Marketing Science Institute Working Paper Series 2020
Report No. 20-126

What Leads to Longer Reads? Psychological Drivers of Reading Online Content

Jonah Berger, WendyW. Moe, David A. Schweidel

What Leads to Longer Reads? Psychological Drivers of Reading Online Content

Jonah Berger

WendyW. Moe

David A. Schweidel

April 2020

Jonah Berger G berger@wharton.upenn.edu) is an associate professor of marketing at the Wharton School at the University of Pennsylvania, Wendy W. Moe (wmoe@rhsmith.umd.edu) is Dean's Professor of Marketing at the University of Maryland's Robert **H.** Smith School of Business, and David A. Schweidel (dschweidel @emory.edu) is Professor of Marketing at the Emory University's Goizueta Business School. The authors that Yuhao Fan, Nir Grinberg, Mor Naaman, and Russel Richie for assistance with various aspects of the project, Jordan Etkin and attendees at the Marketing Science and Association for Consumer Research Conference for helpful feedback on the research, and the company for providing the data.

ABSTRACT

More and more consumers read content online. They scan *Wall Street Journal* articles, catch up on sports, and peruse biogs on tech and celebrity gossip. But sometimes people read the whole piece of content and other times they only read a small portion. What about certain articles encourage sustained attention? Combining natural language processing of a unique dataset of over 700,000 page-reading sessions from over 35,000 articles with three experiments, we examine how textual features (i.e., the words used) shape reading. Results suggest that emotion plays an important role. Importantly, however, not all emotions increase reading. Consistent with research on appraisal and action tendencies, content that evokes anger and anxiety encourage sustained attention while content which evokes sadness discourages it. Textual features that increase processing ease (e.g., concreteness and familiar words) also increase sustained attention. Experimental evidence underscores the causal impact of emotion on reading and demonstrates that these effects are driven by uncertainty and arousal. These findings shed light on psychological drivers of reading and how to design content that will gamer sustained attention.

KEYWORDS: sustained attention, online content, digital marketing, content marketing

People have been telling stories for thousands of years. Our early ancestors sat around campfires sharing stories of the hunt. Epic poems like the *Epic of Gilgamesh,* the *Iliad,* and the *Odyssey* were passed down verbally from generation to generation before eventually being written down. And before there was the *New York Times* or *Wall Street Journal,* early journalists circulated handwritten news sheets.

But why do some stories hold our attention more than others? What leads to sustained attention?

While this question has ancient roots, it is just as relevant in today's digital age. The average American spends 24 hours a week online, and much of that time is spent reading (Cole, Suman. Schramm, and Zhou, 2017). People browse the latest news from the *New York Times,* read about sports at ESPN.com, and peruse tech blogs and celebrity gossip. The rise of content marketing has only exacerbated this trend. More than 86 million blog posts are published every month (Greesonbach 2018), and companies spend millions creating and distributing content to attract and retain customers.

But as anyone who has ever read a article or blog post can attest, not all content generates sustained attention. For some articles, people only read a couple paragraphs before moving on to something else. For others, most people finish the entire piece. Why do certain pieces of content hold people's attention more than others?

Efforts to answer this question have been hampered by data availability. When reading a magazine or physical newspaper, for example, there's no record of which articles people read, let alone how far they got through the article. Further, while online metrics like views provide information on what content gets clicked on, they don't provide any insight into how much of that content, if any, actually gets consumed.

This article fills this gap by investigating how content characteristics shape sustained attention. We use natural language processing (Humphreys and Wang 2017) to analyze a unique dataset of over 700,000 reading sessions from over 35,000 articles from nine major online publishers. This data allows us to examine, for a given person reading a given article, how textual features of a given paragraph (i.e., the words used) shape whether someone keeps reading. In addition to features that should impact processing ease (e.g., concrete language and whether familiar words are used), we examine how the valence of content (i.e., positive or negative) as well as the specific emotions it evokes (e.g., anxiety versus sadness) affect whether users continue reading. To supplement our empirical analysis of field data, we also conduct three experiments. They provide direct evidence that emotions influence reading and demonstrate the underlying processes behind these effects.

Our findings make four main contributions. First, while research has examined some aspects of online behavior (see Berger, 2014; Chen and Stephen 2018 for reviews), there has been less attention to what drives sustained attention (i.e., reading). A rapidly growing literature has begun to examine how emotion shapes word of mouth, the usefulness or helpfulness of online reviews (Chen and Lurie 2013; Yin, Bond, and Zhang 2014; 2017), and whether reviews drive conversion (Ludwig et al. 2013). Other work on narrative transportation (e.g., Busselle and Bilandzic 2009; Escalas 2004; 2006; Green and Brock 2000) finds that narratives can increase self-brand connection and brand evaluations.

But while these outcomes are important, little work has examined sustained attention. What leads people to continue to scroll through a given piece of content rather than stop or switch to something else instead. While "engagement" is often used to describe consumers' relationship with online content, it can refer to a diverse range of behaviors including sharing, visiting, or being attentive (Calder, Malthouse, and Schaedel 2009). We examine a narrower question . Not what causes people to share something or visit a site, but simply given they have opened an article, what leads them to continue reading it. This novel dependent variable is the focus here.

Second, while some work has examined how visual features of content (e.g., pictures, layout, or the presence of ads) impact attention, there has been less research on how textual features might shape sustained attention. Our findings provide insight into how the emotions evoked by content and how easily the content can be processed impact whether people continue reading. Further, they shed light on the underlying psychological processes that drive such effects.

Third, from a practical perspective, our findings help content creators design content that is more likely to receive sustained attention. Content marketing is expected to be a $300 billion industry by 2019 (McCoy 2017). But while many have focused on metrics like clicks or views, clicks don't always translate into reads, and content has more impact if people actually read it. Further, while social shares (e.g., retweeting) increase reach, there's little relationship between shares and reading (Jeffries 2014). While people's comprehension of what they read is generally high (Jacoby and Hoyer 1987), comprehension requires attention. If content creators want people to read their content, they have to understand what drives sustained attention in the first place. This work provides a set of actionable directions (e.g., use more concrete or familiar language, or evoke emotions associated with arousal or uncertainty) that creators can use to craft content that encourages reading.

Note that we focus on features of content rather than individual differences. While particular people may be more or less motivated to read, regardless of idiosyncratic individual tendencies, different articles can also differentially evoke certain emotions and encourage (or discourage) reading. This approach makes our results more useful to content creators and content marketers. While content providers may not have detailed (or any) data on the specific people reading their content, they can influence content features to encourage reading and deepen attention.

Finally, we demonstrate how natural language processing can shed light on consumer behavior. Automated textual analysis provides a rich set of tools for extracting behavioral insight from text, but these tools have just started being adopted by marketing researchers (Netzer, Feldman, Goldenberg, and Fresko 2012; Netzer, Lemaire, and Herzenstein 2018; Moore and Mcferran 2017; Packard, Moore, and Mcferran 2018; Rocklage and Fazio 2015; Rocklage, Rucker, and Nordgren 2018; Tirunillai and Tellis 2014; 2017; see Humphreys and Wang 2017 for a review). We not only illustrate how a range of features can be used, but how they can deepen understanding around the psychological drivers of behavior.

RELATED RESEARCH

Sustained attention is beneficial for a variety of reason. Some work highlights the link between attention and advertising effectiveness. Consumers that found websites more engaging, for example, reported more positive attitudes towards ads on that site and greater intention to click (Calder, Malthouse, and Schaedel, 2009). Other work suggests that attention may increase purchase. Looking across a range of media channels (e.g., television, magazines, and the internet), people that were more engaged in content (e.g., a television show, news article, or website) reported greater willingness to purchase a product advertised on that content (Kilger and Romer 2007).

From an online media outlet perspective, how much consumers read can also directly impact advertising revenue. If a media outlet shows banner ads every quarter of the way through an article (i.e. one quarter of the way through, halfway, three-quarters, and at the end), then how far down consumers read will impact the number of ads they see. For display ads, this, in tum, will impact the revenue the media outlet receives.

Sustained attention also determines whether content has impact. Beyond traditional media outlets like newspapers, other organizations create content to engage their audience. Consulting firms write white papers to attract clients, and nonprofits write articles to educate readers and drive donations . But the impact of these efforts depends on people actually reading the content. If people don't read it, it is unlikely to boost their knowledge or encourage them to hire the firm.

A key question, then, is what drives sustained attention.

Some research has examined how things like media type, layout, ads, or the presence of pictures or videos impact attention (Lagun and Lalmas, 2016; Lagun and Agichtein, 2015). When reading the physical paper, people usually quickly glance at the right page of a spread before moving to the left (driven in part by the page turning process), while in online content, the top of the page gets the most attention (see Leckner 2012 for a review). Within textual content, photos can increase the length of time people spend reading an article (Zillmann, Knobloch, and Yu, 2001). Pictures and graphics can also act as "entry points" that encourage readers to focus on the text around them (Garcia and Stark 1991).

While prior work has provided insight into how visual or channel features (e.g., online or off) shape sustained attention, there has been less attention to how textual elements of content impact reading. Indeed, while many authors have theorized about storytelling, narrative, and related topics (Gergen and Gergen 1988; McKee 1998; Booker 2004; Vonnegut 2005), little empirical work has actually investigated these claims (though see Van Laer et al 2018).

Research on textual features has mainly focused on labeling texts with their appropriate grade level. Some sentences can be read by most people (e.g. "the dog likes food"), while others require more specialized knowledge or vocabulary that comes with education (e.g., "Natural gas is used to heat our homes and run some transportation"). Standard indices such as Flesch-Kincaid Automated Readability Index (Kincaid, 1975), SMOG (McLaughlin, 1969), and Coleman-Liau (Coleman and Liau, 1975) use simple features like sentence and word length to approximate readability. Longer sentences and longer words tend to be more complex and require a higher-grade level to be able to comfortably parse.

But while word and sentence length are certainly important, they provide less insight into how other aspects of content impact sustained attention. Further, while computer scientists have begun to examine a broader set of textual features (Pitler and Nenkova 2008), they have focused mainly on predicting self-report ratings of text-quality on a small set of articles. Similarly, while psychological research has examined reading comprehension (e.g., Just and Carpenter 1980; Freebody and Anderson 1983) or learning (see Hidi 1990 for a review), it been limited by small samples (e.g., 15 passages of 132 words each) and tasks where participants are forced to read an entire passage of text. To truly begin to understand how textual features shape continued reading, it is important to look across a broad range of both people and content types.

THE CURRENT RESEACH

A great deal of prior work suggests that motivation and ability are two major drivers of behavior (see Fogg, 2009 for a recent perspective). The more motivated people are to take action, the more likely they will do so; and the more able they are to take that action, the more likely it is to occur. Applied to reading, these dimensions can be operationalized by the emotion content evokes (i.e., which can provide motivation) and how easy the text is to process (i.e., which shapes people's ability to continue reading). We start by reviewing our theorizing on processing ease and then discuss emotion.

Processing Ease

We suggest that textual features that make passages easier to process should increase reading. These aspects should reduce the effort required to continue, and thus increase the likelihood that people do so. This should play out across a number of content features.

First, at the most basic level, standard measures of readability should influence processing ease. As noted above, readability measures assign texts an appropriate school grade level (e.g., Kincaid 1975). Shorter words and sentences should make content easier to process and thus encourage continued reading.

This approach to readability is rather simplistic, however, and things like syntactic complexity should also play a role. Measures like Flesch-Kincaid only consider the words in a sentence, but they ignore sentence structure. Linguistic and psycholinguistic theory suggests that as people read, they construct syntactic representations of sentences, or parse trees (Chomsky 1957). Deriving meaning depends on the ability to mentally construct such a parse, and, as a result, sentences with more complex parse trees are more difficult to build and understand (Pitler and Nenkova 2008; Schwarm and Ostendorf 2005). This suggests that greater parse tree height (i.e., more complex sentence structure) may decrease reading.

Even beyond word or sentence complexity, however, familiarity should also shape processing ease. A great deal of research suggests that familiar stimuli are easier to process (e.g., Winkielman & Cacioppo 2001). Work on mere exposure (Zajonc 1965), for example, suggests that the more people see something, the more they like it, in part because repeated exposure makes things easier to process. Consequently, one might imagine that passages that contain more familiar words are easier to read and thus people are more likely to continue reading

Finally, concreteness may also play a role. While some things in the world are relatively abstract (e.g., trust or values), others (e.g., birds or shoes) are relatively concrete. Concrete items tend to be easier to visualize or imagine and require less cognitive resources to process (Friederici, Opitz, and von Cramon, 2000; see Connell and Lynott 2012 for a review). This greater processing ease, in tum, may encourage continued reading. Indeed, preliminary work suggests that concreteness increases comprehensibility, interest, and recall (Sadoski, Goetz, and Rodriguez, 2000).

Taken together, we examine how processing ease, as represented by readability, parse tree height, familiarity, and concreteness, influences reading. [1]

Emotion

Beyond processing ease, however, we suggest that content should also impact reading through the emotions it evokes. Indeed, prior work on narrative transportation suggest that transportation may impact experienced emotion (Escalas 2006). Emotion might shape reading in three ways, through (1) emotionality, (2) valence, or (3) specific emotions.

[1] While Nielsen and Escalas (2010) find that disfluency can increase content evaluations under narrative processing because consumers are more likely to be transported into the story, they focus on situations where the task is to consume the full content (rather than being optional). In our case, content consumption is entirely optional. Consequently, while we agree that if consumers are committed to consuming a full piece of content, disfluency can sometimes increase engagement, in situations where content consumption is more optional, and consumers can opt out at any point, we predict that processing ease (rather than disfluency) should be more be more likely to encourage sustained attention.

The most basic possibility is that emotionality, or sheer amount of emotion, influences reading. Emotions can increase attention (Easterbrook 1959; Vuilleumier 2005) and flag that something is important and deserves further processing. This perspective suggests that content that evokes any emotion, regardless of type, should encourage continued reading.

A second possibility is that valence drives reading. The simplest way to organize emotions is by valence, or whether something is positive or negative. Some emotions (e.g., happiness) are associated with positive states, while others (e.g., sadness) are associated with negative ones. Research on word of mouth, for example, finds that positive reviews are valued less than negative ones (Chen and Lurie 2013). One could argue that positive emotion should increase reading. People like feeling good and tend to approach positive stimuli, so positive content might encourage reading. Negative content might discourage reading as people avoid bad news. At the same time, however, one could argue the opposite. Research on negativity bias (Baumeister, Finkenauer, and Vohs 2001; Rozin and Royzman 2001) finds that negative information gamers greater attention. When forming impressions of others, for example, people tend to spend longer looking at negative photographs (Fiske 1980). Similarly, the old news adage, "if it bleeds, it leads," is based on the notion that negative information will grab the viewers' attention. These notions would suggest that negative content should encourage reading.

In contrast, we suggest a third possibility: that the relationship between emotion and reading is more complex than valence alone. In addition to being positive or negative, emotions are also characterized by different cognitive structures (i.e., appraisals) and action tendencies (Smith and Ellsworth 1985, Lerner and Keltner 2000; 2001). While anger tends to be associated with appraisals or assessments of certainty, and lead people to feel certain about things in their environment (even beyond the inciting stimulus), anxiety tends to be associated with appraisals, or assessments, of uncertainty. These appraisals, in tum, can lead different emotions to have different action tendencies, or effects on downstream judgments and behavior (Cavanaugh, Bettman, and Luce 2015; Yin, Bond, and Zhang 2014). Feeling anxious, for example, can lead people to take action to reduce feelings of uncertainty (Ragunathan and Pham, 1999). Consequently, work on the Appraisal Tendency Framework (see Han, Lerner, and Keltner 2007 for a review in marketing) suggests that even among emotions of the same valence, different specific emotions (e.g., anxiety versus sadness) may have different effects on reading.

Applied to the present context, we suggest that the impact of specific emotions on reading may be driven by how they shape (1) uncertainty and (2) arousal. Uncertainty involves not knowing or not being sure about something. While certain emotions (e.g., anger) tend to be characterized by certainty, others (e.g., anxiety or fear) tend to be characterized by a state of uncertainty and uncertainty reduction (Ragunathan and Pham, 1999; Lerner and Keltner 2001). These feelings of uncertainty, in tum, increase attention, search, and information processing as people they try to resolve predictions about what will happen next (Tiedens and Linton 2001; Weary and Jacobson 1997). If someone feels anxious about whether it's going to rain, the accompanying uncertainty might lead them search for and carefully process information that resolves that uncertainty (e.g., checking the weather). Applied to reading, we suggest that because uncertainty encourages attention to, and processing of, relevant stimuli, emotions associated with uncertainty (e.g., anxiety and sadness) should encourage reading.

Beyond uncertainty, emotions are also characterized by differences in arousal. Arousal is a state of being physiologically alert, awake, and attentive (Yin, Bond, and Zhang, 2017; see Heilman 1997 for a review). While some emotions (e.g., anger and anxiety) are characterized by high arousal, others (e.g., sadness) are characterized by low arousal. A great deal ofresearch finds that emotionally arousing stimuli attract attention (see Mather, 2007 for a review). Arousal-biased competition theory (Mather and Sutherland 2012) suggests that arousal particularly increases attention for high priority stimuli, or those that are relevant to the task at hand. Further, arousal may lead to an increased state of vigilance (Pham 2004) which should encourage sustained attention. Applied to the context of reading, given arousal can increase attention and vigilance toward relevant stimuli, we suggest that emotions characterized by high arousal (e.g., anger and anxiety) should encourage people to continue reading.

EMPIRICAL TESTS

We test these predictions regarding emotion and processing ease in both the field and the lab. First, we use natural language processing to analyze over 700,000 page read events from over 35,000 online articles (Study 1). We examine whether people are more likely to continue reading articles whose content evokes certain emotions, or is easier to process. Second, to directly test specific emotions' causal impact, we conduct three experiments (Studies 2-4). We manipulate specific emotions and measure the impact on reading. We also measure arousal and uncertainty to test whether they can explain the observed effects.

Note, while this work relates to Berger and Milkman (2012), it differs in some important ways. First, we look at a completely different dependent variable. While Berger and Milkman (2012) investigated whether people share articles, we examine whether or not people continue to read them. These are two completely different actions, so it is not clear a priori whether the effects would be the same or different. Indeed, other work finds that sadness increases donations (Small and Verrochi 2009) so one could just as easily argue that sadness might increase reading. Second, we examine a different process. While Berger and Milkman (2012) focus on the role of emotional arousal in driving effects of specific emotions on sharing, we also examine the effect of uncertainty. Third, we examine additional textual features. While Berger and Milkman (2012) just examined specific emotions, we also examine features that should impact processing ease (e.g., concreteness and familiarity).

STUDY 1: EMPIRICAL ANALYSIS OF OVER 700,000 PAGE READ EVENTS

We worked with a major content intelligence company that tracks sustained attention for online publishers. For the last two weeks of October 2014, they provided a representative random sample of page-read events (i.e., cases where a given page with an article on it was accessed by a user) from nine popular online news sites. To allow for data privacy, we will not disclose the exact outlets, but give a sense of the type of content. We selected sites to cover a wide range of topics, including global news and business (think New York Times or Wall Street Journal), sports (think ESPN.com), local news (think San Francisco Chronicle), technology (think Gizmodo), and celebrity news, fashion, and lifestyle content (think Jezebel). We selected these sites in particular because they used fixed layouts (i.e., content is laid out the same way across articles), do not have ads within the text, and are not responsive, meaning the page shows up the same way regardless of the device. This last point means that regardless of whether an article was read on a phone, desktop, or other device, the content was not reformatted based on viewport size and the line breaks are exactly the same.

We focus only on page-read events that involve some interaction. In some cases, readers may click on an article only to leave right away. In others, they may open an article but then open another browser tab and do something else. To avoid such "bounce backs" and other cases where users are unlikely to be reading much, if any, of the article, we focus on page-read events that involve some interaction. We rely on the company's definition, which involves anyone who had at least two interactions with the page (e.g., mouse scrolls or clicks). Someone who opens an article only to close it a few seconds later is unlikely to have been exposed to much of the article's content. Further, while they may have stopped reading based on the article title or topic, it is unlikely they stopped due to deeply processing textual features of the article itself.

Given our interest in textual features of content, we focus on articles rather than other content types (e.g., videos). Given most of the textual features we consider are based on the English language, we focus on English language articles only. The final dataset involved 717,288 page-read events from 35,448 articles (see Appendix Table 1 for summary statistics).

For each read-event, we have information about the page (collected based on the URL) and the device through which someone accessed the page. While each site assigns a user a unique ID each time they visit the site, these user IDs are not tracked over time or across sites. Consequently, while the data allows us to compare behavior across users for a given article, it limits our ability to examine repeat user behavior across articles.

Dependent Variable. We are interested in how content shapes reading. While one could imagine modeling the average reading depth across users as a function of the average textual features of an article, such an aggregate approach ignores key micro-level details. First, behavior can't be influenced by content people haven't read, so it's important to only predict a given users behavior based on content that appeared before where they stopped. Second, an aggregate analysis ignores paragraph-to-paragraph variation. In a story about the economy, for example, some paragraphs may evoke anxiety while others may not. Further, even if the entire article evokes little anxiety overall, if a certain paragraph evokes some, that may encourage sustained attention.

Consequently, we take a more fine-grained approach, examining how the text of each paragraph relates to whether a user continues to read into the next paragraph. In other words, we conceptualize reading as a process where at the end of each paragraph the reader either continues reading or stops.

This approach also helps address several alternative explanations. As discussed in greater detail below, we include various controls (e.g., publisher, device, and article topics) to try to rule out selection concerns (e.g., certain types of people tend to read certain types of articles). That said, one could still argue that some unobserved feature of the content is what is driving sustained attention. To address this concern, we use article-level content features as a control. This allows us to provide a more fine-grained test. Even controlling for anxiety in the rest of the article, for example, whether paragraphs that evoke more anxiety still increase reading.

To capture paragraph-to-paragraph reading, we measure how far down the page a user scrolls (see Appendix for a more detailed description and example). This is determined using JavaScript code that is embedded on the publishers' sites and executed on the user's browser when an article page is loaded (i.e., for each read event). The code records the pixel position a user scrolls to on the page. This is specifically defined as the top position that is visible on the user's screen. Pixel length starts at 0 for every page and increases up to the length of a given article. If a page was 1000 pixels long, for example, the user would start at 0 and pixel depth would increase, potentially up to 1000.

For each page-read event, we map pixel length to a position within the article. The conversion from pixel length helps us get a measure of reading depth that is independent of site layout. Further, it ensures that our measure is consistent across devices and that we only examine text as opposed to other elements like ads and comments. For this conversion, we developed a custom CSS selector, unique for each site, to identify the page content. We downloaded each article and visually rendered the page using the PhantomJS JavaScript library. From this rendering we extracted the pixel location of the top of each paragraph to know whether the user read past this point during a given page-read event. By selecting sites that have non-responsive layouts, we ensure that the pixel depth remains consistent across devices, screen resolutions and window sizes. The content is not scaled and text is not re-flowed based on screen resolution or browser window size. For example, the same page would be 1000 pixels on both a low-resolution mobile device or high-resolution screen. We validated the results of the pixel conversion process to by manually verifying it for a small set of articles. [2]

While one could be concerned that differences in how people interact with desktop vs. mobile might drive our effects, we address this two ways. First, we control for device. Second, we perform the same analyses on just the desktop and tablet data and find the same effects (Table A5). Thus, while differences between device interactions may certainly shape reading, it does not seem to be able to explain our effects.

Across all page-read events, readers read an average of 48.43% of the article text (SD= 28.97%), with an inter-quartile range of [24.34%, 73.48%]. This suggests a high degree of variation across articles and page-read events. Readers read an average of 8.63 paragraphs (SD= 7.53), with an inter-quartile range of [3,12]. The distribution of page-read events across articles is heavily skewed, with an average number of page-read events per article of 20.23 (SD = 117.42), but as noted below, all the main results are robust to focusing only on articles that have

[2] One could wonder whether our measure truly captures reading. For example, some people may quickly scroll without reading, so one could argue that our measure simply captures how far someone scrolled and nothing else. If that was the case, however, it should make it harder to find effects of textual features on scrolling. If someone simply scrolled, without reading any of the words, then words should have no impact on how far they scrolled. Thus, unless skimming or deeper scrolling is somehow driven by some alternative feature that is also correlated with the textual features we examine, if anythin g, this measure provides a conservative test of our hypotheses. The less people are reading the weaker any relationships between textual features and reading should be.

a significant number of page-read events in the data (e.g., 100 or 250)

Independent Variables

As noted in the introduction, we examine how textual features linked to emotion and processing ease shape reading (See Table A2 for summary statistics for the main features and Table A3 for examples of paragraphs that score highly on each key dimension).

Emotion. We test the relationship between emotion and reading in two ways. First, we measure the sentiment or valence of the text using Linguistic Inquiry and Word Count (LIWC, Pennebaker et al 2015). Following prior work (Berger and Milkman 2012) we separately measure both positive and negative emotions to examine how each relates to reading.

Second, we measure specific emotions. Compared to positive emotion, specific negative emotions are easier to distinguish (Keltner and Lerner 2010) and tools exist to extract them from text, so the analyses focus there. Mohammad and Bravo-Marquez (2017) asked people how much different online content evoked anger, anxiety, and sadness and then used machine learning to extrapolate these responses to new content. [3] We use this approach to measure the amount of anger, anxiety, and sadness evoked by each sentence and then average those to get a score for each paragraph. Pre-test data suggests that these measures are able to reliably distinguish between articles that make readers feel more or less of each of these emotions. [4]

[3] We used Mohammad and Bravo-Marquez (2017)'s approach rather than LTWC for several reasons. First, LTWC simply measures the presence or absence of individual words, which is quite coarse. Mohammad and Bravo-Marquez (2017) use a more continuous method that allows for more accurate variation. Second, while the validity of LIWC's positive and negative emotion categories have been tested in several studies, the specific emotion categories have less empirical support. Third, a small-scale comparison showed that Mohammad and Bravo-Marquez (20 l 7)'s approach was more predictive of manual coding of the articles. That said, to ensure that the results are robust to different methodologies, we also run a version where we use LIWC inste ad. Results are almost identical.

[4] To test whether this method captured differences in felt emotion, we conducted a pretest. We randomly selected articles that scored highly on the different emotion dimensions. Three articles that scored high on anger, three that scored high on anxiety, three that scored high on sadness, as well as three control articles that did not high on any

Processing Ease. To begin to measure processing ease, we use *Flesch-Kincaid (1975) grade level.* This measure combines elements like word and sentence length to get a sense of the number of years of education generally required to understand a text.

To provide a deeper measure of syntactic or sentence complexity, as discussed in the theory section, we also use *parse tree height* (Pitier and Nenkova 2008; Schwarm and Ostendorf 2005). This counts the number of steps it takes to get from the root, or top node, to the bottom most node. While both "The cat on the hot tin roof meowed at my parent's house" and "The cat on the hot tin roof at my parent's house meowed" are similar sentences and involve the same words, the second one has a taller parse tree (i.e., 8 edges vs. 6 edges tall).

To measure the *familiarity* of the words in each paragraph, we used ratings from Paetzold and Specia (2016). They used bootstrapping with word embeddings to extend the MRC Psycholinguistic Database (Coltheart, 1981) from around 9,000 words to over 85,000. Participants in the original studies (e.g., Gilhooly and Logie 1980) rated words based on how familiar they were, where 1 = never seen, heard, or used and 7 = seen, heard, or used every day.

A similar approach was used to measure *concrete* language. We used Paetzold and Specia (2016)'s bootstrapped ratings with word embeddings building on the MRC database. Participants in the original studies (e.g., Spreen and Schulz 1966) were asked to rate concreteness on a 7-point scale (1 = least concrete, 7 = most concrete). Words referring to objects, materials, or people received high concreteness rating. Words referring to abstract concepts that could not be experienced by the senses (e.g., the word "facts") received low

---

dimension . Then, participants (N = 330) were given one of these articles and asked to rate the degree to which the article made them feel the focal emotions (i.e., "how angry did reading the article make you feel?" "how anxious did reading the article make you feel?" and "how sad did reading the article make you feel?" 1 = Not at all, 7 = Extremely). As predicte d, results indicate that the articles evoked the expected emotions. Articles scored as high anger evoked more anger than other articles ($M = 4.59$ vs. 2.50, $t(326) = 8.99$, $p < .001$), articles scored as high anxiety evoked more anxiety than other articles ($M = 3.14$ vs. 2.51, $t(326) = .99$, $p < .001$), and articles scored as high sadness evoked more sadness than other articles ($M = 4.23$ vs. 2.64, $t(326) = 2.75$, $p = .006$).

concreteness ratings. Note that this measure is highly correlated with imagery (r = .93), consistent with our suggestion that concrete words are easier to imagine. We find the same results using Brysbaert, Warriner, and Kuperman (2014)'s concreteness measure.

Controls

As discussed, various factors may affect reading that have little to do with the content itself. Consequently, we control for such features to rule out alternative explanations and test the robustness of the effect.

Publisher. Different publishers may attract different types of readers, attract readers when they have more or less time to read, or publish types of articles that encourage longer or shorter reads. People who read CNBC may have longer attention spans than those that read ESPN, for example, or the same person may read the CNBC (ESPN) when they have more (less) time. Similarly, the *Wall Street Journal* tends to publish business news, which may hold people's attention less than sports. Thus, we use dummy variables to control for the site on which a given article was published.

Reading Device. Similarly, the device on which an article is read (e.g., mobile vs. desktop) should impact reading. Different types of people may use different devices, people may use different devices at different times, and different devices may themselves impact behavior (Ransbotham, Lurie, and Liu 2018). Younger people may read on their phones, for example, while older people read on their desktop. People may skim online content on their phones on their morning commute while lazily browsing on their tablets when they have more time on the weekend. Given the smaller screen size, reading on a mobile device may itself encourage shorter reads. To address these possibilities, we use dummy variables to control for whether users read an article on mobile, desktop, or tablet (0.5% of page-reads are from an unknown device).

Article Topic. Article topic may also influence reading. Articles about certain topics may attract different types of readers, attract readers when they have more time, or impact reading in other ways. To control for this, we control for article topic. Rather than divide articles into discrete categories, we take a more fine-grained approach, performing topic modeling across the entire set of articles and allowing each to be represented as a proportion of different topics.

We use latent Dirichlet allocation (e.g., Blei et al. 2003), a common topic modeling framework that assumes each article can be represented as a mixture of topics (e.g., Berger and Packard 2018; Tirunillai and Tellis 2014). The data generating process assumes that for each word position in a document, a topic is drawn and a word is then drawn conditional on the topic. The posterior distribution of topics can be used to characterize the content of an article. We estimate a 25-topic solution and calculate the posterior topic distribution across topics (see Appendix Table A4 for the distribution of topics across articles). We are interested in controlling for the distribution of topics, and not the exact topics themselves, but example topics include things like government (i.e., words like state, law, and govern$), sports (i.e., words like game, team, and sport) and personal technology (i.e. words like app, google, and file, see Appendix Table A4 for example words from each topic). Posterior topic probabilities for a given article sum to 1, so we include the posterior topic probabilities for 24 of the topics as control variables. These variables represent the relative prevalence of each topic compared to others.

Paragraph Length. Given limited attention spans, the longer one paragraph is, the less likely people may be to read the next one. Consequently, we control for paragraph length using the number of words in each paragraph.

Position in Article. People may be less likely to continue reading the longer they have read already. Alternatively, the more someone has read already, the more invested they may be in the content and the more likely they are to continue. Either way, we control for how long someone has been reading using the article length in words up to that point. We use both a linear and quadratic term to allow for non-linearities.

Percentage Read. Continued reading may also depend on where someone is in the article. Work on goal gradient (Hull 1932), for example, might suggest that readers are more likely to read a paragraph if they are almost done with an article. Other work suggests that motivation is lowest in the middle of goal pursuit (Bonezzi et al. 2011). Consequently, we control for percentage read so far using both a linear and quadratic term.

Article Level Feature Controls. Finally, to further control for content differences across articles, as discussed previously, for each feature we examine at the paragraph level, we also control for it at the article level.

Analysis Strategy

We conceptualize each reading session, *i,* as a sequence where at the end of each paragraph the reader either continues reading or stops. We denote the action made after paragraph *j* of reading session *i* as Y_{ij}, in which $Y_{ij}=1$ if the reader continues to the next paragraph and $Y_{ij}=0$ if they do not. We assume that the probability of continuing past paragraph j in reading session *i* is a function of the paragraph-level content variables and control variables. Formally, we estimate the following logistic regression:

$$Y_{ij} \sim Bernoulli(p_{ij}) \text{ where}$$

$$logit(p_{ij}) = \beta_0 + \sum_k \beta_k x_{ijk} + \sum_c \gamma_c z_{ijc}$$

where Xijk denotes the kth independent variable that characterizes the content of paragraph} in reading event *i* and Zijc denotes the dh control variable.

This analysis is consistent with survival analysis and can be viewed as a discrete time hazard model, similar to methods used to model binge viewing behavior (Schweidel and Moe 2016), clickstream behavior (Sismeiro and Bucklin 2004, Moe 2006), and multiple item purchasing (Harlam and Lodish 1995). Effectively, we model the time (as measured in paragraphs read) until an individual stops reading, recognizing that some individuals may read the entire article. For an individual who is observed to stop reading after seeing paragraph T, this likelihood is given by decisions to continue reading after paragraphs 1,2,... ,T-1 and to stop reading after paragraph T. If Pit=pi for all t, this would be equivalent to assuming a geometric distribution (the discrete analog of the exponential distribution) for the number of paragraphs in an article that an individual reads.

Results

Emotion. We start by examining emotional valence. At first glance, the results seem to support the emotionality hypothesis (i.e., that any emotion encourages sustained attention). People are more likely to continue reading when content evokes either more positive emotion (B = *0.005,p* < .001) or more negative emotion W = 0.102,p < .001, Table 1, Model 1).

Unpacking negative emotion into different specific emotions, however, suggests the picture is more complex (Table 1, Model 2). People are more likely to continue reading when content evokes more anxiety (B = 0.518,p < .001) or anger (B = *0.375,p* < .001). They are *less* likely to continue reading, however, when content evokes more sadness (B = -0.585,p < .001).[5]

[5] Using an alternate measure of specific emotions (i.e., LIWC) shows almost identical results for both anxiety and sadness. The results for anger are largely consistent albeit weaker.

The fact some negative specific emotions increase reading while others decrease it suggests that the effect of emotions are driven by more than mere emotionality (i.e., amount of emotion) or valence. Instead, the results are more consistent with the notion that specific emotions are characterized by different appraisal or action tendencies that spillover to impact behavior. We expand on this point in more detail in the discussion of the study.

Ancillary analyses also begin to shed light on why specific emotions might influence reading. We suggested that emotions shape reading through evoking arousal and uncertainty. To begin to test these possibilities, we measured arousal using ratings from Warriner, Kuperman, and Brysbaert (2013)[6] and uncertainty using LIWC's certainty measure (i.e., words like "always" and "never"). When people are uncertain, they use more tentative language (e.g., "maybe" and "perhaps," Tausczik and Pennebaker 2010) so we measure this LIWC category as well.

Results are consistent with the notion that arousal and uncertainty increase reading. People were more likely to continue reading when content evoked greater arousal(= 0.085, $p <$ .001) or used less certain(= -0.007, $p <$ .001) or more tentative language W = 0.001, $p <$ .001). We test these points more directly in the experiments.

Processing Ease. Next, we examine processing ease. As predicted, a range of content features that should impact processing ease are linked to sustained attention (Table 1, Model 2). First, the two variables that increase complexity hurt sustained attention. People are less likely to continue reading when content is written at a higher Flesch-Kincaid Grade Level W = -0.014,$p <$

[6] People rated nearly 14,000 words based on how aroused they felt after reading them. Words like "insanity" and "lover" were rated as high arousal while words like "dull" and "librarian" were rated as low.

.001) or with greater average parse-tree height (= -0.259, $p < .001$).[7] This suggests that even beyond word and sentence length, syntactic complexity may also discourage continued reading.[8]

Second, the two variables that should increase processing ease have positive effects. People are more likely to continue reading when content uses more familiar(= 0.001,p < .001) or concrete(= 0.001,p < .001) words. This suggests that words that are easier to process encourage reading.

Control Variables. The results are robust to various controls (Table 1, Model 2). These variables are not our theoretical interest, but may be practically relevant, so we report briefly. People are less likely to continue reading if they are on a mobile device. Interpretation of topic effects, which we included as control variables, can be complex. Since topic probabilities sum to one across topics, an increase in one topic necessitates a decrease in others. Thus, topic coefficients cannot be interpreted in isolation. That said, to get a sense of which topics matter more and in which direction, we calculate the average impact of an article's topic by considering a one standard deviation increase in each topic and assuming that the remaining 24 topics are reduced by 1/24 of this amount so that the topic proportions sum to 1. This approach suggests that Topic 20 (words including "app," "googl" and "window" and seems to be associated with personal technology) and Topic 25 (words including "earth," "space" and "universe" and seems to be associated with science) are associated with some of the greatest decreases in reading.

[7] While Pitier and Nenkova (2008) found that things like words per sentence and parse tree height did not predict readability ratings, it involved only 30 articles, and may thus have been underpowered. It is also possible that readability ratings don't fully reflect actual reading behavior.

[8] Note while one could wonder whether the relationship between complexity and sustained attention might be an inverted-U, additional analyses casts doubt on this possibility. Rather than complexity simply being bad, for example, one might wonder whether some complexity is good, but too much becomes bad. But this does not seem to be the case. Including a quadratic effect of parse tree height or Flesch-Kincaid Grade Level, for example, shows that while the quadratic effects are positive (parsetree height stat, FK stat), rather than being an inverted Uthe relationship is strictly decreasing for FK and decreasing for 90% of the data for parsetree height (it only has a positive effect in the last 1% of the data). Thus complexity is almost always negatively linked to reading.

People are more likely to continue reading articles that have a higher proportion of topic 9 (which includes words such as "season," "yard" and "touchdown" and seems to be associated with football), topic 6 (including words like "food," "eat", and "drink" and seems to be associated with food) and topic 5 (including words like "movie," "film", and "character" and seems to be associated with entertainment).

Word count of the current paragraph has a negative relationship, indicating that people are less likely to continue reading after reading a longer paragraph. The number of words read so far in an article exhibits a positive relationship with continued reading, suggesting that people are more likely to continue reading longer articles. This could indicate that people tend to not to open longer pieces until times when they know they have the bandwidth to read them. Percentage read has a negative relationship, indicating that people are less likely to continue reading the further they have gotten in to an article.

Robustness checks. The results are consistent with the notion that emotions and processing ease shape reading, but one could wonder whether they are truly driven by the features identi fied. To provide further evidence, we conduct various robustness checks. We control for (1) the same textual features at the article-level, (2) heterogeneity across users, (3) other major linguistic features, and (4) textual features of the prior paragraph. We also (5) build a stock model to allow for carryover from all prior paragraphs, (6) test whether the results hold focusing on articles that have at least a certain number of readers (e.g., over 250), and (7) examine the results for just desktop. In all cases the results still hold.

First, while we attempted to address selection by controlling for site, article topic, and platform, one could still wonder whether the results are driven by different articles attracting different types of readers. To further address this concern, we include article-level feature controls. While articles that evoke certain emotions may attract certain readers for example, or attract them when they have more or less time to read, controlling for article-level features allow us to test whether even controlling for, say the amount of anxiety an article evokes, whether evoking that anxiety in a particular paragraph encourages people to read to the next one. Our results hold even using this more stringent test (Table 1, Model 3). This provides further support for the notion that the features identified have a causal impact on sustained attention.

Second, to rule out explanations that may be related to differences that exist across users, we incorporate unobserved heterogeneity into our model. Specifically, we assume that users may differ in their tendencies to continue reading paragraphs. However, given the size of our data, many methods typically used to account for continuous, unobserved heterogeneity (e.g., Ansari and Mela 2003) are infeasible as they require significant computing power. For example, when we ran a model with continuous unobserved heterogeneity using Bayesian methods, a handful of iterations took over 2 weeks to complete. Thus, we considered alternative approaches to accommodate heterogeneity and allowed for heterogeneity in the intercept via a finite number of latent classes across users (e.g., Kamakura and Russell 1989). We estimated the model with up to four latent classes, and find that our results hold.

Third, one might wonder whether other major linguistic features could explain the results. To test this, we ran a model including baskets of words empirically linked to other social or psychological constructs from LIWC dictionaries (e.g., cognitive processes, emotion, sociality, perception, motivation, time, relativity, and formality). Even after including these factors, however, the effects of our key textual features remained significant (Table 1, Model 3). The only feature that changed even slightly was parse tree height.

Fourth, we test whether the results are robust controlling for textual features of the prior paragraph (Table 1, Model 4). Results remain the same.

Fifth, we built a stock model that allows for all previous paragraphs to affect the decision to continue reading. For emotion in particular, for example, one might imagine that it builds across paragraphs, and so the total amount is important to consider. To attempt to test this, we consider the cumulative effect of linguistic features since the start of the article and assume a given linguistic feature accumulates from one paragraph to the next but that the prior amount of the stock decays. The stock model is analogous to a leaky bucket in which the previous level will decline over time, but that the level may increase with the presence of the linguistic feature in the latest paragraph. There are various formulations such a model could take, but given computational demands, we used a version where the decay parameters for key predictors are kept the same. The stock model is formulated as follows:

$$y_{ij} \sim Bernoulli(P_{ij}) \text{ where}$$

$$logit(P_{ij}) = \beta_0 + \sum_k \beta_k \; Stock_{ijk} + \sum_c \gamma_c \; z_{ijc}$$

$$Stock_{ijk} = \delta \cdot Stock_{i(j-1)k} + X_{ijk}$$

Given computational demands, this model could only be run on a 20% sample of the articles, but the results are generally the same. Fear ($\beta = 0.05$, $p < .001$), positive emotion ($\beta = 0.0002$, $p = .36$), and anger ($\beta = 0.10$, $p < .001$) all remain positively associated with continued reading, though the effect of positive emotions is no longer significant. Sadness remains negatively associated with continued reading ($\beta = -0.10$, $p < .001$).

Sixth, one could wonder whether the results are somehow driven by some of the articles having very few readers. But this is not the case. Even looking at articles with over 250 reads (Table 1, Model 5, 1.4% of all articles, 48% of all reading sessions), the results remain the same.

Finally, one could be concerned that the way users interact with different devices might be driving the results. Mobile has a smaller screen, for example, which likely impacts reading. While we controlled for device in our prior analyses, to further address this point, we also restrict our analysis to articles reads on just desktops. We find the same effects (Table A5).

Discussion

Natural language processing of over 700,000 page-read events from over 35,000 articles suggests how textual features of content shape reading. First, the results suggest that evoking emotion shapes reading, but that the effects are more complex that simple emotionality or valence alone. Consistent with a specific emotion based perspective, different negative emotions had different relationships with reading. While readers were more likely to continue reading after paragraphs that evoked more anxiety or anger, they were *less* likely to continue reading after paragraphs that evoked more sadness. Second, the results suggest that people are more likely to continue reading after passages that are easier to process.

Second, the findings suggest that while the topic of the article (e.g., sports vs. politics) certainly shapes reading, textual features are also of importance. Figure 1 illustrates the fitted probability of reading the entire article based on a one standard deviation increase in each of the main textual features for each paragraph. For example, increasing the amount of anxiety each paragraph evokes by one standard deviation increases the odds that readers will finish the entire article by more than 15%. Similarly, a one-standard-deviation decease in Flesch-Kincaid score increases the odds that readers will finish the article by over 20%.

In an effort to illustrate the how an article's content (i.e., its topics) versus its writing style (i.e., processing ease, emotionality and ambiguity) affects an individual's reading behavior, we compare the effects of each textual feature to the average impact of each topic in a particular scenario (Figure 1). We calculate the average impact of an article's topic by considering a one standard deviation increase in each topic (assuming the remaining 24 topics are reduced by 1/24 of this amount) and averaging the net absolute[9] effects across all 25 topics. This analysis depicts a point of comparison for the effect of textual features, if they each increased by one standard deviation. As Figure 1 illustrates, the effects of emotions, ambiguity and processing ease can be similar to (or even exceed) the effect of topics on continued reading.

FIGURE 1: LIKELIHOOD OF READING ENTIRE ARTICLE BASED ON EACH TEXTUAL FEATURE

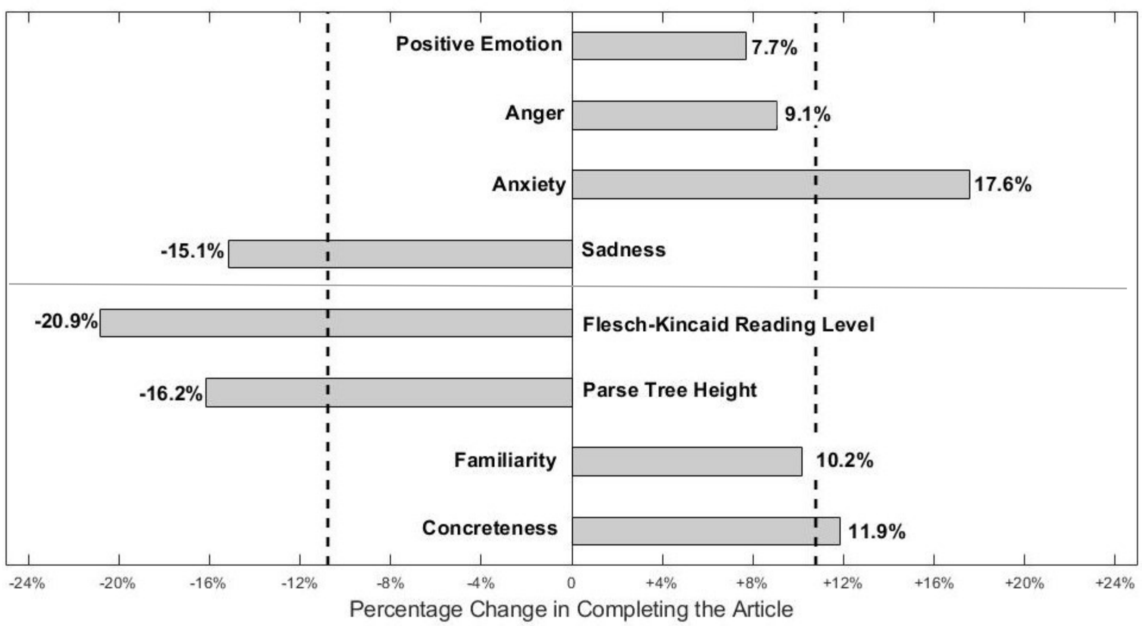

Note: Each bar represents the effect of a one standard deviation increase in the specified textual feature on continued reading, relative to a "baseline" article (defined as an article in the Wall Street Journal read on a desktop computer with average posterior probabilities assumed for each topic, and average emotion and control measures). Given our interest in the magnitude of topic effects (rather than direction), we average the absolute effects as some topics increase continued reading while others decrease it and illustrate both the positive and negative impact in the figure.

[9] Given our interest in the magnitude of topic effects (rather than direction), we average the absolute effects as some topics increase continued reading while others decrease it.

Ancillary Analyses. One might wonder whether surprise plays a role. We did not include this variable in our main analysis because there is no measure for it in either emotion dictionary we, and it is not frequently examined in papers on specific emotion. That said, some prior research suggests that surprise is low in certainty (Lerner, Li, Valdesolo, and Kassam, 2015; Tiedens and Linton 2001) and high in arousal, consequently, if our theorizing about the role of these two factors is correct, surprise should be associated with sustained attention. To test this possibility, we used the NRC Lexicon (Mohammed and Tumey, 2013) to measure surprise and included as a predictor variable in the model. As expected, people are more likely to continue reading when content is more surprising($= 0.009, p < .001$). Though exploratory, these results are consistent with our theorizing about how arousal and certainty might shape reading.

STUDY 2: MANIPULATING SPECIFIC EMOTIONS -ANGER

Studies 2-4 have three main goals. First, while the results of Study 1 are supportive, as is often the case with observational data, it is difficult to definitely demonstrate causality. Including various controls casts doubt on many alternative explanations, but one could still wonder whether the features identified truly influenced reading. Consequently, to more directly test the causal nature of the effects, we conduct experiments. Simultaneously testing all the different features would be challenging, so we focus on the three specific emotions. We manipulate how much anger (Study 2), anxiety (Study 3), and sadness (Study 4) an article evokes and measure the resulting impact on reading. Consistent with the field data, we predict that anxiety and anger will both increase reading and sadness will decrease it.

Second, the studies test the underlying process. We suggested that specific emotions may influence reading through arousal and uncertainty. Anxiety, for example, tends to increase arousal (which should encourage reading) and make people feel uncertain (which should also encourage reading). If this theorizing is correct, then these underlying dimensions should mediate any of specific emotions' effects. We measure both arousal and uncertainty and test whether they can explain the results.

This is particularly important because looking at Study 1 in isolation, one could wonder whether arousal alone is sufficient to explain the results. After all, while anger (high arousal) and anxiety (high arousal) were associated with increased reading, sadness (low arousal) was associated with decreased reading. Thus, arousal seems sufficient to explain the results. Further, on the surface, the results seem less consistent with an uncertainty-based explanation given that some emotions characterized by uncertainty (i.e., anxiety) increased reading while others characterized by uncertainty (i.e., sadness) decreased reading. Emotions differ on multiple dimensions simultaneously, however, and without knowing how specific articles effect specific readers, it is hard to fully test the underlying mechanisms. Further, while one could argue that specific emotions impact though arousal might overwhelm their effects through certainty, anxiety should have positive effects on reading through both arousal and uncertainty, making it difficult to know which is truly having the larger effect. The experiments directly tests whether uncertainty effects reading above and beyond any effects of arousal.

Third, one could wonder whether measuring emotions through text analysis truly captures whether those emotions were evoked in article readers. A great deal of psychological research demonstrates that reading words can evoke emotion (see Quigley, Lindquist, and Feldman Barrett 2014 for a review) and prior work in marketing (e.g., Berger and Milkman 2012) uses text analysis to measure the emotions evoked by content. Further, marketing researchers (e.g., Coleman, Williams, Morales, and White 2017) often use reading a passage as an emotion induction (e.g., reading a story about an intruder to introduce fear). That said, to underscore that people are actually experiencing the emotions in question, we measure how much participants felt key emotions as a manipulation check.

Method

Participants ($N = 179$), recruited from an East Coast behavioral lab, completed a "News Article Study" experiment as part of a larger group of experiments. They were told experimenters were interested in perceptions of the news and that they would read the beginning of a news article and respond to some questions.

First, we manipulated how much anger the initial part of the article induced. All participants read the beginning of an article, adapted from an article in Study 1, about how women in El Salvador have been arrested for terminating their pregnancies (see Appendix for full text). In the low anger condition, the article talked about how women's right advocates were trying to fix the problem and bills had been proposed in the legislature. In the high anger condition, the article used more anger inducing language (e.g., talked about how women had been "imprisoned" rather than "arrested") and focused on how the situation was unfair (i.e., that the women were "convicted based on a 'test' that has been ruled unsound for over a century"). Pre-test data ($N = 99$) confirmed that the anger manipulation worked. Participants read one of the two versions of the article and reported how angry they felt. As expected, the high anger version of the article induced more anger ($M = 5.04$ vs. 3.73, $t(97) = 3.21$, $p = .002$).

Second, we measured the key dependent variable (i.e., "how much do you want to read the rest of the article," 1 = Not at All, 7 = Very Much).

Third, we measured the hypothesized underlying processes, arousal and uncertainty, using measures from prior work. Uncertainty was measured using three items from Faraji-Rad and Pham (2017) on 7-point scales: "how do you feel about your environment?" anchored by unsure/sure, hesitant/determined, and don't feel confident/feel confident (a= .94, reverse-scored and averaged to an uncertainty index). Arousal was measured using three items from Berger (2011) on 7-point scales: "how do you feel right now" anchored by very low energy/very high energy, very passive/very active, very mellow/very fired up (a= .90, averaged to arousal index).

Results

As predicted, making the initial part of the article more anger inducing increased participants' interest in reading further ($M = 4.46$ vs. 3.64), $t(177) = 2.99$, $p = .003$).

Further, as expected, making the content more anger inducing increased arousal ($M = 4.80$ vs. 4.18), $t(177) = 3.06$, $p = .003$) but decreased uncertainty ($M = 1.76$ vs. 2.27), $t(177) = 2.24$, $p = .026$).

Finally, as predicted, a bias corrected simultaneous mediation model (Hayes 2017) found that both arousal [*ab*= .20, 95% CI .06 to .45] and uncertainty [*ab*= -.12, 95% CI -.35 to -.01] drove the effects of anger on reading. Anger increased arousal which increased interest in reading (coeff = .33, se = .10, t = 3.11, p = .002). Anger also decreased uncertainty, but because uncertainty has a positive effect interest in reading (coeff = .23, se = .09, t = 2.47, p = .014), the indirect effect is negative (see Figure A2 for a mediation diagram).

Discussion

Results of Study 2 bolster the findings of the field data in a controlled setting and provide evidence for the hypothesized underlying process. First, taking the same article, and manipulating the beginning to make it more anger inducing, led people to be more interested in reading further. Second, as predicted, these effects were driven by arousal and uncertainty. When the article was more anger inducing it boosted arousal but decreased uncertainty, each of which shaped interest in reading further.

These findings have clear implications for content creators. Even when the topic of an article is fixed, changing the text of the article to evoke more on certain emotions, in this case anger, can boost sustained attention.

STUDY 3: MANIPULATING ANXIETY

Building on the findings of study 2, study 3 manipulated anxiety. We again measure the impact on interest in continuing to read, and whether any observed effects are driven by arousal and uncertainty. Further, to provide tight experimental control, and demonstrate how subtle changes in wordings can impact sustained attention, in Study 3 we changed only a few key words of the beginning of the article.

Method

The experiment followed a similar setup to Study 2, but we used different content and manipulated anxiety.

Participants (N = 145) completed a similar "News Article Study" where we manipulated how much anxiety the initial part of the article induced. All participants read the beginning of an article about data breaches, but to provide greater control, the stimulus was only one sentence long. In the low anxiety condition, the article noted that "Every year, some people lose data to data breaches, according to credit bureau TransUnion." In the high anxiety condition, the article used more anxiety inducing language "Every minute in the U.S., 19 people fall victim to identity theft, according to credit bureau TransUnion." Rather than talking about "data breaches", it uses the more specific and anxiety inducing version ("identity theft"). Rather than talking in general about how some people are affected every year, it noted that every minute 19 people are affected. Pre-test data (N = 99) confirmed that the anxiety manipulation worked. Participants read one of the two versions of the article and reported how anxious they felt. As expected, the high anxiety version of the article induced more anxiety (M = 4.08 vs. 3.02, t(97) = 2.54, p = .013)

Second, we measured the key dependent variable (i.e., "how much do you want to read the rest of the article," 1 = Not at All, 7 = Very Much).

Third, we measured the hypothesized underlying processes, arousal and uncertainty, using the same measures from Study 2.

Results

As predicted, making the initial part of the article more anxiety inducing increased participants' interest in reading more (M = 4.20 vs. 3.30), t (143) = 2.56, p = .012). Further, as expected, making the content more anxiety inducing increased both arousal (M = 4.38 vs. 3.89), t (143) = 2.13, p = .035) and uncertainty (M = 2.66 vs. 1.96), t (143) = 2.59, p = .011).

Finally, as predicted, a bias corrected simultaneous mediation model (Hayes 2017) found that both arousal *[ab*= .20, 95% CI .03 to .49] and uncertainty *[ab*= .16, 95% CI .01 to .44] drove the effects of anxiety on reading. Anxiety increased both arousal and uncertainty and each increased interest in reading more (arousal coeff= .36, se = .12, t = 3.06, p = .003; uncertainty coeff= .23, se = .11, t = 2.04, p = .042). See Figure A3 for a mediation diagram.

Discussion

Results of Study 3 further bolster the findings of the field data in a controlled setting and provide evidence for the hypothesized underlying process. First, taking the same article, and manipulating the beginning to make it more anxiety inducing, led people to be more interested in reading further. Second, as predicted, these effects were driven by arousal and uncertainty. When the article was more anxiety inducing it boosted arousal and uncertainty, each of which encouraged interest in further reading.

STUDY 4: MANIPULATING SADNESS

While the results of Study 2 and 3 are consistent with our theorizing, one could argue that they simply demonstrate that increasing the amount of *any* emotion increases interest in reading further. To test this possibility, in study 4 we manipulate sadness. We predict that making the beginning of an article more sadness inducing will make people *less* interested in continuing to read. Further, consistent with study 2 and 3, we predict that these effects will be driven by arousal and uncertainty.

Method

Study 4 followed a similar setup to study 2, expect that we manipulated sadness. Participants (N = 136) completed a similar "News Article Study" where we manipulated how much sadness the initial part of the article induced. All participants read the beginning of an article about a local diner. In the low sadness condition, the article noted that after 68 years a local diner might be closing "Last call? After 68 years of serving up hearty omelets and mouthwatering burgers, the historic Salem Diner may have close. "It's heartbreaking," said one customer "I've been coming every month with my Dad since I was a kid. Losing such an iconic, local restaurant would be a big loss for the community." In the high sadness condition, the article was almost exactly the same except that it was made more sadness inducing because the diner was definitely closing "It's finally the last call. After 68 years of serving up hearty omelets and mouthwatering burgers, the historic Salem Diner is closing. "It's heartbreaking," said one customer "I've been coming every month with my Dad since I was a kid. Losing such an iconic, local restaurant is a big loss for the community. A decrease in customers and rising food costs meant the economics just didn't work." Pre-test data (N = 99) confirmed that the sadness manipulation worked. Participants read one of the two versions of the article and reported how sad they felt. As expected, the high sadness version of the article induced more sadness (M = 4.67 vs. 2.74, t(97) = 2.49, p = .015)

Second, we measured the key dependent variable. To ensure that the results of the prior two studies were not somehow driven by the wording of the dependent variable, participants were asked "how interested are you in reading the rest of the article," 1 = Not at All, 7 = Very Much).

Third, we measured the hypothesized underlying processes, arousal and uncertainty, using the same measures from Study 2.

Results

As predicted, and in contrast to Study 3a and 3b, making the initial part of the article more sadness inducing made participants *less* interested in continuing to read (M = 3.10 vs. 4.25), t (134) = -3.47, p = .001).

Further, as expected, making the content sadder decreased both arousal (M = 3.70 vs. 4.30), t (134) = 2.43, p = .016) and uncertainty (M = 1.69 vs. 2.25), t (134) = 2.16, p = .032).

Finally, as predicted, a bias corrected simultaneous mediation model (Hayes 2017) found that both arousal *[ab*= -.22, 95% CI -.56 to -.04] and uncertainty *[ab*= -.13, 95% CI -.42 to -.01] drove the effects of sadness on reading. As shown in Figure 4, sadness decreased both arousal and uncertainty and each increased interest in reading (arousal coeff= .37, se = .12, t = 3.05, p = .003; uncertainty coeff= .25, se = .12, t = 2.11, p = .036).

Discussion

Results of Study 4 underscore the findings of the field data are Studies 2 and 3. First, taking the same article, and manipulating the beginning to make it induce more sadness, led people to be less interested in reading further. Second, as predicted, these effects were driven by arousal and uncertainty. When the article induced more sadness, it reduced arousal and uncertainty, which in tum drove interest in reading further.

Additional Study. While the results of Study 2-4 support our theorizing, one could argue that in addition to manipulating the focal emotion, they may have also manipulated other textual features and these features, rather than the emotion in question is what drove the observed effect. Though such alternatives have difficulty explaining the mediation results, to provide an even cleaner test, we conducted a follow-up experiment (Study 5, Appendix) where we manipulated emotion exogenously. We kept the focal content itself the same, but by having participants write about a time they felt either sad, angry, or anxious, depending on condition, we manipulated emotion incidentally (see Berger 2011 for similar design). We find the same results.

GENERAL DISCUSSION

Online content has become a key way to engage audiences. Media companies, like newspapers and magazines, depend on readers to attract advertisers, and through content marketing, a much broader set of companies and individuals now use content as a strategy to engage an audience, grow awareness, and change attitudes.

But all of this depends on people actually consuming or reading the content. Why do some articles generate sustained attention while others don't?

This paper attempts to provide at least a preliminary answer. Analyzing over 700,000 reading sessions encompassing over 35,000 articles from nine major online publishers, we use natural language processing to shed light on how textual features might shape continued reading. Results suggest that both emotion and processing ease shape sustained attention.

First, rather than any emotion increasing reading, or reading being driven by valence alone, the results suggest a more complex picture. Anger, anxiety, and sadness are all negative emotions, but while evoking anxiety and anger encouraged reading, evoking sadness discouraged it. Ancillary analyses (Study 1) and mediational analyses (Study 2-4) highlight the important role that arousal and uncertainty play in these effects. Specific emotions shape how uncertain and physiologically aroused people feel, which in tum, shape reading

Second, various features that should shape processing ease seem to influence reading. People are less likely to continue reading content that is more complex (i.e., high Flesch-Kincaid reading level and parse tree height) and more likely to continue reading after content that uses more familiar and concrete language.

Demonstrating these findings across thousands of articles from a variety of different content providers speaks to their generalizability. Experimental evidence underscores the causal nature of the effects observed in the field and provides evidence for the underlying processes.

Contributions

These findings make several contributions. First, on a theoretical level, they provide insight into what drives reading. While some work has looked at how visual features shape attention (e.g., Zillmann, et al., 2001), or explored what characteristics of visual and auditory stimuli drive attention (e.g., whether people fixate more on novel images, Berlyne 1971; 1974), there has been less attention to how textual features shape sustained attention. This work begins to shed light on this issue, demonstrating the important role of emotion and processing ease. It is hoped that these preliminary findings will encourage further work in the area.

An analogy can be made to burgeoning interest in behavioral drivers of word of mouth. People have been sharing things for thousands of years, and many papers have demonstrated the causal impact of word of mouth, but until recently, *why* people share word of mouth in the first place was less clear. The same can be said of continued reading or sustained attention. Reading is not new. But the availability of better data on how people are reading provides a new opportunity to study behavioral drivers of this important dependent variable in greater detail.

Second, on a more applied level, the results have clear practical takeaways for designing content that generates sustained attention. For journalists, media outlets, or marketers trying to get people to pay sustained attention to their content, these findings suggest that certain linguistic features may be useful. Using more concrete language and familiar words, for example, or replacing abstract words with more concrete ones, and less familiar words with more familiar synonyms. Similarly, while we are not advocating making the reader anxious unnecessarily, particularly for topics that may already evoke some anger and anxiety, writing in ways that allow those emotions to come through should encourage sustained attention. Even evoking uncertainty more generally may be useful. As shown in Studies 2, 3, and 4, simple shifts in language can have important impacts on interest in reading further.

Further, the results demonstrate that reading depends on more than topic alone. Organizations often lament that it is easier to get people to care about some topics than others. A common refrain is that while people love reading about sports and celebrity gossip, it's harder to get them to pay attention to heavier topics like policy discussions and environmental appeals. But while topic certainly plays a role in driving sustained attention, our results suggest that they are not the only factor. Even controlling for what an article is about (i.e., its topic or topics), how that topic is discussed plays an important role in whether people continue reading.

This provides a hopeful note for organizations trying to generate sustained attention for less "engaging" topics. While the topic itself may not engender continued reading, writing about it in a way that generates uncertain emotions and processing ease should deepen sustained attention. Writing style can compensate for topic.

Future Research

As with any preliminary effort, much more remains to be done. While we focused on continued reading, or sustained attention, several other related dependent variables might be interesting to examine. Thought it would be tough to do at scale, subsequent work might use eye-tracking measures to get a more fine-grained sense of how textual features shape attention. Similarly, work could examine how different textual features shape memory or comprehension.

Future work could also examine *when* different textual features might have greater impact. Might the effects of concreteness, for example, vary for different types of content? To begin to examine this question, we took the topic modeling results and used it to cluster articles into different types of content (e.g., politics, sports, and lifestyle). Concreteness, however, did not seem to have different effects on sustained attention across the groupings. Future work might look at other textual features though, as well as different types of articles within content groups. Maybe concreteness is more beneficial for certain types of political articles than others, for example. Alternatively, work could examine how the effects of different textual features vary by time of day. Prior work suggests that people may have more resources at certain times of day, and structurally, likely have more time to read after work than during. One could also make predictions about certain emotions having greater impact at different times of day.

Another interesting question is how where readers came from previously might impact behavior. While our data does not allow us to examine this in detail, how might coming in from social media versus a media outlet's page impact sustained attention? Similarly, does reading one article effect reading behavior on a subsequent article? If people leave one article because it is not engaging enough, for example, might that increase their impatience or sustained attention threshold, such that it decreases their likelihood of sustained attention in a subsequent article? If so, that would suggest that media outlets may want to feature content that generates longer reads to encourage people to more deeply engage with whatever they read next.

Similarity or atypicality would also be an interesting feature to examine. How does the similarity between an article and its outlet shape reading? Take a media outlet, for example, that focuses primarily on sports. Are people more likely to deeply read a typical (i.e., sports) article, or one that is more atypical (e.g., about player's personal lives)?

Similarity could also be examined within content. Any story, article or long form narrative (e.g., book or movie) can be broken down into small chunks. How might the similarity between those chunks impact sustained attention? One could argue that similarity is beneficial. Content that flows between similar chunks should be easier to process and follow, which might deepen sustained attention. At the same time, however, chunks being too similar might feel repetitive or like the plot is not advancing fast enough.

This brings up the broader question of narrative development. Emotional trajectories might be one area to consider. While we focused how emotion in one paragraph related to reading the next, across an article or longer narrative, emotion can take various trajectories. These trajectories, in tum, may shape things like narrative transportation (Busselle and Bilandzic 2009; Green and Brock 2000; Escalas 2004; 2006 van Laer, Escalas, Ludwig, and van den Rende 2018), which could encourage sustained attention. Great stories can suck people in, transporting them into the narrative world and helping them see things through the eyes of key characters. Work in this area has demonstrated key dimensions of narrative engagement (e.g., emotional engagement, Busselle and Bilandzic 2009), shown how narrative processing can enhance self-brand connection (Escalas 2004), and demonstrated how narrative processing interacts with processing ease to impact brand evaluations. But while there has been a great deal of theorizing about narrative structure and what makes a good story (e.g., Bruner 1986; see Escalas 1998 for a review), there has been less empirical attention to what textual features actually drive narrative transportation (though see Laer, Escalas, Ludwig, and Van den Rende 2019). Further, while there is some suggestion that lack of mind wandering is a dimension of narrative engagement (Busselle and Bilandzic 2009), it remains to be seen whether this is the same as sustained attention in terms of reading.

Future work might examine the connection between sustained attention and narrative transportation. Some of the articles examined here certainly seem like stories (i.e., narratives that have a beginning, middle and end and relationship between key elements, Escalas 1998), but other are more about conveying information (e.g., product reviews). Consequently, not all may engender narrative transportation. Further, recent work (Pennebaker, Chung, Frazee, Lavergne, and Beaver 2014) suggest a text-based measure of analytic versus more narrative thinking that may help automatically measure some of the key aspects.

Research might also examine how sustained attention affects return visits. One would imagine that the more attention readers give an article, the more likely they will be to return to that content provider in the future. If most articles only hold people's attention for a couple paragraphs, they're less likely to keep coming back. This highlights the downsides of overly attention-grabbing headlines. Clickbait may be great for attracting views, but to maintain long term brand value, deeper sustained attention with the content itself may be needed.

It would also be interesting to examine whether easy to process language has positive effects in other domains. Take customer service calls, for example. Might customer service agents been seen as more helpful if they use concrete language or more familiar words? Such actions might make the interaction feel easier which might increase customer satisfaction.

In conclusion, while these results shed light on the behavioral drivers of continued reading, they also highlight the value of using natural language processing to extract behavioral insight from textual data. From articles, social media posts, and transcribed customer service calls, to movie scripts, song lyrics, and millions of digitized books, technology has made more and more textual data accessible. With the right tools, these data can provide insight into a range of behaviorally interesting and managerially relevant questions.

TABLE 1: READING AS FUNCTION OF CONTENT CHARTACTERISTICS

		Valence (I)	Specific Emotions (2)	Additional Features (3)	Last Paragraph (4)	250+ Reads (5)
Emotion	Positive Emotion	0.005***	0.006***	0.008***	0.006***	0.008***
		(0.000)	(0.000)	(0.000)	(0.000)	(0.001)
	Negative Emotion	0.003***				
		(0.000)				
	Anger		0.375***	0.228***	0.372***	0.315***
			(0.025)	(0.028)	(0.025)	(0.037)
	Anxiety		0.518***	0.254***	0.463***	0.872***
			(0.025)	(0.028)	(0.025)	(0.037)
	Sadness		-0.585***	-0.326***	-0.539***	-1.240***
			(0.027)	(0.030)	(0.027)	(0.039)
Processing Ease	FK	-0.013***	-0.014***	-0.020***	-0.012***	-0.014***
		(0.000)	(0.000)	(0.000)	(0.000)	(0.001)
	Parse Tree Height	-0.025***	-0.026***	-0.000	-0.024***	-0.014**
		(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
	Familiarity	0.001***	0.001***	0.00 I***	0.001***	0.003***
		(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
	Concreteness	0.005***	0.001 ***	0.002***	0.001 ***	0.008***
		(0.000)	(0.000)	(0.000)	(0.000)	(0.001)
Controls	Platform=Desktop	0.027	0.027	0.030	0.028	0.023
		(0.019)	(0.019)	(0.019)	(0.019)	(0.030)
	Platform=Mobile	-0.141***	-0.142***	-0.144***	-0.141***	-0.191***
		(0.019)	(0.019)	(0.019)	(0.019)	(0.030)
	Platform=Tablet	0.019	0.018	0.021	0.018	0.021
		(0.019)	(0.019)	(0.020)	(0.019)	(0.030)
	Para Word Count	-0.010***	-0.010***	-0.010***	-0.011***	-0.011***
		(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
	% Article Read	-3.099***	-3.116***	-3.235***	-3.721***	-3.560***
		(0.020)	(0.020)	(0.020)	(0.023)	(0.029)
	squared	-2.492***	-2.481***	-2.456***	-1.980***	-2.384***
		(0.020)	(0.020)	(0.020)	(0.022)	(0.029)
	Amount Read	0.004***	0.004***	0.004***	0.004***	0.004***
		(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
	squared	0.000***	-0.000***	-0.000***	0.000***	-0.000***
		(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
	Site dummies	Yes	Yes	Yes	Yes	Yes
	Topic Controls	Yes	Yes	Yes	Yes	Yes
	Features@article level	No	No	Yes	No	No
	Additional features	No	No	Yes	No	No
	Last Paragraph	No	No	No	Yes	No
	Observations	6,994,372	6,994,372	6,994,372	6,994,372	3,644,358

***p < .001, ** p < .01, *p < .05, "p <.10

REFERENCES

Baumeister, Roy F., Bratslavsky, Ellen, Finkenauer, Catrin, and Kathleen D. Vohs (2001), "Bad is Stronger than Good," *Review of General Psychology,* 5 (4), 323.

Berger, Jonah (2011), "Arousal Increases Social Transmission of Information," *Psychological Science,* 22 (7), 891-93.

Berger, Jonah (2014), "Word of mouth and interpersonal communication: A review and directions for future research", *Journal of Consumer Psychology,* 24 586-607. 10.1016/j.jcps.2014.05.002.

Berger, Jonah and Katherine Milkman (2012), "What Makes Online Content Viral?" *Journal of Marketing ,* 49 (2), 192-205.

Berger, Jonah, and Grant Packard (2018), "Are Atypical Things More Popular?" *Psychological Science,* 29(7), 1178-1184.

Berlyne, D. E. (1971), "Aesthetics and psychobiology", East Norwalk, CT, US: Appleton-Century-Crofts.

Berlyne, D. E. (1974), "Studies in the new experimental aesthetics: Steps toward an objective psychology of aesthetic appreciation", Oxford, England: Hemisphere.

Blei, David M., Ng, Andrew Y., and Michael I. Jordan (2003), "Latent Dirichlet Allocation," *Journal of Machine Learning Research,* 3 (Jan), 993-1022.

Bonezzi, A., Brendl, C. M., and M. De Angelis (2011), "Stuck in the Middle: The Psychophysics of Goal Pursuit", *Psychological Science,* 22(5), 607-612.

Booker, Christopher (2004), *The Seven Basic Plots: Why We Tell Stories.* London: A&C Black.

Brooks, Alison Wood, and Maurice E. Schweitzer (2011), "Can Nervous Nelly Negotiate? How Anxiety Causes Negotiators to Make Low First Offers, Exit Early, and Earn Less Profit," *Organizational Behavior and Human Decision Processes,* 115 (1), 43-54.

Bruner, Edward M., "Experience and its expressions," *The anthropology of experience* 3(1986): [redacted]

Busselle, Rick and Helena Bilandzic (2009), "Measuring Narrative Engagement", *Media Psychology,* 12 (4), 321-347.

Calder, Bobby J., Malthouse, Edward C., and Ute Schaedel (2009), "An Experimental Study of The Relationship Between Online engagement and Advertising Effectiveness," *Journal of Interactive Marketing,* 23 (4), 321-331.

Cavanaugh, Lisa A., Bettman, James R., and Mary Frances Luce (2015), "Feeling Love and Doing More for Distant Others: Specific Positive Emotions Differentially Affect Prosocial Consumption," *Journal of Marketing Research,* 52 (5), 657-673.

Chen, Zoey and Nicholas H. Lurie (2013), "Temporal Contiguity and Negativity Bias in the Impact of Online Word-of-Mouth," Journal of Marketing Research, 50 (4), 463-76.

Chen, Zoey and Andrew Stephen (2018), "Digital and Social Media Research," In Frank R. Kardes, Paul M. Herr, and Norbert Schwarz (eds). *Handbook of Research Methods in Consumer Psychology.* Routledge.

Chomsky, Noam (1957), *Syntactic Structures.* The Hague/Paris: Mouton.

Coleman, Meri and Ta Lin Liau (1975), "A Computer Readability Formula Designed for Machine Scoring," *Journal of Applied Psychology,* 60 (2), 283.

Cole, Jeffery, Suman, Michael, Schramm, Phoebe, and Liuning Zhou (2017), *Surveying the Digital Future 2017,* Los Angeles, Calif.: Center for the Digital Future at USC Annenberg.

Coltheart, Max (1981), "The MRC Psycholinguistic Database," *The Quarterly Journal of Experimental Psychology,* 33 (4), 497-505.

Connell, Louise and Dermot Lynott (2012), "Strength of Perceptual Experience Predicts Word Processing Performance Better than Concreteness or Imageability," *Cognition,* 125 (3), 452-465.

Easterbrook, James A. (1959), "The Effect of Emotion on Cue Utilization and The Organization of Behavior," *Psychological Review,* 66 (3), 183.

Escalas, Jennifer Edson (1998), "Advertising Narratives: What Are They and How Do They Work?" in *Representing Consumers: Voices, Views, and Visions,* Barbara Stem, ed., London: Routledge, 267-289.

Escalas, Jennifer Edson (2004), "Imagine Yourself in the Product: Mental Simulation , Narrative Transportation, and Persuasion", *Journal of Advertising,* 33 (2), 37-48.

Escalas, Jennifer Edson (2006), "Self-Referencing and Persuasion: Narrative Transportation versus Analytical Elaboration", *Journal of Consumer Research,* Volume 33, Issue 4, March 2007, 421-4 29.

Faraji-Rad, Ali and Michel T. Pham (2017), "Uncertainty Increases the Reliance on Affect in Decisions," *Journal of Consumer Research,* 44 (1), 1-21.

Fiske, Susan T. (1980), "Attention and Weight in Person Perception: The Impact of Negative and Extreme Behavior," *Journal of Personality and Social Psychology,* 38 (6), 889.

Fogg, B. J. (2009), "A behavior model for persuasive design", In *Proceedings of the 4th international Conference on Persuasive Technology* (p. 40). ACM

Friederici, Angela D., Opitz, B., and D. Yves Von Cramon (2000), "Segregating semantic and syntactic aspects of processing in the human brain: an fMRI investigation of different word types," *Cerebral cortex, 10(7),* 698-705.

Garcia, Mario R., Stark, Margaret Mary, and Ed Miller (1991), *Eyes on the News*. St. Petersburg, FL: The Poynter Institute.

Gergen, K. J., and M. M. Gergen (1988), "Narrative and the self as relationship", *Advances in experimental social psychology, Vol. 21. Social psychological studies of the self: Perspectives and programs* (pp. 17-56). San Diego, CA, US: Academic Press.

Gilhooly, Ken J. and Robert H. Logie (1980), "Age-of-Acquisition, Imagery, Concreteness, Familiarity, and Ambiguity Measures for 1,944 Words," *Behavior Research Methods & Instrumentation,* 12 (4), 395-427.

Green, Melanie C., and Timothy C. Brock (2000), "The Role of Transportation in the Persuasiveness of Public Narratives," *Journal of Personality and Social Psychology,* 79 (5), 701-721.

Greesonbach, Sarah (2018), "Slow Your Content Marketing Down," New York, Content Marketing Institute.

Han, S., Lerner, J. S., and D. Keltner (2007), "Feelings and consumer decision making: The appraisal-tendency framework", *Journal of consumer psychology, 17(3),* 158-168.

Harlam, B. A., and L.M. Lodish (1995), "Modeling Consumers' Choices of Multiple Items", *Journal of Marketing Research,* 32(4), 404-418.

Heilman, K. H. (1997), "The Neurobiology of Emotional Experience," *The Neuropsychiatry of Limbic and Subcortical Disorders,* 133-142.

Hidi, Suzanne (1990, "Interest and Its Contribution as a Mental Resource for Leaming," *Review of Educational Research,* 60(4), 549-571.

Hull, Clark L. (1932), "The Goal-Gradient Hypothesis and Maze Leaming," *Psychological Review,* 39 (1), 25.

Humphreys, Ashlee and Rebecca Jen-Hui Wang (2017), "Automated Text Analysis for Consumer Research," *Journal of Consumer Research,* 44 (6), 1274-1306.

Jacoby, Jacob and Wayne D. Hoyer (1987), *The Comprehension and Miscomprehension of Print Communications ,* Advertising Educational Foundation, Lawrence Earlbaum Associates, New Jersey .

Jeffries, Adrianne (2014), "You're Not Going to Read This, But You'll Probably Share It Anyway," *The Verge ,* (February 14), [available at https://www.theverge.com/2014/2/14/5411934 /youre-not-going-to-read-this].

Just, Marcel A. and Patricia A. Carpenter (1980), "A Theory of Reading: From Eye Fixations to Comprehension," *Psychological Review,* 87 (4), 329.

Kamakura, Wagner A., and Gary J. Russell (1989), "A Probabilistic Choice Model for Market Segmentation and Elasticity Structure", *Journal of Marketing Research, 26(4),* 379-390 . doi:10.2307/3172759

Keltner, Dacher and Jennifer S. Lerner (2010), "Emotion," *in The Handbook of Social Psychology,* 5th ed., D. Gilbert, S. Fiske, and G. Lindsey, eds. New York: McGraw-Hill.

Kincaid, J. Peter, Fishburne Jr., Robert P., Rogers, Richard L., and Brad S. Chissom (1975), "Derivation of New Readability Formulas (Automated Readability Index, Fog Count and Flesch Reading Ease Formula) for Navy Enlisted Personnel," *Institute for Simulation and Training,* 56.

Kilger, Max and Ellen Romer (2007), "Do Measures of Media engagement Correlate with Product Purchase Likelihood?" *Journal of Advertising Research,* 47 (3), 313-325.

Lagun, Dmitry, and Eugene Agichtein (2015), "Inferring Searcher Attention by Jointly Modeling User Interactions and Content Salience." *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval,* Santiago, (August 9-13), 483-492.

Leckner, Sara (2012), "Presentation Factors Affecting Reading Behaviour in Readers of Newspaper Media: An Eye-tracking Perspective," *Visual Communication,* 11 (2), 163-184.

Lerner, Jennifer S., and Dacher Keltner (2001), "Fear, Anger, and Risk," *Journal of Personality and Social Psychology,* 81 (1), 146.

Mather, Mara and Matthew Sutherland (2012), "The Selective Effects of Emotional Arousal on Memory," *Psychological Science Agenda,* 26, 2.

Mather, Mara (2007), "Emotional Arousal and Memory Binding: An Object-Based Framework," *Perspectives on Psychological Science,* 2 (1), 33-52.

McCoy, Julia (2017), "The Big 2017 Content Marketing Spend: Content Budgets by the Numbers, & How to Set Your CM Budget (Infographic)," *Express Writers,* (January 26), [available at https://expresswriters.com /the-big-2017-content-marketing-spend-infographic/].

McLaughlin, G. Harry (1969), "Clearing the SMOG," *Journal of Reading,* 639-645 .

McKee, Robert (1998), *Story: Substance, Structure, Style, and the Principles of Screenwriting,* HarperCollins, New York.

Moe, Wendy W. (2006), "A Field Experiment to Assess the Interruption Effect of Pop-Up Promotions," *Journal of Interactive Marketing ,* 20 (1), 34-44.

Mohammad, Saif, M. and Felipe Bravo-Marquez (2017), "Emotion Intensities in Tweets," *Proceedings of the Sixth Joint Conference on Lexical and Computational Semantics, Vancouver, (August 3-4), 65-77.*

Moore, Sarah G. and Brent Mcferran (2017), "She Said, She Said: Differential Interpersonal Similarities Predict Unique Linguistic Mimicry in Online Word of Mouth," *Journal of the Association for Consumer Research,* 2 (2), 229-245.

Netzer, Oded, Feldman, Ronen, Goldenberg, Jacob and Moshe Fresko (2012), "Mine Your Own Business: Market Structure Surveillance through Text Mining," *Marketing Science,* 31 (3), 521-543.

Netzer, Oded, Lemaire, Alain and Michal Herzenstein (2018), "When Words Sweat: Identifying Signals for Loan Default in the Text of Loan Applications," working paper, Columbia Business School.

Packard, Grant, Moore, Sarah and Brent Mcferran (2018), "(I'm) Happy to Help (You): The Impact of Personal Pronoun Use in Customer-Firm Interactions," *Journal of Marketing Research,* 55 (4), 541-555.

Paetzold, Gustavo and Lucia Specia (2016), "Inferring Psycholinguistic Properties of Words," *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies,* San Diego, (June 12-17), 435-440.

Pennebaker, James W., Boyd, Ryan L., Jordan, Kayla, and Kate Blackbum (2015), *The Development and Psychometric Properties of LIWC2015.*

Pennebaker James W., Chung, C.K., Frazee, J., Lavergne, G.M., and D.I. Beaver (2014), "When Small Words Foretell Academic Success: The Case of College Admissions Essays" PLoS ONE 9(12).

Pham, Michel T. (2004), "The Logic of Feeling", *Journal of Consumer Psychology,* 14 (4), 360-
[redacted]

Pitler, Emily and Ani Nenkova (2008), "Revisiting Readability: A Unified Framework for Predicting Text Quality," *Proceedings of The Conference on Empirical Methods in Natural Language Processing,* Honolulu, (October 25-27), 186-195.

Quigley, Karen S., Lindquist, Kristen A., and Lisa Feldman Barrett (2014), "Inducing and Measuring Emotion: Tips, Tricks, and Secrets," in *Handbook of Research Methods in Social and Personality Psychology,* Harry T. Reis and Charles M. Judd, eds. New York: Cambridge University Press, 220-250.

Ranghunathan, Rajagopal and Michel Tuan Pham (1999), "All negative moods are not equal; motivational influences of anxiety and sadness on decision making," *Organizational Behavior and Human Decision Processes,* 79 (1), 56-77.

Ransbotham, Sam, Lurie, Nicholas H., and Hongju Liu (2018), "Creation and Consumption of Mobile Word of Mouth: How are Mobile Reviews Different?" *Marketing Science,* forthcoming.

Rocklage, Matthew D. and Russell H. Fazio (2015), "The Evaluative Lexicon: Adjective Use as a Means of Assessing and Distinguishing Attitude Valence, Extremity, and Emotionality," *Journal of Experimental Social Psychology,* 56, 214-227.

Rocklage, Matthew D., Rucker, Derek D., and Loran F. Nordgren (2018), "Persuasion, Emotion, and Language: The Intent to Persuade Transforms Language via Emotionality," *Psychological Science,* 29 (5), 749-760.

Rozin, Paul and Edward B. Royzman (2001), "Negativity Bias, Negativity Dominance, and Contagion," *Personality and Social Psychology Review,* 5 (4), 296-320.

Sadoski, Mark, Goetz, Ernest T., and Maximo Rodriguez (2000), "Engaging Texts: Effects of Concreteness on Comprehensibility, Interest, and Recall in Four Text Types," *Journal of*

Educational Psychology, 92 (1), 85-95.

Schwarm, Sarah E., and Mari Ostendorf (2005), "Reading Level Assessment Using Support Vector Machines and Statistical Language Models," *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics,* Ann Arbor, (June 25-30), 523-530.

Schweidel, D. A., and Wendy W. Moe (2016), "Binge Watching and Advertising", *Journal of Marketing,* 80(5), 1-19. https://doi.org/10.1509/jm.15.0258

Sismeiro, Catarina and Randolph E. Bucklin (2004), "Modeling Purchase Behavior at an E-Commerce Web Site: A Task-Completion Approach," Journal of Marketing Research, 41(August), 306-323.

Smith, Craig A. and Phoebe C. Ellsworth (1985), "Patterns of Cognitive Appraisal in Emotion," *Journal of Personality and Social Psychology,* 48 (4), 813.

Tausczik, Yla R. and James W. Pennebaker (2010), "The Psychological Meaning of Words: LIWC and Computerized Text Analysis Methods," *Journal of Language and Social Psychology,* 29 **(1),** 24-54.

Tiedens, Larissa Z. and Susan Linton (2001), "Judgment Under Emotional Certainty and Uncertainty: The Effects of Specific Emotions on Information Processing," *Journal of Personality and Social Psychology,* 81 (6), 973.

Tirunillai, Seshadri and Gerard J. Tellis (2014), "Mining Marketing Meaning from Online Chatter: Strategic Brand Analysis of Big Data Using Latent Dirichlet Allocation," *Journal of Marketing Research,* 51(4), 463-479.

Van Laer, Tom, Escalas, Jennifer Edson, Ludwig, Stephan, and Ellis A. van den Rende (2018), "What Happens in Vegas Stays on TripAdvisor? A Theory and Technique to Understand Narrativity in Consumer Reviews", *Journal of Consumer Research,* 46 (2), 267-285.

Vonne ut, Kurt and Lacey Rose (2005), "Kurt Vonnegut on Telling a Story", *Forbes* Accessed tt s://www.forbes.com/2005/10/19/vonne ut-kurt-writin -literac -comm05-cx lr 1024vonnegut.html#3b3a4f48702a

Vuilleumier, Patrik (2005), "How Brains Beware: Neural Mechanisms of Emotional Attention," *Trends in Cognitive Sciences,* 9 (12), 585-594.

Winkielman, Piotr and John T. Cacioppo (2001), "Mind at Ease Puts a Smile on the Face: Psychophysiological Evidence that Processing Facilitation Elicits Positive Affect," *Journal of Personality and Social Psychology,* 81 (6), 989.

Weary, Gifford and **Jill** A. Jacobson (1997), "Causal Uncertainty Beliefs and Diagnostic Information Seeking," *Journal of Personality and Social Psychology,* 73 (4), 839.

Warriner, Amy Beth, Kuperman, Victor, and Marc Brysbaert (2013), "Norms of Valence, Arousal, and Dominance for 13,915 English Lemmas," *Behavior Research Methods,* 45 (4), 1191-1207.

Yin, Dezhi, Bond, Samuel D., and Han Zhang (2014), "Anxious or Angry? Effects of Discrete Emotions on the Perceived Helpfulness of Online Reviews," *MIS Quarterly,* 38 (2), 539-
[redacted]

Yin, Dezhi, Sabyasachi Mitra, and Han Zhang (2016), "When Do Consumers Value Positive Vs. Negative Reviews? An Empirical Investigation of Confirmation Bias in Online Word of Mouth," Information Systems Research, 27 (1), 131-44.

Zajonc, R. B. (1965), "Social Facilitation," *Science,* 149 (3681), 269-274.

Zillmann, Dolf, Knobloch, Silvia, and Hong-sik Yu (2001), "Effects of Photographs on The Selective Reading of News Reports," *Media Psychology,* 3 (4), 301-324.

WEB APPENDIX

TABLE Al: PREDICTOR VARIABLE SUMMARY STATISTICS

	Distribution of Read Events	Distribution of Paragraphs Read
Site1	22.23%	23.83%
Site2	0.46%	0.67%
Site3	10.01%	9.33%
Site4	7.39%	6.81%
Site5	10.79%	8.81%
Site6	21.72%	18.28%
Site7	13.46%	12.84%
Site8	6.53%	7.83%
Site9	7.44%	11.60%
Desktop	62.60%	64.91%
Mobile	27.91%	25.25%
Tablet	8.99%	9.33%
Unknown	0.49%	0.51%

TABLE A2: PREDICTOR VARIABLE SUMMARY STATISTICS

			Mean	**SD**
Primary Predictor Variables	**Emotion**	**Positive Emotion**	2.72%	4.21%
		Negative Emotion	0.40%	0.06%
		Anger	0.42%	0.07%
		Fear	0.39%	0.08%
		Sadness	0.40%	0.08%
	Processing Ease	**FK**	9.20	5.66
		Parse Tree Height	6.30	2.15
		Familiarity	560.30	27.43
		Concreteness	367.71	37.99
Control Variables		WC	54.44	46.83
		%Read	0.33	0.25
		Total WC	1172.08	722.33

Means and standard deviations are averages across the paragraphs

APPENDIX: ILLUSTRATION OF HOW THE DATA CAPTURES READING BEHAVIOR

Consider the article depicted in four panels of Figure Al. When a reader clicks on a link or headline, their browser window initially displays just the header and the first few paragraphs of the article. For example, the initial window (Panel A) displays the header, paragraph 1 and part of paragraph 2. After reading the content displayed in window A, the reader then decides to either stop reading or to continue reading and scroll down.

Figure Al: Illustration of article and reading data collection for a single article.

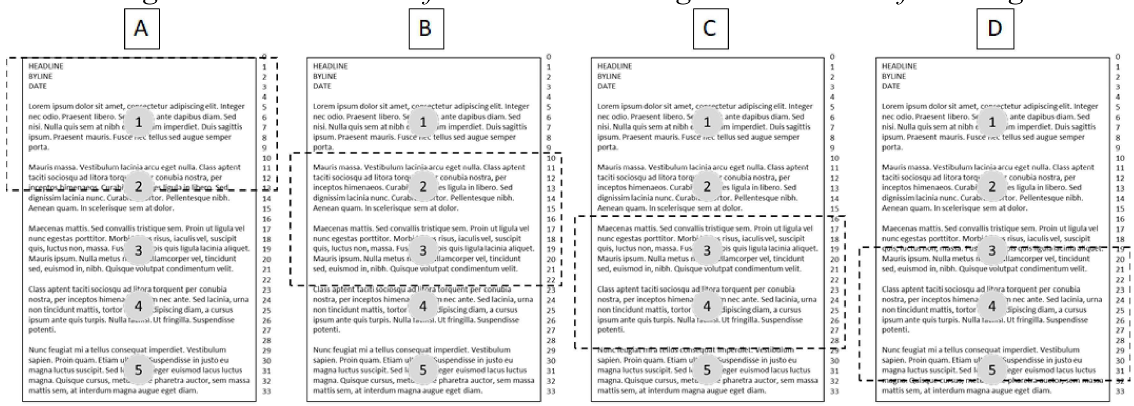

Note: The dotted line boxes indicate what is displayed in the browser window, the line numbers on the right indicate pixel positions, and the circled numbers indicate the paragraph number.

Our data is collected from JavaScript code embedded on the publishers' sites that records the pixel position associated with the top of the user's browser window. In other words, when the user sees window A, the data recorded indicates a pixel position of 0.

If the user decides to scroll down a bit (panel B), so that the top of the browser window is at pixel position 10, our data would indicate that the top of the browser window has exceeded the end of paragraph 1 (which goes from pixel position 5 to 9), and thus treat paragraph 1 as having been consumed.

As the user continues to scroll (e.g., panel C), the data continues to record the pixel position of the browser window, allowing us to infer the last paragraph consumed (e.g., paragraph 2). Note that if they scroll past the first half of the third paragraph so that the second half is showing but the first half is not (panel D), they are not treated as having consumed that paragraph. This only occurs once the user scrolls past the entire paragraph.

This process is similar regardless of device. When a reader clicks on an article on their desktop, for example, several paragraphs of content usually appear on the screen. A user may not easily be able to see as many paragraphs on their mobile device, given the smaller size of the window, but we still record what paragraphs they have scrolled past. Further while one might be concerned that differences in window size might lead to different reading experiences on mobile and desktop, the fact that results are the same even controlling for device, and that the results are the same looking at just desktop data casts doubt on this concern.

Rather than focusing on full paragraphs, one could imagine looking more finely at individual lines of text. But while this avoids concerns about readers consuming half paragraphs, it introduces other issues. Given our interest in looking at how textual features shape reading, for example, only a single line or two of text may not provide enough data to accurately measure the independent variables.

Consequently, while our approach is not perfect, it seems the best given the available data and we believe it can yield important insights. We also use more controlled experiments that allow us to better control for reading.

TABLE A3: PARAGRAPHS THAT SCORE HIGHLY ON DIFFERENT DIMENSIONS

	Emotion
Positive Emotion	*High Scoring:* Amazing mural artist Mona Caron has created a number of wall-si zed paintings in San Francisco. But this painted utility box is one of her standout works -- and it's a great optical illusion, too.
Negative Emotion	*High Scoring:* Hundreds of English Defence League (EDL) members made a trip to Birmingham on Saturday to continue doing what they love : getting drunk in public and shouting about how much they hate Muslims
Anger	*High Scoring:* After several truly infuriating accounts of the Columbia administration mishandling and neglecting sexual assault reports became public, the University pledged to clean up its act. But, as is so often and so alarmingly the case, this seems to have just been a PR smokescreen: the administration's first concern was salvaging its reputation, not the health and safety of its students.
Anxiety	*High Scoring:* The bishop of the Fargo Catholic Diocese exposed some parishioners at North Dakota churches in Fargo, Grand Forks and Jamestown to the hepatitis A virus in late September and early October.
Sadness	*High Scoring:* "When 17-year-old Zach Sobiech learned that the rare form of bone cancer he's battling likely has left him with only a few months to live, he wrote and recorded this song as a way of saying goodbye to his family and friends."
	Processing Ease
FK	*High Reading Level:* Last week we highlighted a few retro gaming wallpapers for some Friday fun. That was nothing compared to DeviantArt member Orioto's high-resolution , you-won't-believe-your-eyes paintings of classic video games for your desktop. A must-see for any gaming enthusiast. *Low Reading Level:* You go to a convention, and you come home with 55 cards in your pocket. If one or two cards have photos, you'll remember those people .
Parse Tree Height	*High Parse Tree Height:* Goldman's move is a blow to junior analysts who put in legendarily long hours for a shot at moving up the ladder. It is the latest sign that the financial industry is grappling with issues such as pay and perks amid an uneven economy and tight new rules limiting profits. A typical investment-banking analyst class at a Wall Street firm numbers about 100 people. *Low Parse Tree Height:* Attention Black Friday shoppers: You're probably wasting your time.
Familiarity	*High Scoring:* If you're lucky enough to win the upcoming $600 million Powerball, you'll also be putting a big smile on Uncle Sam's face. *Low Scoring:* If you listen to the architect Kengo Kuma, the craze for kyosho jutaku, that distinctly Japanese variant of the micro home, started in the thirteenth century, when the poet Kamo no Chomei penned an essay about the joys of living in a shack called An Account of My Hut. Contemporaneously speaking, though, micro homes became a thing in the 1990s, when rising real estate prices and a nagging recession spurred many young Tokyo residents to reconsider suburbia.
Concreteness	*High Scoring:* On October 17, 1814, over a quarter million gallons of beer were unleashed onto London's streets. The 15-foot tall tidal wave of booze crashed into buildings and flooded cellars, even killing eight especially unfortunate souls. The culprit? A bursting vat. *Low Scoring:* When US-Soviet relationships were at their frostiest in the 1980s, there was no telling what sort of exotic threat was about to come roaring through Russia's Tron Curtain. That's where the Defense Intelligence Agency came in.

STUDY 2 MATERIALS

Low Anger

In El Salvador, where abortion and termination of pregnancy is illegal, a few women have been arrested. Fortunately, now women's rights groups and their allies in congress believe they may be able to assemble a majority of votes to approve abortion under certain conditions. Two bills have been proposed in the legislature, opening up debate on the issue for the first time since the wholesale ban was passed two decades ago.

High Anger

In El Salvador, where termination of pregnancy whether abortion, miscarriage, or stillborn, is illegal, 17 women are currently sitting in prison, convicted of assassinating their children, that is, convicted of not giving birth to healthy, live babies. They are known as Las 17, and eight of these women were convicted based on a "test" has been ruled unsound for over a century.

FIGURE A2: ANGER AND SUSTAINED ATTENTION

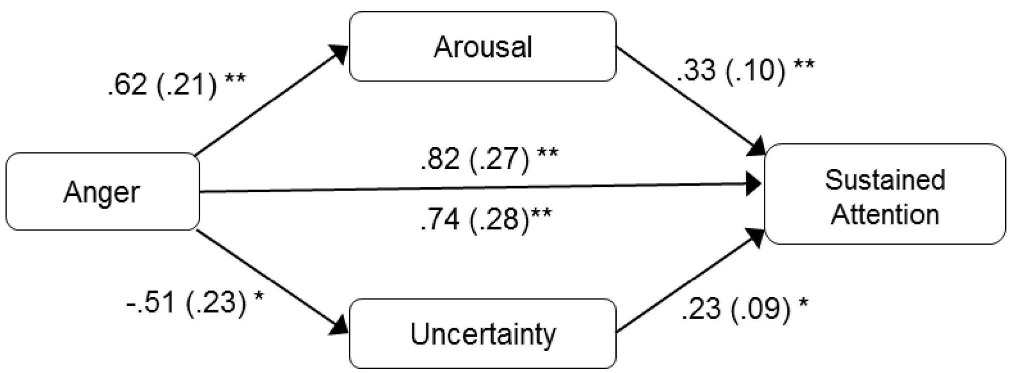

FIGURE A3: ANXIETY AND SUSTAINED ATTENTION

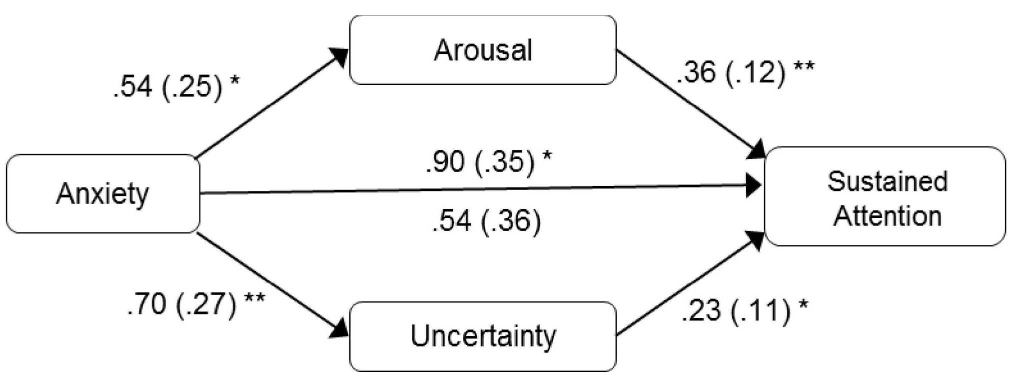

FIGURE A4: SADNESS AND SUSTAINED ATTENTION

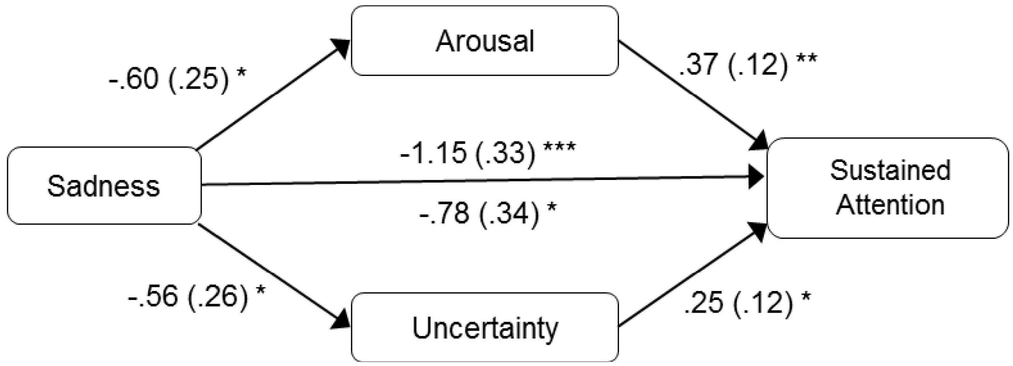

STUDY 5: MANIPULATING INCIDENTALEMOTIONS

While Studies 2, 3, and 4 manipulate integral emotion, through the content of the article itself, to even more cleanly manipulate specific emotions, study 5 manipulates emotion exogenously (see Cavanaugh, Bettman, and Luce 2015, Berger 2011, for similar approaches). Depending on condition, participants wrote about a time they felt either angry, anxious, or sad. Participants in the control condition wrote about a neutral stimulus (i.e., office products). Then, as part of an ostensibly unrelated experiment, participants read a couple paragraphs of a neutral article and reported their interest in continuing to read more. If emotion impacts reading, as we suggest, then the emotion induced in the first task should spill over into the second. Even though everyone read the same article, the incidental emotion manipulation should impact reading. By keeping the actual article the same across conditions, and manipulating emotion incidentally, we ensure that any observed difference between conditions is driven by emotion rather than some other factor.

Further, similar to studies 2-4, we measure arousal and uncertainty and test whether they mediate the effects.

Method

Three hundred and two Mechanical Turk participants completed two ostensibly unrelated studies in exchange for a small payment. To keep the study short, no demographic information was collected. Participants were randomly assigned to one of four between subject conditions (control vs. anger vs. anxiety vs. sadness).

First, we manipulated specific emotions. Prior literature reviews (Quigley, Lindquist, and Barrett 2014) suggest that recalling an emotional experience is one of the cleanest ways to manipulate emotions, so we rely on this approach (see Cavanaugh et al 2015 and Berger 2011 for similar designs). Participants were asked to list three to five things that made them feel a specific emotion, and then to write about one of those things in greater depth. In the anger condition, for example, participants were asked to list three to five things that make them angry and then "Describe in more detail a situation that made you feel particularly angry," writing "in a way that someone reading it might even get angry just from learning about the situation." The prompt was identical in the anxiety and sadness conditions except that the word angry was replaced with the word anxious or sad. To make sure all participants engaged in a similar task, participants in the control condition completed the same writing task but wrote about a neutral topic (i.e., office products). Manipulation checks suggested the manipulations worked as intended. [10]

Second, we measured the hypothesized process using the arousal and uncertainty measures from Study 2

Third, we measured the dependent variable. After completing the first "study" participants moved on to the "second." This involved reading content from a news article (i.e., about wireless phone changing, adapted from the *New York Times)* and answering some follow up questions. After reading the first paragraph, participants were asked "how much do you want to read the rest of the article" (1 = not at all, 7 = very much).

Finally, all participants were shown more of the article and the study concluded.

Results

Emotion's Effect on Reading. As predicted, a one-way ANOVA found that emotion shaped reading ($F(3, 299) = 7.23$, $p < .001$). Consistent with our theorizing, while anxiety ($M = 4.52$, $t(299) = 2.37$, $p = .018$) and anger ($M = 4.24$; $t(299) = 1.56$, $p = .12$) made people want to read more (compared to the control condition, $M = 3.75$) sadness decreased their interest in reading more ($M = 3.09$, $t(299) = -2.02$, $p = .045$).

Underlying Process. A series of bias-corrected simultaneous mediation models (Hayes 2017) tested whether arousal and certainty can explain the effects of emotion on reading.

First, we examined the anxiety manipulation. As predicted, compared to the control condition, boosting anxiety increased arousal (coeff = .42, se = .21, t = 2.04, p= .04) which, in tum, increased reading (coeff = .90, se = .11, t = 8.13, p < .001). Increasing anxiety also increased uncertainty (coeff = 1.48, se = .24, t = 6.28, p < .001) which, in tum, increased reading (coeff = .22, se = .08, t = 2.60, p = .01). A simultaneous mediation shows that both arousal *[ab*= .38, 95% CI .02 to .80] and uncertainty *[ab*= .32, 95% CI .07 to .63] drove the effect of anxiety on reading.

Second, we examined the sadness manipulation. As predicted, compared to the control condition, increasing sadness decreased arousal (coeff = -.50, se = .24, t = -2.07, p= .04) which led to decreased reading (coeff= .66, se = .10, t = 6.55, p < .001). Boosting sadness also increased uncertainty (coeff = .86, se = .22, t = 3.90, p < .001) which, in tum, increased reading (coeff= .38, se = .11, t = 3.14, p < .001). Finally, a simultaneous mediation shows that both arousal *[ab*= -.33, 95% CI -.70 to -.02] and uncertainty *[ab*= .32, 95% CI .14 to .61] drove the effect of sadness on reading.

Third, we examined the anger manipulation. As predicted, compared to the control condition, increasing anger increased arousal (coeff= .64, se = .20, t = 3.11, p = .002) which, in tum, increased reading (coeff= .82, se = .11, t = 7.58, p < .001). Boosting anger also increased uncertainty (coeff= .47, se = .17, t = -2.81, p = .006) which, in tum, increased reading (coeff= .26, se = **.11,** t = 2.42, p = .02). Finally, a simultaneous mediation shows that both arousal *[ab*=

[10] Using a similar natural language processing approach to the field analysis, we calculated the amount of anxiety, anger, and sadness indicated in each participants' writing. As expected, the manipulations changed the amount of anxiety ($F(3, 299) = 36.568$, $p < .001$), anger ($F(3, 299) = 31.40$ $p < .001$), and sadness ($F(3, 299) = 13.48$, $p < .001$) participants expressed. The anxiety manipulation, for example, evoked more anxiety (M = 4.08%) than any of the other conditions (Ms < .5%, ps < .001). Similarly, the anger manipulation evoked more anger (M = 3.44%) than any of the other conditions (Ms< .7%, ps < .001) and the sadness manipulation evoked more sadness (M = 2.27%) than any of the other conditions (Ms< 1.14%, ps < .003). This suggests that each manipulation evoked the key emotion of interest and not other specific emotions.

.52, 95% CI .19 to .93] and uncertainty *[ab*= .13, 95% CI .02 to .32] drove the effect of anger on reading.[11]

Discussion

Study 5 extends the findings of Studies 2-4 through manipulating emotion incidentally. Anxiety and anger made people more interested in reading the rest of an article. Sadness made them less interested. Further, consistent with our theorizing, specific emotions influenced arousal and uncertainty, which, in tum, increased reading.

[11] Anger may not have increased certainty due to the specific control condition used. We tried to select an innocuous control condition (i.e., writing about office products) but participants still reported an extremely high level of certainty (M = 6.1 on a 7-point scale) making it difficult for any emotion condition to boost certainty even further. Ancillary data we collected using a control condition where participants didn't write about anything however, showed a decreased level of certainty in the control condition, such that anger did in fact increase certainty.

TABLE A4. DISTRIBUTION OF TOPICS ACROSS ARTICLES

Topic	Average Posterior Probability
1	2.4%
2	3.6%
3	3.3%
4	3.4%
5	6.3%
6	2.8%
7	4.2%
8	4.7%
9	4.3%
10	2.9%
11	2.4%
12	5.6%
13	3.9%
14	3.0%
15	4.0%
16	0.9%
17	3.5%
18	3.3%
19	4.0%
20	8.2%
21	8.1%
22	4.7%
23	2.8%
24	4.7%
25	3.2%

TABLE A4: TOP JO STEMMED WORDS FOR EACH LDA TOPIC

Topic 1	china, year, oil, countri, product, chines, energi, industri, global, price
Topic 2	studi, health, brain, peopl, patient, drug, medic, doctor, test, bodi
Topic 3	compani, store, sale, product, price, custom, year, market, buy, retail
Topic 4	human, war, anim, peopl, year, live, centuri, histori, robot, dog
Topic 5	movi, film, charact, stori, episod, time, comic, star, seri, scene
Topic 6	food, eat, drink, water, cook, coffe, wine, restaur, tast, beer
Topic 7	state, law, govern, offici, case, polic, report, presid, court, public
Topic 8	camera, screen, design , batteri, iphon, phone, display , devic, video, power
Topic 9	game, week, season, yard, defens, team, play, point, pass, offens
Topic 10	work, job, school, student, colleg, year, univers, worker, employe, program
Topic 11	build, citi, air, flight, project, design, airlin, space, area, plane
Topic 12	time, hand, thing, work, day, sleep, clean, water, minut, bodi
Topic 13	year, citi, car, day, hous, famili, time, live, room, travel
Topic 14	money, pay, tax, cost, year, plan, insur, card, credit, save
Topic 15	book, read, work, design, art, write, imag, artist, model, publish
Topic 16	discussion, materi, warn, graphic, repli, unapprov, map, imag, counti
Topic 17	compani, busi, invest, fund, firm, manag, investor, billion, year, market
Topic 18	game, team, player, play, year, state, sport, footbal, coach, fan
Topic 19	percent, market, year, bank, rate, price, stock, growth, expect, quarter
Topic 20	app, googl, window, file, set, phone, comput, android, featur, work
Topic 21	peopl, time, thing, good, work , feel, person, lot, talk, help
Topic 22	video, peopl, music, facebook, twitter, photo, song, post, news
Topic 23	data, servic, site, network, secur, number, internet, onlin, user, call
Topic 24	women, men, sex, girl, woman, sexual, femal, man, parent, children
Topic 25	earth, space, univers, light, time, planet, imag, scienc, scientist, system

TABLE AS: EFFECTS FOR JUST DESKTOP

	Estimate	**SE**	**P-Value**
Intercept	2.723	0.055	<0.001
Paragraph Word Count	-0.011	<0.001	<0.001
Percent Read	-3.146	0.025	<0.001
Percent Read squared	-2.650	0.025	<0.001
Amount Read	0.004	<0.001	<0.001
Amount Read squared	-0.000	<0.001	<0.001
FK	-0.015	<0.001	<0.001
Parsevec	-0.025	<0.001	<0.001
Concreteness	0.001	<0.001	<0.001
Familiarity	0.001	<0.001	<0.001
Positive	0.006	<0.001	<0.001
Anger	0.492	0.032	<0.001
Fear	0.4333	0.031	<0.001
Sadness	-0.636	0.035	<0.001