlight that will attract more user attention. This leads to the following hypothesis:

Hypothesis 3b: *When several articles about the same story compete for user attention, controlling for position, readers are, on average, more likely to click on articles that have an accompanying image*

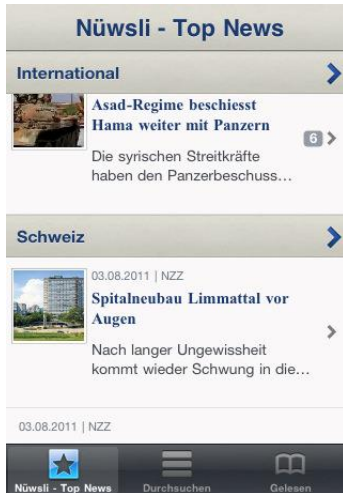
4. Field Experiment Setting

Our field experiments are conducted on a Swiss news aggregator application called Newscron. The front-end of the app consists of two separate client versions (for iPhones and iPads respectively) that can be freely downloaded from Apple’s App Store. The back-end of the app is a server that collects and organizes news articles. The server collects all news articles published online by every major newspaper in Switzerland (in all three national languages: German, French, and Italian) on a daily basis (see Appendix A.1). The server performs a semantic analysis of article texts to group them together into topics (stories). Topics are, further, classified as belonging to one of 9 categories: international, local, business, technology, entertainment, sports, life, motors, and culture. This leads to the following data structure: every article belongs to a topic; a topic contains one or more articles and is assigned to a category.

The iPhone and iPad client versions of Newscron provide distinct user interfaces with different features and different strengths and limitations vis-à-vis the research questions that motivate this work. We have, therefore, conducted separate experiments on each version of the app to obtain complementary insights. In the rest of the section we describe each client version, the experiments we conducted on it, and the properties of the resulting data sets.

4.1. iPhone Client And Experiment

User interaction with the iPhone version of Newscron is designed as a three step process (Figure 2). First, the user is presented with a list of topics (news stories), organized by category (Figure 2a). When the user clicks on a topic, she sees all articles related to the particular topic, sorted by their publication dates (i.e., the most recent articles are displayed first). Only an outline (headline, snippet, and - if available - image) of each article is displayed (Figure 2b). Snippets in Newscron are simply the first characters of each article. By clicking on the article’s dedicated and labeled button at the bottom of the article’s outline (the button is labeled “Ganzen Artikel Lesen” in Figure 2b), the user is directed to the newspaper’s website to read the full article (Figure 2c).



(a) First level: Topics



(b) Second level: Article outline



(c) Third level: Full article

Figure 2: Newsron iPhone user interface

To test our hypotheses, we manipulated the length of article snippets at the second level of the user interface (Figure 2b). The default snippet length used in our app is equal to 245 characters, which is the average number of characters of snippets at Google News. We reduced/increased this default snippet length in increments of 20%, which is twice the standard deviation of snippets in Google News. We, thus, defined six different snippet lengths ranging from -60% to +40% of the default length (see Figure 3). We chose -60% because it is the shortest length that is supported by the user interface and +40% because it is the longest snippet possible subject to copyright agreements we have with the news providers. During our experiment the snippet length that was displayed when user i accessed article j was randomized. This means that different users might encounter the same article with different snippet lengths and the same user may encounter different articles with different snippet lengths. Furthermore, different articles within the same topic group could be displayed with different snippet lengths.⁷

⁷ As can be seen in Figure 3, snippets of different lengths were displayed in a way that does not significantly change the user's experience and, thus, does not confound the experimental results. All snippet lengths pretty much take up an entire iPhone screen. Longer snippets do not crowd out the page and shorter snippets do not have "extra white space."

98 characters (-60%)



147 characters (-40%)



196 characters (-20%)



245 characters (default)



294 characters (+20%)



343 characters (+40%)



Figure 3: Preview of different snippet lengths

Our main variable of interest is the click-through rate, which is the probability that a user will click through to an article linked to through the aggregator and will proceed to read it in its entirety at the content producer's site. We are interested in measuring two types of click-through: individual and group. An individual click-through rate stands for the click-through rate of a single article and is defined as the ratio of the number of times users click the button at the bottom of an article's outline (Figure 2b) and move to reading the full article at the publisher's site (Figure 2c) over the number of times that the article's outline is displayed to the users. Popular stories typically have multiple newspaper articles written about them. Newsron collects such articles together under a topic and displays their outlines next to each other. We denote group click-through rate the probability that a user will click at least one article from among a group of related articles. In such cases we are, additionally, interested in understanding which article(s) users choose to read.

The field experiment lasted for two weeks in the Spring of 2012 during which we had 2,057 users interacting with the Newsron app, generating 32,986 article outline display events (Figure 2b). A typical user opens the application on average 1.9 times per day, accumulating 2.7 minutes of average daily usage. Users select on average 3.44 topics per day, containing around 1.21 articles per topic. The field experiment data set is organized in topics (first level, Figure 2a), each topic containing one or several articles (second level, Figure 2b). An article can belong to only one topic throughout the experiment. During the two week period of the experiment, each user opened 12.21 topics on average. Decision time is the elapsed time between the time an article’s outline is displayed on a user’s display and the time the user either clicks-through to the publisher’s page or goes back to the list of topics. On average, users clicked-through 52% of article outlines with an average decision time of 12.41 seconds. Conditional on clicking, the average full article reading time was 82.77 seconds. Table 1 summarizes the key parameters of the iPhone data set.

Measurement	Value
Total unique users	2,057
Total unique topics	3,454
Total unique articles	4,953
Total articles having an image	3,677
Total article outline display events (Fig. 2b)	32,986
Total article click-through events (Fig. 2c)	17,220
Average number of topic access per user	12.21
Average decision time (in seconds)	12.41
Average reading time (in seconds)	82.77

Table 1. iPhone data set parameters

4.2. iPad Client and Experiment

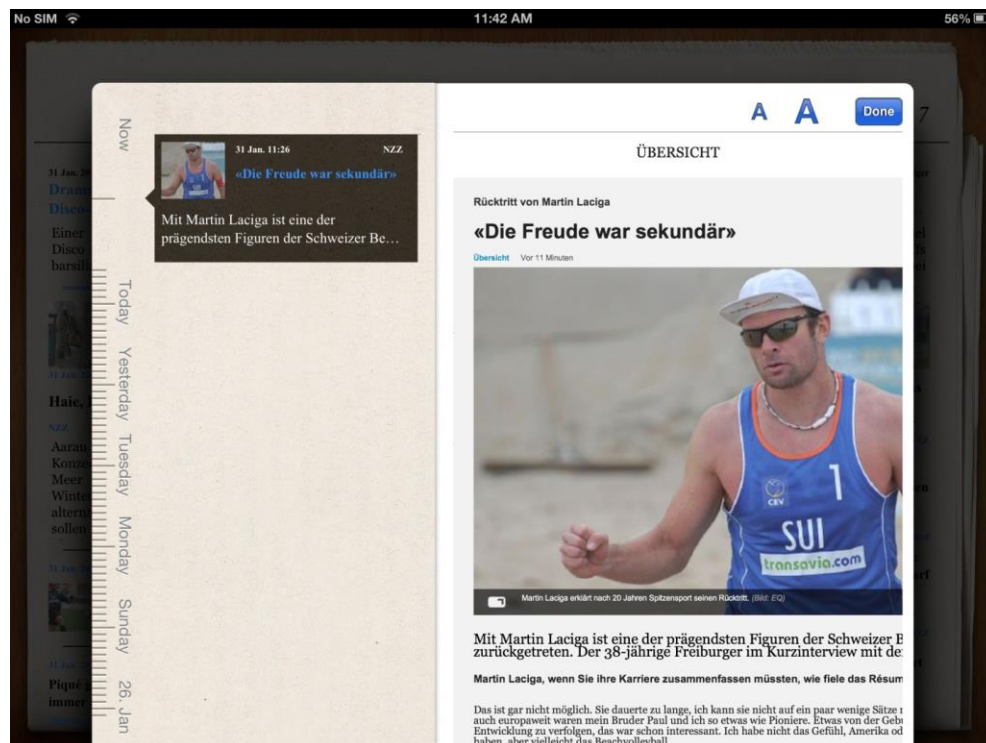
On the iPad, user interaction is designed as a two-stage process (Figure 4) that attempts to mimic the process of reading a traditional newspaper. The app’s entry page aims to mimic the front page of a traditional newspaper: the user is presented with the outline (i.e. headline, snippet and image) of a lead article at the center of the page. To the left of the lead article, a secondary article outline is displayed. Around these two article outlines, the app lists the headlines and images of 6-10 more articles (Figure 4a). Each section is displayed as a separate page, with a structure that is very similar to that of the front page. Upon clicking on one of the articles, a pop up window covers the screen displaying the publisher’s website with the full content of the article on the right, and any related articles on the left, on a timeline (see Figure 4b).

To test our hypotheses, we manipulated the lead article’s snippet length at the front page as well as at every category page (Figure 4a). During our experiment, the lead article snippet length that was displayed when user i accessed a page of the app was randomized. A random number between 1 and 7, drawn from a uniform distribution, was generated on each access. Depending on the value of this number, the snippet length shown to the user was either zero (i.e., no snippet was displayed, only the image and the title of the news) or one of the following lengths (in characters): 98, 147, 196, 245, 294 and 343, as reasoned in the iPhone experiment case. Compared to the iPhone experiment, where the smallest snippet length was 98 characters, the iPad experiment adds the possibility of articles without snippets.

The second manipulation investigates the effect of images on the click-through. The secondary article outline (Figure 1a) is manipulated to randomly display or hide its image. Before displaying secondary article i to user j , if that article had an available image, our system generated a random number between 1 and 7 drawn from a uniform distribution. If that number was between 1 and 3, the image was displayed, whereas if the number was between 4 and 7, it was not. The application then logs whether the article was displaying an image or not and whether it was clicked by the user. On the iPhone version this manipulation was not possible.



(a) First stage - Article outlines



(b) Second stage - Full article text and timeline of related articles

Figure 4: Newsron iPad user interface and manipulations

The iPad experiment ran for 16 weeks in 2012, during which 1,399 users interacted with the application, generating 65,906 topic display events. A topic display event represents displaying the first level page of a certain news category (see Figure 4a) which gathers data for both a snippet length manipulation (on the lead article) as well as an image display manipulation (on the secondary article). The average user launched the application 1.43 times a day, each time scrolling through 8.25 categories (and thus seeing 8.25 lead and 8.25 secondary articles), on which she clicked only 2.87 times (1.98 times on the lead article and 0.87 times on the secondary article) after spending 15.58 seconds deciding; average reading time for clicked articles was 65.48 seconds. Table 2 summarizes the key parameters of the iPad data set.

Measurement	Value
Total unique users	1,399
Total unique lead articles	15,920
Total unique secondary articles	13,613
Total ‘first stage’ display events (Fig. 4a)	65,906
Total lead article click-through events	2,783
Total secondary article click-through events	1,109
Average decision time (in seconds)	15.58
Average reading time (in seconds)	65.48

Table 2. iPad data set parameters

4.3. Why We Used Both Clients

Each of the two client apps allows us to investigate complementary aspects of user news reading behavior in the presence of aggregators. The iPhone app is the most mature of the two and has the largest user base. Of the two apps, only the iPhone app allows us to investigate how aggregating snippets of related articles affects user choice (Hypotheses 2 and 3). On the iPad app, when a user clicks on an story outline at the top level, even when there are multiple articles associated with the story, the app automatically displays the full text of the topmost (most recent, at the time of access) article. On the other hand, technical limitations on the iPhone app’s architecture do not allow us to reduce snippet lengths below 98 characters or to manipulate the presence of images.

The iPad app allows us to reduce snippet length to zero and to manipulate the presence/absence of an image associated with an article headline. It also offers a richer interface that is closer to that of a web browser, and can, therefore be used as a robustness check to make sure that the effects observed on the iPhone app are not due to idiosyncrasies or limitations of mobile interfaces.

Overall, performing similar experiments on two substantially different user interfaces and finding similar results increases our confidence that our findings represent fundamental aspects of online news consumption behavior.

5. Baseline Results

Before we proceed to analyzing the data, it is important to summarize here our variables of interest. Our independent variables are snippet-length, whether an image is displayed together with a snippet (has-image), and the number of related article outlines displayed about a story (related-snippets). In addition, we control for the following article characteristics: article language, article category, topic age (time elapsed between the publication of the earliest article on a topic and the timestamp of an access record) and time of day when an article was accessed (morning=5-8am, lunch=11am-1pm, afterwork=3pm-6pm, afterdinner=8pm-11pm). Most of our variables (except topic age) are categorical variables.

5.1. Impact of Snippet Length and Accompanying Images on Click-Through Rates

To analyze the impact of snippet length and images on click-through rates and filter out any side effects from other articles on the same story or topic, we restricted this analysis to sessions that only have one article (at the time they were accessed) in the iPhone data set. Some of these sessions correspond to topics for which additional related articles were later displayed, as the topic evolved. The data set used in this section contains 21,261 observations. Key descriptive statistics of this filtered data set are shown in Appendix A.2.

We performed model-free comparisons of the manipulated snippet lengths in two steps. First, we compared the click-through probability associated with different snippet lengths. As shown in Figure 5, click-through probabilities decrease as snippet lengths increase. All 95% confidence intervals, except those for lengths 294 and 343, are disjoint.⁸ Second, we applied a permutation test to all pairs of treatments to determine whether the population mean of any of the treatments differs from that of any other treatment. The results (p-values) are depicted in Table 3 and provide further evidence that all treatments, except lengths 294 and 343, result in statistically different average click-through probabilities, which, according to Figure 5, are monotonically decreasing.

⁸ All confidence intervals reported in this paper are 95% bias-corrected and accelerated (BCa) confidence intervals calculated using bootstrapping.




Figure 5: Click-through probabilities associated with different snippet lengths on iPhone data.

		Snippet length				
		147	196	245	294	343
Snippet length	98	0	0	0	0	0
	147		0.0002	0	0	0
	196			0.007	0	0
	245				0.001	0.0002
	294					0.687

Table 3: p-values of permutation test applied to all pairs of snippet lengths on iPhone data.

We performed a similar analysis to explore the impact on article click-throughs of having (not having) an accompanying image. Our results (Table 4) indicate that the presence of an image is associated with lower click-through rates. The corresponding 95% confidence intervals are disjoint. To obtain further insights, we performed logistic regression with user-level fixed effect to account for any systematic differences in the click-through rates of individual users.⁹ As mentioned above, in this regression, we also controlled for article characteristics: article language, article category, topic age, and time of day when an article was accessed. The regression result is shown in Table 5.¹⁰

⁹ We conducted a Hausman test to check if there is any significant difference between fixed-effect model and random-effect model and the result showed that there is no significant difference between the two models (p value of Hausman test : 0.3145). Based on this result, we decided to apply fixed-effect model instead of random-effect model. Every iPhone has a unique ‘device id’ that is recorded when it is accessing our app. We assume that each iPhone is used by a single individual user and include one fixed effect for each ‘device id’ present in our data.

¹⁰ Some of the predictors of our regressions are article-level predictors, such that they are the same for each impression of an article, and so there might be several observations in the data that share this value of the predictor in virtue of displaying the same article; likewise, they may share a common error component. This has the consequence of making inference for these quantities anticonservative. For that reason, throughout the paper we report cluster-robust standard errors, estimated using a sandwich method (Zeileis 2006).

Has image	Observations	Click-through probability	Lower CI	Upper CI
No	4,182	64.17%	62.70%	65.65%
Yes	17,079	60.78%	60.05%	61.51%

Table 4. Click-through probabilities associated with the presence (absence) of an image on iPhone data.

Consistent with our model-free analysis, regression analysis results in monotonically decreasing coefficients for click-through rates associated with increasing snippet lengths. Overall, H1a is not supported. The presence of an accompanying image results in a negative and significant coefficient. Using permutation testing, we compared the click-through probabilities associated with each snippet length without and with the presence of accompanying image (Table 6). The click-through probabilities are uniformly higher when no image is displayed; however they are only statistically different for some snippets (snippet-98, 147, and 245). Overall, we find partial support for H1b.

	Coeff.	Std. Error	Pr(> Z)
Main variables			
snippet-98 (baseline)			
snippet-147	-0.291	0.061	***
snippet-196	-0.559	0.059	***
snippet-245	-0.687	0.059	***
snippet-294	-0.879	0.059	***
snippet-343	-0.901	0.058	***
has-image	-0.302	0.043	***
Control variables			
topic-age	-0.013	0.004	**
category-int. (baseline)			
category-local	0.048	0.048	0.316
category-business	-0.193	0.064	**
category-technology	0.488	0.068	***
category-entertain	0.379	0.061	***
category-sports	0.224	0.081	**
category-life	0.531	0.086	***
category-motors	-0.122	0.136	0.369
category-culture	0.247	0.092	**
time-other (baseline)			
time-morning	0.115	0.060	0.055
time-lunch	0.128	0.062	*
time-afterwork	-0.025	0.063	0.694
time-afterdinner	-0.015	0.061	0.804
language-german (baseline)			
language-french	-0.317	0.269	0.238
language-italian	0.071	0.291	0.806
AIC		27126	

***Significant at < 0.001; **Significant at < 0.01; *Significant at < 0.05

Table 5: iPhone individual click-through rate regression (21,261 single article sessions).

Our combined results suggest that displaying more information (in terms of longer snippets or inclusion of an accompanying image) in article outlines results in lower click-through rates until snippet lengths reach 294 characters. Beyond that point, further increasing snippet length or displaying an image does not have a statistically significant effect on click-through rates.

Snippet lengths	Click-through probability	p-value of the Difference in Click-through probability
98 – no image	0.754	0.010
98 – with image	0.708	
147 – no image	0.701	0.013
147 – with image	0.656	
196 – no image	0.643	0.100
196 – with image	0.615	
245 – no image	0.634	0.003
245 – with image	0.578	
294 – no image	0.562	0.320
294 – with image	0.551	
343 – no image	0.569	0.119
343 – with image	0.544	

Table 6: Click-through probability of the manipulated snippet lengths (without and with image)

With regard to control variables, we found that: 1) click-through rates generally decline as news topics get older, which is an intuitive result, 2) click-through rates are higher during lunch time relative to other times of the day, and 3) click-through rates for technology, entertainment, sports, life and culture news are higher than those for local, international, business and motor news. Our explanation for the findings of the click-through for different news categories is that the former group consists of stories that people tend to personally identify and engage more with. Readers are, then, more likely to seek additional information (beyond what is listed in the article outline) for such stories. On the other hand, local, international and business news tend to be more “impersonal” – most people are interested in knowing that an event happened nationally or internationally but not in finding out more details about it.

Going back to our first variable of interest, i.e., snippet length, and relating it to the theoretical arguments we advanced in Section 3.1, one possible reason for the finding that click-through rates monotonically decrease, is that a snippet length of 98 characters (the smallest length used in the iPhone experiment) is already longer than the ‘optimal’ snippet length that provides just enough information to help users determine if an article is of interest to them. Therefore, any additional information provided via longer

snippets only serves to satiate the appetite of some users for the full story, resulting in lower population-level click-through rates. Repeating the above analysis on the iPad data set provides the benefit of examining what happens when snippet length goes down to zero. In addition, recall that the iPhone snippet length manipulation takes place at the second level of the user interface (Figure 2b), when users have already expressed an interest for the topic (by clicking through the top level, Figure 2a). In contrast, the iPad snippet length manipulation takes place at the lead article at the top level of the interface (Figure 4a), at which point users have seen nothing else about the article. The iPad data set used in this section contains 65,500 observations, each corresponding to a lead article display event. Appendix A.3 lists some descriptive statistics of this data set.

As before, we performed model-free comparisons of the manipulated snippet lengths of the iPad lead articles. First, we compared the click-through probabilities associated with different snippet lengths. As shown in Figure 6, click-through probabilities decrease as snippet lengths increase.¹¹ The associated 95% BCa confidence intervals are no longer disjoint, therefore, further analysis is required to establish whether any of the treatments results in statistically different average click-through probabilities.




Figure 6: Click-through probabilities associated with different snippet lengths on iPad lead article's data.

We applied a permutation test to all pairs of treatments to determine whether the population mean of any of the treatments differs from that of any other treatment. The results (p-values) are depicted in Table 7.

¹¹ The reader will notice that click-through probabilities on the iPad experiment (Figure 6) are much lower than those listed on the iPhone experiment (Figure 5). We attribute the difference to the fact that on iPad, there are several articles displayed on the same screen that users can choose to read (Figure 4a) and we only track the user's decision to click on the lead article.

According to these results, snippet length 0 results in click-through probabilities that are significantly higher from those that result from snippet lengths 245 and higher; snippet length 98 results in click-through probabilities that are significantly higher from those that result from snippet lengths 294 and higher; and snippet length 147 results in click-through probabilities that are significantly higher from those that result from snippet length 343.

	Snippet length					
	98	147	196	245	294	343
0	0.626	0.378	0.062	0.021	0.003	0.001
98		0.685	0.151	0.067	0.015	0.003
147			0.283	0.151	0.036	0.011
196				0.689	0.307	0.139
245					0.53	0.286
294						0.661

Table 7: p-values of permutation test applied to all pairs of snippet lengths on iPad lead article’s data

One concern of the iPhone experiment is that the result regarding the presence of an image could be biased, as an image is displayed only when it is available (in other words, it is not randomized). In the iPad experiment, we randomized the presence of an image at the secondary article that was displayed to the left of the lead article in each section (Figure 4a) as follows: Before displaying secondary article i to user j , if that article had an available image, our system generated a random number between 1 and 7 drawn from a uniform distribution. If that number was between 1 and 3, the image was displayed, whereas if the number was between 4 and 7, it was not. Only sessions where secondary articles had available images (whether these were displayed or not) are included in this data set, resulting in 37,703 observations, of which 16,130 (roughly 3/7^{ths}) had an image displayed and 21,573 (roughly 4/7^{ths}) did not. Appendix A.4 shows key descriptive statistics of the iPad secondary article’s data.

Table 8 shows model-free results of comparing click-through rates for secondary articles with (without) images displayed. It is clear that the presence of an image substantially reduces click-through rates and that the corresponding 95% confidence intervals are disjoint. The results confirm the finding that the presence of an image reduces the probability that users will click on the article, supporting the second part of Hypothesis 1b.

Has image	Observations	Click-through probability	Lower CI	Upper CI
No	21,573	2.96%	2.75%	3.17%
Yes	16,130	2.03%	1.82%	2.23%

Table 8. Click-through probabilities associated with the presence (absence) of an image on iPad secondary article data.

5.2. Impact of the Presence of Multiple Snippets on Group Click-Through Rate

An important function of aggregators is to group together related articles on the same story/topic. Since the aggregator displays a snippet for each article, it is important to know how the presence of multiple snippets affects the probability that any snippet in the group is clicked on. To answer this question we collapsed each topic (story) access session into a single record, resulting in a data set with 25,520 observations. Appendix A.5 shows some key descriptive statistics of this data set. Our independent variable is the number of related snippets displayed during each topic access. Our dependent variable records whether at least one snippet within that topic was clicked by the user during that session.

We performed model-free comparisons to get a sense of what the data looks like. First, we compared the click-through probabilities for story accesses with 1, 2, 3 and 4 or more related articles.¹² As shown in Figure 7, the click-through probability increases as the number of related articles goes from one to two, and then decreases as the number of related articles increases further.

A permutation test provides further evidence that all treatment pairs, except 1 vs. 3 related articles, result in statistically different average click-through probabilities. Our data, therefore, seems to support an inverse U-shaped relationship between group click-through rates and the number of related articles, which peaks at two related articles.

In real life, news stories evolve and publishers report updates on important stories, usually on a daily basis. This means that the same topic can grow in the number of associated articles over time. Our data set contains several examples where a user has opened a topic to find that it contains 1 snippet, while another user has opened the same topic a short while later to find that it contains 2 or more snippets. Our data set, therefore, naturally contains variation in the number of articles snippets within the same topic. Of

¹² Due to the small number of occurrences of story accesses with 4 or more related articles in our data, we binned them together.

course, topics with many snippets also tend to be older and topic age, as opposed to number of snippets, might be the true driver of click-through patterns. We control for this possibility by including a variable topic-age that captures the time between the publication of the first article on the topic and the time of the current user access.




Figure 7. Group click-through probabilities and p-values of associated permutation test for different numbers of related articles.

The regression result is shown in Table 9. Our principal finding is that, even after controlling for topic age, we find results that are compatible with those depicted in Figure 7. Relative to the baseline case when there is only one article about a topic, click-through rates increase when the number of articles increases to two, then decrease to baseline levels for three related articles, and further decrease below baseline levels when there are more than three related articles. Hypothesis 2 is therefore supported.

With regard to the control variables, as expected, topic age is negatively associated with group click-through rates. For the other control variables, we found that click-through for technology, entertainment, sports, life and culture news are higher than those for local, international, business and motor news, and click-through are higher in the morning relative to other times of the day. These findings are consistent with those of Section 5.1.

	Coeff.	Std. Error	Pr(> Z)
Main variables			
related-snippets-1 (baseline)			
related-snippets-2	0.537	0.040	***
related-snippets-3	-0.005	0.100	0.958
related-snippets-4+	-0.260	0.106	*
Control variables			
topic-age	-0.012	0.004	**
category-int. (baseline)			
category-local	0.043	0.043	0.320
category-business	-0.130	0.056	*
category-technology	0.433	0.060	***
category-entertain	0.299	0.055	***
category-sports	0.219	0.070	**
category-life	0.508	0.079	***
category-motors	-0.096	0.131	0.464
category-culture	0.278	0.084	**
time-other (baseline)			
time-morning	0.138	0.054	*
time-lunch	0.099	0.056	0.074
time-afterwork	-0.011	0.056	0.844
time-afterdinner	-0.064	0.055	0.245
language-german (baseline)			
language-french	-0.157	0.239	0.510
language-italian	0.050	0.255	0.984
AIC		32774	
***Significant at < 0.001; **Significant at < 0.01; *Significant at < 0.05			

Table 9: iPhone group click-through rate regression (25,520 topic access sessions).

5.3. Impact of Snippet Length and Accompanying Images on an Article's Choice Probability

Table 10 shows that, across all instances where there were two or more articles on the same story, if users read any articles, 95% of the time they read exactly one article, and only 5% they read more than one.

Number of articles read	Count	%	% (conditional on reading ≥ 1 article)
0	1,287	30.22%	
1	2,828	66.40%	95.15%
2	131	3.08%	4.41%
3	13	0.30%	0.44%

Table 10. When two or more articles are available on the same story, very few users read more than one.

Since, in the majority of cases, users read at most one article per topic, besides knowing how the presence of multiple snippets of the same topic affects the probability that at least one snippet is clicked on, it is equally important to find out what factors determine which snippet is chosen. Specifically, we are interested to find out whether, in the presence of competition from other articles on the same story, an

article's choice probability is affected by the length of the snippet and the presence of an accompanying image. To answer this question, we performed discrete choice analysis. Specifically, we looked at article topic groups containing two or more snippets and where exactly one snippet was clicked. For this analysis, we replaced absolute snippet lengths with dummy variables that indicated whether an article's snippet was longer than (snippet-longer), equal to (snippet-average) or shorter than (snippet-shorter) the average snippet length of the article topic groups. We also added control variables that indicated a snippet's position in the group (position-top, position-second, position-low; the latter variable indicating third or lower position). Key descriptive statistics of this data set are summarized in Table 11.

Total number of samples: 2,828 choice sets (sessions), 6,519 article outlines		
Categorical Variables		
Variable	Frequency	%
snippet-shorter than average	2,702	41.45
snippet-average	1,114	17.09
snippet-longer than average	2,703	41.46
position-top	2,828	43.38
position-second	2,828	43.38
position-lower	863	13.24
has-image	5,507	84.48
(no-image)	1,012	15.52

Table 11. Descriptive statistics of iPhone data set for analyzing snippet choice probability

If we frame a user's article choice as a discrete choice problem, we observe that the utility for each alternative (article) depends on attributes of that alternative. We, therefore, use a conditional logit specification (McFadden 1973), which is a special case of a multinomial logit regression model. The regression result (Table 12) shows that an article's snippet length relative to the group's average snippet length was a significant predictor on it being chosen. The baseline case corresponds to the case where the snippet length is shorter than average. Compared to the baseline case, we notice that having longer than average snippets has a small but statistically significant positive effect on the choice probability. Thus, Hypothesis 3a is (weakly) supported. Moreover, the presence of an accompanying image increases a snippet's within-group choice probability. The effect of having an image is strong and comparable to moving from second to first position on the list of related articles. Hypothesis 3b is, thus, supported. This finding is interesting and should be contrasted with the earlier finding that the presence of an image is associated with a decrease in a snippet's absolute click-through rate when there is no other snippet of the same story. Regarding the control variables, as expected, an article's position has a very important effect on it being chosen, with the topmost article being chosen most often.

	Coeff.	Std. Error	Pr(> Z)
Main variables			
snippet-shorter than average (baseline)			
snippet-average	-0.185	0.480	0.705
snippet-longer than average	0.266	0.112	*
has-image	2.076	0.258	***
Control variables			
position-top (baseline)			
position-second	-2.078	0.112	***
position-lower	-2.022	0.179	***
Log-likelihood		-320.3	

***Significant at < 0.001; **Significant at < 0.01; *Significant at < 0.05

Table 12: iPhone’s snippet choice probability (2,828 topic sessions; 6,519 article snippets).

6. Additional analyses and robustness tests

In this section we report additional analyses and robustness checks that offer further insight into how users allocate their attention in news aggregators.


6.1. Impact of Snippet Length and Accompanying Images on Time Spent at the Aggregator

Our first set of results (Section 5.1) established the presence of a negative relationship between an article’s snippet length and/or the presence of an accompanying image at the aggregator and the probability that a user will click on the link and visit the original article site.

We obtained additional perspective on these results by conducting an analysis of the amount of time that users spend at the aggregator, browsing snippets related to a single topic. The dependent variable of this additional analysis (decision time) is the amount of time elapsing between the initial display of snippets of a topic group (group of related article outlines, Figure 2b) and a user's decision to either click a snippet of that topic (Figure 2c) or to move back to the top level interface (Figure 2a) of the aggregator. We first performed a model-free comparison of decision times associated with different snippet lengths and obtained evidence of a monotonically increasing relationship with non-overlapping 95% confidence intervals (Figure 8).

We obtained more depth by performing Poisson regression, since our dependent variable is a time-to-event quantity and we did not find evidence of overdispersion (decision time mean = 12.41 seconds, variance = 13.02 seconds). The regression results (Table 13) are almost a mirror image of the results of Table 3: we find that people spend more time on the aggregator when snippets are longer, when there are images (has-image=1), and when they do not click on the article link (no-click = 1). This analysis, together with the click-through analysis, provides evidence for the substitution effect of news aggregators

on the content ecosystem: the more information is provided by the aggregator, the more attention users allocate to the aggregator vs. the original article.



Snippet length	Observations	Avg. decision time (secs)	Lower CI	Upper CI
98	3,463	7.16	6.93	7.43
147	3,427	9.10	8.81	9.41
196	3,468	11.02	10.71	11.35
245	3,619	12.27	11.95	12.63
294	3,546	13.72	13.38	14.09
343	3,738	14.73	14.38	15.10

Figure 8. Time spent on the aggregator monotonically increases with snippet lengths.

	Coeff.	Std. Error	Pr(> Z)
Main variables			
snippet-98 (baseline)			
snippet-147	0.211	0.025	***
snippet-196	0.391	0.021	***
snippet-245	0.511	0.023	***
snippet-294	0.604	0.022	***
snippet-343	0.688	0.018	***
no-click	0.185	0.014	***
has-image	0.244	0.018	***
Control variables			
topic-age	-0.001	0.002	0.703
category-int. (baseline)			
category-local	-0.026	0.018	0.161
category-business	-0.117	0.026	***
category-technology	-0.040	0.025	0.120
category-entertain	-0.046	0.024	0.057
category-sports	-0.112	0.030	***
category-life	-0.141	0.035	***
category-motors	-0.234	0.052	***
category-culture	-0.112	0.031	***
time-other (baseline)			
time-morning	0.054	0.021	**
time-lunch	-0.010	0.022	0.669
time-afterwork	0.010	0.023	0.676
time-afterdinner	0.021	0.021	0.334
language-german (baseline)			
language-french	0.028	0.093	0.761
language-italian	-0.018	0.106	0.867
AIC			181452

***Significant at < 0.001; **Significant at < 0.01; *Significant at < 0.05

Table 13: iPhone decision time regression (21,261 single article sessions).

6.2 Impact of Topic Importance on Access Patterns

As mentioned in Section 5.1, Hypotheses 1a and 1b were tested on a data set that only includes sessions that had one article at the time they were accessed. Some of these sessions correspond to topics to which additional related articles were later added, as the story evolved. Each topic can, thus, be associated with the maximum number of related articles (*max-articles*) at the end of its lifespan. *Max-articles* is a plausible measure of that topic's importance.

Max-articles	Sessions	Percentage
1	16,434	77.30%
2	2,432	11.44%
3	745	3.50%
4 or more	1,650	7.76%

Table 14. Distribution of max-articles across all iPhone data sessions with one article at time of access.

Table 14 depicts the distribution of max-articles in the data set used to derive the results of Section 5.1. Over 77% of sessions never had more than one article written about them. It is reasonable to assume that these sessions dominate the statistical results. It is, therefore, interesting to explore whether the impact of snippet length on click-through rates is different for topics with a higher number of max-articles (that presumably correspond to more significant news stories).

Figure 9 provides some evidence to answer this question. The graphs show how, in sessions of a single article, average click-through rates vary as a function of snippet length for topics with different max-article values. The figure shows that, for a given value of max-articles, click-through rates decline as snippet lengths grow. This general trend remains true for all values of max-articles, even though, most likely due to lack of statistical power, the relationship is not as clearly monotonic for higher values. The addition of linear trend lines shows that slopes tend to get a bit less steep for higher values of max-articles. However, the difference is not dramatic.¹³

¹³ Figure 9 also shows that, for a given snippet-length, higher values of max-article are associated with lower click-through rates. This is a somewhat unexpected finding that falls outside the scope of this paper. One tentative interpretation is that users have more opportunities to learn about “important” stories from multiple sources, so when they encounter them in the news aggregator, they have already heard about them and might be less inclined to click through to the entire article. Further analysis is needed to interpret this finding.




Figure 9. Click-through probabilities as a function of snippet length for different values of max-articles.

Overall, our conclusion from this exercise is that the importance of a story, as proxied by the number of articles written about it, does not affect the negative relationship between snippet lengths and click-through rates. This strengthens the evidence for the robustness of our findings related to Hypothesis 1a.

To test the robustness of Hypothesis 1b, we compare click-through rates for articles with/without images for different values of max-articles. As previously, we perform a permutation test with 10,000 iterations to determine whether each pair of click-through rates are statistically different and report the corresponding p-values (Table 15). Our results show that a statistically significant negative impact of images on click-through rates persists for all values of max-articles.

	Max-articles =1		Max-articles = 2		Max-articles >= 3	
	Observations	Click-through rate	Observations	Click-through rate	Observations	Click-through rate
Has image						
No	3,010	68.46%	466	66.91%	692	56.92%
Yes	13,424	64.94%	1,966	57.41%	1,703	50.19%
p-value		0		0.0008		0.0114

Table 15. Click-through probabilities with (without) images for different values of max-articles.

As a final robustness check, we added the average snippet length as an independent variable to the group click-through rate regression of Table 9 (Section 5.2) and redid the regression only for sessions with two or more related articles. We find (Table 16) that the average snippet length is significant and negative, suggesting that, even in settings where there are multiple related articles (that presumably correspond to

important stories), the probability that the user will click through to at least one article has a negative relationship with the average length of snippets displayed. These additional analyses strengthen the evidence for the robustness of findings related to Hypotheses 1a and 1b.

	Coeff.	Std. Error	Pr(> Z)
Main variables			
related-snippets-2 (baseline)			
related-snippets-3	-0.603	0.108	***
related-snippets-4+	-0.843	0.117	***
avg-snippet-length	-0.004	0.000	***
Control variables			
topic-age	0.004	0.016	0.817
category-int. (baseline)			
category-local	0.235	0.093	*
category-business	0.167	0.107	0.119
category-technology	0.204	0.145	0.159
category-entertain	0.112	0.147	0.448
category-sports	0.474	0.121	***
category-life	0.452	0.235	0.074
category-culture	0.045	0.241	0.852
time-other (baseline)			
time-morning	0.113	0.108	0.293
time-lunch	-0.088	0.108	0.418
time-afterwork	-0.019	0.108	0.857
time-afterdinner	-0.125	0.103	0.224
language-german (baseline)			
language-french	-0.259	0.095	**
language-italian	-0.176	0.085	*
AIC			4878.5

***Significant at < 0.001; **Significant at < 0.01; *Significant at < 0.05

Table 16. iPhone group click-through rate regression with average snippet length added. Only article groups with two or more related articles are included (4,259 sessions).

6.3. Understanding the Drivers of Group Click-through Rate Patterns

In Section 5.2 we found an inverse U-shaped relationship between the number of snippets in a topic group and the probability that readers will click on at least one snippet from that group. We found that click-through attained its maximum value when a topic had 2 related snippets displayed but decreased beyond this point. This is a previously unnoticed side effect of news aggregators that deserves more attention.

As discussed in Section 3.2, there are two possible explanations for this finding: complementary information among related snippets, and/or choice overload. According to the complementary information hypothesis, the more snippets are displayed on a topic, the more likely it is that a user will find at least one of them appealing and/or assess that this is an important story worth knowing more about. At the

same time, the combined presence of multiple article snippets potentially provides more information about the story itself. This reduces the residual utility of the original articles and may satisfy the curiosity of readers who then may not feel the need to click on any of the linked articles.

We proceeded to test this hypothesis by quantifying the degree of content complementarity of snippets belonging to every one of the topics that were displayed to users during our two-week iPhone experiment. We employed two student coders for snippets in German, Italian, and French language respectively (six coders in total). Coders were tasked with characterizing each snippet belonging to a topic as being either a replica (containing identical or essentially the same information), an alternative (containing complementary information about the same event), or an update (containing new developments in the story) relative to the information contained in the snippet that was displayed immediately before it. The average inter-coder reliability was 76%.

We consider the fraction of alternatives and updates contained in a topic group as proxies of the degree to which the snippets of the group collectively reveal more information about the story, relative to any single snippet. A correlation check attests that these variables contain distinct information, not captured by other variables in our regression ($\text{Cor}(\text{related-snippets}, \text{alternatives}) = 0.01$, $\text{Cor}(\text{related-snippets}, \text{updates}) = 0.18$, $\text{Cor}(\text{alternatives}, \text{updates}) = 0.10$). We added the fraction of alternatives and updates of each topic group as additional independent variables and repeated the regression of Table 9 (Section 5.2). We found a significant positive relationship between the fraction of alternatives on the group click-through rates (see Table 17). However when we split the analysis between topics with 2 snippets and topics with more than 2 snippets, we found that the number of alternatives and updates was not significant on group click-through rate when there are 3 or more snippets displayed about a topic.

The findings of this additional analysis provide no evidence to support our speculative hypothesis of complementary snippet information being responsible for the reduction in group click-through rates when groups have more than 2 snippets. Support for that hypothesis would have required a *negative* effect of alternatives and updates on click-through rates. Instead we found a positive effect that loses significance when there are more than 2 snippets per group. We must, therefore, conclude that the reason for the drop in group click-through rates past 2 snippets is due to other factors, such as choice overload. Aggregator designers must take this effect into consideration when designing their interfaces.

	Coeff.	Std. Error	Pr(> Z)
Main variables			
related-snippets-1 (baseline)			
related-snippets-2	0.469	0.049	***
related-snippets-3	-0.092	0.105	0.378
related-snippets-4+	-0.318	0.110	**
Control variables			
frac-alternatives	0.229	0.095	*
frac-updates	0.039	0.073	0.590
topic-age	-0.012	0.004	**
category-int. (baseline)			
category-local	0.043	0.043	0.302
category-business	-0.129	0.056	*
category-technology	0.434	0.060	***
category-entertain	0.313	0.055	***
category-sports	0.234	0.070	**
category-life	0.516	0.079	***
category-motors	-0.136	0.132	0.301
category-culture	0.289	0.085	***
time-other (baseline)			
time-morning	0.141	0.054	**
time-lunch	0.113	0.056	*
time-afterwork	-0.003	0.056	0.955
time-afterdinner	-0.034	0.055	0.532
language-german (baseline)			
language-french	-0.113	0.241	0.638
language-italian	0.060	0.257	0.813
AIC			32367

***Significant at < 0.001; **Significant at < 0.01; *Significant at < 0.05

Table 17: Added complementarity analysis in iPhone group click-through rate regression (25,520 topic sessions).

7. Implications and Research Opportunities

News aggregators have emerged as an important component of the digital content ecosystem. A better understanding of how their design parameters affect the allocation of reader attention is useful, both in terms of informing aggregator design, as well as in terms of informing the current controversy that exists between aggregators and content producers. In this study, we conducted field experiments with that objective in mind.

Our first set of experimental results provides evidence of a substitution effect of news aggregators on the content ecosystem: the more information is provided by the aggregator, in the form of longer snippets and accompanying images, the more time users spend on the aggregator and the less likely they are to click-through to original articles. One way of interpreting our findings, especially the results of the iPad experiment, is that anything that aggregators display in addition to an article’s headline, decreases click-

through rates. Based on the above results, content producers have a point in terms of challenging the currently prevailing view that the reproduction of headlines and article snippets by aggregators falls under the “fair use” provisions of U.S. copyright laws (Copyright Act of 1976, 17 U.S.C. § 107).¹⁴ Specifically, one of the factors for fair use set forth by current law is “the effect of the use upon the potential market for or value of the copyrighted work.” Our study shows that such an effect, indeed, exists. Furthermore, this effect is very sensitive to the amount of information that is reproduced by aggregators – practically every character makes a difference to click-through rates.

With respect to the competition among related snippets, we examined what factors determine which snippet(s) in a topic group are more likely to be chosen by readers. As expected, snippets positioned at the top of the list were chosen most often. Controlling for position, snippets whose length was longer than the average snippet length of related articles and, especially, snippets that had an accompanying image, were more likely to be chosen. This is in line with Cutrell and Guan’s (2007) findings that as the snippet length of search engine results gets longer, the relative attention to the snippet increases, and the accuracy of the choice made by search engine users improved. These results reinforce the difficult position in which content producers are placed by aggregators – on one hand, our findings on individual click-through rates suggest that content producers might want to place limits on the amount of article text/graphics that an aggregator is allowed to reproduce. On the other hand, our findings on competition click-through rates show that, doing so unilaterally may place a publisher’s content at a disadvantage vis-à-vis the content of competing publishers who choose to not impose such limits. This argument is consistent with the theoretical predictions of Dellarocas et al. (2013) regarding the prisoner’s dilemma situation that competing content producers are facing in their negotiations with aggregators. The above discussion suggests the need for industry-wide (as opposed to one-on-one) negotiations between content producers and news aggregators with respect to the terms of their relationship.

As with any study, the findings of this paper should be viewed with regard to the study limitations. Although this work offers novel insight into the relationship between news aggregators and content producers, its objective is not to provide answers to the question of whether aggregators are, on balance, beneficial or harmful to content ecosystems. What the current work establishes is that aggregators extract an “attention tax” from content producers, in the form of users who never click through to the original articles. What is outside the scope of the current research is the impact that aggregators have on increasing the overall consumption of content (e.g. because they reduce search costs by organizing

¹⁴ See http://en.wikipedia.org/wiki/Fair_use

content). Despite notable recent attempts to provide answers to the latter question (Athey and Mobius 2012; Chiou and Tucker 2011), a balanced examination of the cumulative impact of aggregators, that takes into consideration both the search cost reduction and the attention tax effects, still remains an elusive and interesting empirical question for future research.

This study is based on the currently standard news aggregator practice of using the first few words of newspaper articles as “snippets” and shows that this practice erodes the audience’s willingness to click-through to read the entire article. An interesting avenue for future research is, thus, to understand how one can construct article snippets that optimize click-through rates by striking the right balance between offering users enough information about the article’s relevance to their interests, while not giving away the article’s content.¹⁵

References

- Agarwal, D., B-C. Chen, P. Elango. 2009. Spatio-temporal models for estimating click-through rate. *Proc. of the 18th Int. Conf. on World Wide Web (WWW'09)*, Madrid, Spain, 21-30.
- Athey, S., M. Mobius. 2012. The impact of news aggregators on internet news consumption: The case of localization. Working Paper, Microsoft Research.
- Bawa, K., R. Shoemaker. 2004. The Effects of Free Sample Promotions on Incremental Brand Sales. *Marketing Sci.* 23(3) 345-363.
- Chi, E.D., Pirolli, P., Chen, K., Pitkow, J. 2001. Using information scent to model user information needs and actions on the web. *Proc. of ACM SIG Human Computer Interactions*, Seattle, USA, 1-8.
- Chiou, L., C. Tucker. 2011. Copyright, digitization, and aggregation. NET Institute Working Paper No. 11-18. Available at SSRN: <http://ssrn.com/abstract=1864203> or doi:10.2139/ssrn.1864203
- Cutrell, E., Guan, Z. 2007. Eye tracking in MSN Search: Investigating snippet length, target position and task types. *Proc of ACM Conf. on on Human Factors in Computing Systems*, San Jose, California, USA, 407-416.
- David P. 1998. Visual-Verbal Association. Do News Pictures Narrow the Recall Gap Between Concrete and Abstract News? *Human Communication Research*, 25(2) 180-201.

¹⁵ Some news aggregators have been known to engage in practices of using “tempting, vacuous, ‘curiosity gap’ headlines” to attract clicks. Such practices are referred to as “clickbait.” (See also, <http://www.theatlantic.com/entertainment/archive/2014/11/clickbait-what-is/382545/>) The authors believe that there is room for legitimate research on the generation of snippets that optimize click-through rates while not misleading readers.

- Dellarocas, C., Z. Katona, W. Rand. 2013. Media, aggregators and the link economy: Strategic hyperlink formation in content networks. *Management Sci.* 59(10) 2360–2379.
- Dewan, R. M., M. L. Freimer, A. Seidmann, J. Zhang. 2004. Web portals: Evidence and analysis of media concentration. *J. Management Info. Sys.* 21(2) 181-199.
- Di, W., Bhardwaj, A., Jagadeesh, V., Piramuthu, R., Churchill, E. 2014. When relevance is not Enough: Promoting Visual Attractiveness for Fashion E-commerce. arXiv:1406.3561 [cs.HC].
- Ghose, A., S. Yang. 2009. An empirical analysis of search engine advertising: Sponsored search in electronic markets. *Management Sci.* 55(10) 1605-1622.
- Greene, S., G. Marchionini, C. Plaisant, B. Shneiderman. 2000. Previews and overviews in digital libraries: Designing surrogates to support visual information seeking. *J. of the American Society for Info. Sci.*, 51(4) 380–393.
- Halbheer, D., S. Florian, O. Koenigsberg, D.R. Lehmann. 2014. Choosing a Digital Content Strategy: How Much Should be Free? *International Journal of Research in Marketing*, 31(2), 192-206.
- Heiman, A, B. McWilliams, Z. Chen, and D. Zilberman. 2001. Learning and Forgetting: Modeling Optimal Product Sampling over Time. *Management Sci.* 47(4), 532-546.
- Hong, S. 2011. A simple model of news aggregators, information cascades, and online traffic. Working Paper. Available at SSRN: <http://ssrn.com/abstract=1927035>
- Isbell, K. 2010. The rise of the news aggregator: Legal implications and best practices. Research Publication No. 2010-10, Berkman Center.
- Iyengar, S. S. and M. R. Lepper (2000) When Choice Is Demotivating: Can One Desire Too Much of a Good Thing? *Journal of Personality and Social Psychology*, 79 (6) 995–1006.
- Jarvis, J. 2008. The link economy vs. the content economy. Buzzmachine blog, June 18, 2008. (<http://www.buzzmachine.com>)
- Kahn, B., and D. Lehmann. 1991. Modeling Choice Among Assortments. *Journal of Retailing* 67(3) 274-99.
- Kannan, P.K. 2013. Designing and Pricing Digital Content Products and Services: A Research Review. *Review of Marketing Research*, 10, 97-114.
- Karp, S. 2007. The web's link-driven attention economy. Publishing 2.0, December 15, 2007. (<http://publishing2.com/>)
- Knobloch, S., Hastall, M., Zillmann, D., and Callison, C. 2003. Imagery effects on the selective reading of internet newspapers. *Communication Research*, 30(1) 3-29.
- Kuksov, D. and Villas-Boas, J-M. 2010. When more alternatives lead to less choice. *Marketing Science* 29(3) 507-524.

- McFadden, D. 1973. Conditional logit analysis of qualitative choice behavior. In P. Zarembka, ed., *Frontiers in Econometrics*: 105-135. New York: Wiley.
- Pirolli, P., Card, S.K. 1999. Information foraging. *Psychological Review* 106(4) 643-675.
- Roos, J., C. F. Mela, R. Shachar. 2011. Hyper-media search and consumption. Working Paper, Duke University.
- Sargent, S. L. 2007. Image effects on selective exposure to computer-mediated news stories. *Computers in Human Behavior* 23(1) 705-726.
- Scheibehenne, B., Greifeneder, R., Todd, P.M. 2010. Can there ever be too many options? A meta-analytic review of choice overload. *Journal of Consumer Research* 37(3) 409-425.
- Shapiro, Carl and Hal R. Varian 1998. *Information Rules: A Strategic Guide to the Network Economy*. Boston, MA: Harvard Business School Press.
- Wu, F., B. Huberman. 2008. Popularity, novelty and attention. *Proc. of the 9th ACM Conf. on Electronic Commerce (EC'08)*, Chicago, USA, 240-245.
- Xiang, Y., D.A. Soberman. 2011. Preview provision under competition. *Marketing Sci.* 30(1) 149-169.
- Zeileis A (2006). Object-oriented Computation of Sandwich Estimators. *Journal of Statistical Software*, 16(9), 1–16.
- Zillmann, D., Knobloch, S., and Yu, H.-S. 2001. Effects of photographs on the selective reading of news reports. *Media Psychology* 3(4) 301-324.