2015

Herd Journalism: Investment in Novelty and Popularity in Markets for News

Benjamin Ho
Vassar College

Peng Liu
Cornell University School of Hotel Administration, pl333@cornell.edu

Herd Journalism: Investment in Novelty and Popularity in Markets for News

Abstract
Consumers of news care both about the novelty of the news they read, as well as how popular that news topic is with others. Editors choose what to report on based on consumer preferences and the coverage of their competitors. We build a continuous time model that predicts whether news providers invest in covering novel news stories or instead report on popular. We construct a dataset of cover stories to test the model and find that both novelty and popularity are associated with increased sales, but popular stories only have a lifespan of one to two weeks. We find that the impact of novelty has declined since 2006, and the lifespan of a story has shortened. While theory predicts that editors strategically alternate between reporting on novel stories and popular stories, we find evidence of positive serial correlation in the popularity of cover stories. Finally, news outlets often cover the same story. Theory predicts that competitive news outlets are more likely to pool on the stories they cover while news outlets with market power are more likely to separate.

Keywords
news media, popularity, novelty, cover stories, herd journalism, JEL L82, JEL L86

Disciplines
Economic Policy | Journalism Studies

Comments
Required Publisher Statement

Herd Journalism: Investment in Novelty and Popularity in Markets for News

Benjamin Ho

Vassar College

Peng Liu

Cornell University

Final version published 2015 in *Information Economics and Policy*, 31, 33-46

Benjamin Ho, Economics Department, Vassar College, 124 Raymond Ave, Poughkeepsie, NY 12604, United States. Tel.: (650) 867-8270. E-mail: beho@vassar.edu.

Peng Liu, School of Hotel Administration, Cornell University, Ithaca, NY 14853. Email: pl333@cornell.edu.

We would like to acknowledge excellent research assistance by Elly Dembo, Andrew Joung, Steve Choi, Haoyu Li, Chris Leung and excellent comments from seminar participants at Cornell, UT Houston, Vassar College and KIEP. Also thanks to discussion from Lisa George and excellent comments from the participants at the Workshop on Media Economics and Public Policy.

Abstract

Consumers of news care both about the novelty of the news they read, as well as how popular that news topic is with others. Editors choose what to report on based on consumer preferences and the coverage of their competitors. We build a continuous time model that predicts whether news providers invest in covering novel news stories or instead report on popular. We construct a dataset of cover stories to test the model and find that both novelty and popularity are associated with increased sales, but popular stories only have a lifespan of one to two weeks. We find that the impact of novelty has declined since 2006, and the lifespan of a story has shortened. While theory predicts that editors strategically alternate between reporting on novel stories and popular stories, we find evidence of positive serial correlation in the popularity of cover stories. Finally, news outlets often cover the same story. Theory predicts that competitive news outlets are more likely to pool on the stories they cover while news outlets with market power are more likely to separate.

Keywords: news media, popularity, novelty, cover stories, herd journalism

JEL classification: L82, L86

Herd Journalism: Investment in Novelty and Popularity in Markets for News

Consumer preferences for the novelty and the popularity of news stories affect the investment decisions by news providers. Editors must decide which news events they should cover, and how that coverage depends on the coverage of their competitors. A novel story is characterized as breaking news, a story on a topic about which the news consumer was previously unaware. A popular story is a story that many people already know about due to past coverage by other news outlets. A news story is valuable to consumers both because it provides new information, but also because it provides information that everybody is talking about—i.e. it follows the herd. We show that novel stories and popular stories are associated with higher sales, with the optimal lifespan of a story being just one to two weeks. We model how news firms choose which stories to cover in relation to the novelty and popularity of stories they have covered in the past and in relation to the novelty and popularity of the news covered by their competitors.

Our model of news production is a continuous time model where competing news providers make investment decisions about which news events to cover at regular time intervals. Consumers base their consumption decisions on the novelty and the popularity of the stories being offered. We then take the model to the data. The main approach will be to look at the choice of cover stories of the most popular weekly newsmagazines and how that choice is associated with newsstand sales. Specifically we measure the novelty and popularity of magazine cover stories and the similarity of coverage between competitors.

Much recent research on the news media has asked whether the industrial organization of the news industry affects the bias in the news consumers receive (Gentzkow and Shapiro, 2008;

Anderson and Coate, 2005), providing policy implications on the effects of media consolidation. Theoretical models (Mullainathan and Schleifer, 2005; Chiang and Knight, 2011) and empirical studies (Gentzkow and Shapiro, 2006; Galvis et al., 2012) find that media bias can be explained by profit maximizing firms adapting their product to satisfy the political leanings of their customers.

We ask a complementary question: what determines the content of news coverage in terms of the trade-off between covering something novel[1] and covering something popular. Franceschelli (2011) shows that a one-day delay in reporting of a breaking news story by newspapers in Argentina led to a 65% reduction in readership, and that this preference for speed increased as news transitioned online. On the other hand, news outlets may instead choose to cover an older topic simply because it is currently popular in the public mindset. When Time magazine covered a feature story on the morality of Jesus one week, the same story was covered in other outlets in following weeks. Stories involving the kidnapping of one particular girl often take a life of their own. This practice of covering what is popular is widely termed herd journalism or pack journalism.

The study of the popularity of herds has a substantial theoretical provenance in economics (Bernheim, 1994; Bikhchandani et al., 1992; Pesendorfer, 1995; Karni and Schmeidler, 1990) with substantial empirical evidence through field experiments showing that people are more likely to consume something when it is more popular (e.g. Duflo and Saez, 2003; Beshears et al., 2011). While much of this literature has focused on the incentive to conform to what is popular, a more recent literature has noted that individuals sometimes choose

[1] As a precursor to the current study, we conducted a preliminary survey of 63 people chosen from a national sample through the Study Response Project and found that out of seven news provider attributes (local relevant, style, alignment of political views, speed of coverage, quantity of advertising, accuracy, other), speed of coverage (rated 3.2 out of 4) was second most important to consumers, second only to accuracy (3.7 out of 4).

to diverge from the herd (Heath et al., 2006). Amaldoss and Jain (2010) calls the desire to be unique a snob effect, while Brewer (1991) discusses the innate preference for optimal distinctiveness. Both of these papers discuss a similar trade off to the one that newspapers face: how to be distinct while also conforming to what is popular.

Our empirical approach follows the content analysis approach that has grown out of computer science. Wu and Huberman (2008) observes user behavior on *digg.com*, a website that aggregates news stories from other news providers to estimate the influence of novelty, popularity, and attention on *digg.com* users' demand for news. Leskovec et al. (2009) employ a computational intensive approach to identify how certain linguistic phrases (''memes'') spread from news outlets to blogs. Their approach has been widely adopted in the computer science literature to examine diffusion in online social networks (Gomez Rodriguez et al., 2010) or the influence of Twitter on the news (Kwak et al., 2010; Hong, 2012). Our focus here is on competition and market structure and employ a more straightforward and less computationally demanding approach by using the keyword descriptors provided by the news aggregator service *Proquest* to classify news stories.

The ultimate goal of this project is to understand how consumer demand for novelty and popularity affects the amount of innovation for news. Coverage of novel news is particularly important for maintaining an informed electorate in a democracy. While recent economic research has shown that political bias in news media has little effect on voting (Gentzkow et al., 2011; Prior, 2013), we are interested in how well the news media keeps the public informed.

As in the literature on media bias, we are interested in the effect of market structure, particularly as the rise of online news has led to the collapse of traditional print media. The baseline theoretical framework models ''small'' media firms that take public attention on a topic as exogenous, just as small firms take prices as exogenous.[2] Additionally, in the appendix, we consider extending the model to firms with market power in the market for attention. Empirically we find that the role of novelty in generating magazine sales has largely faded since 2006, and the lifespan of how long a popular story can generate higher sales has fallen. The lifespan of a story was one to two weeks before 2006, falling to under one week after. How much of this decline can be attributed to the rise of online news (Hong, 2010; Simon and Kadiyali, 2007), versus market consolidation within the news magazine industry (Anderson and Coate, 2005; George, 2007) cannot be separately identified.

Part of the difficulty in our analysis is measuring the novelty and popularity of news events. We approach this problem in two ways: through an analysis of keywords, and by having present day readers evaluate the titles of cover stories in the past. While each measure has its own problems, the combination of the two offers a more complete picture. To assess measures of novelty, popularity, and similarity, we evaluate keywords by comparing the keywords the news magazines chose to cover, with the frequency of coverage of those same words in a national sample of 64 newspapers in the same week. We validated these measures by having workers—recruited from Amazon's Mechanical Turk—rate each article (or pair of articles) on the same measures of novelty, popularity and similarity.

[2] Here, we use the term ''small'' in the technical sense as it is used in the theoretical international trade literature, which for convenience sake sometimes even models the United States as ''small.'' We acknowledge that the firms we are analyzing are the two largest news magazines (that is why they were chosen), and therefore they arguably do have agenda setting power in deciding what is popular. We discuss how this would change the model in the appendix but leave a fuller analysis to future research. Also, assuming a larger market definition of all news providers (the internet, newspapers, cable news), the market share of these magazines are small.

Theoretical framework

We motivate our empirical findings with a novel theoretical framework to better understand the economic motivations for how editors select stories vis-a-vis their competitors. While most theories of news content focus on the choice of political bias, we focus on the firm's investment in either novel or popular news coverage. Doing so helps us understand the motivations of the news providers and the welfare consequences for consumers.

We begin by assuming that firms can only publish one news story in any given period (for example, magazines only print one cover story) and firms take popularity to be exogenous (as if they behaved as if they were a small player in the marketplace of media attention). Consumes care about both the novelty and popularity of the stories they consume, and can only choose to consume a single news magazine due to limits to their attention as in Falkinger (2007). For simplicity we assume the price and marginal cost of a magazine is equal to zero and that firms are supported by advertising. Thus firms simply seek to maximize the number of consumers who choose their product.[3]

We assume the novelty of a given story decays over time. Firms therefore find it is optimal to either publish the newest news, or to publish older news at the point where the decline in novelty stops being offset by a story's high popularity. Whether a firm chooses the most novel story or the most popular story depends on whether the rate of change of popularity in the current period is above or below a given threshold that depends on the rate of novelty decay and the consumer's elasticity of substitution between novelty and popularity.

The model also predicts that a news provider will tend to follow a novel story in one week with a popular story the following week and vice versa. This is because when the firm

[3] Of course papers may have other preferences. See for example Fan (2013), Gal-Or (2012), Geylani and Yildirim (2011) for models with multi-channel revenue streams.

chooses a novel story, regression to the mean implies that the rate of change in popularity will likely be low the next week. Alternatively, when the firm chooses the most popular story to cover, the rate of change in popularity will likely be high in the following week.

General setup of the market for news game

Consider a continuous time model where there are two firms who face a continuum of consumers. Each firm offers a single bundle of news stories. Consumers choose to consume one firm's bundle. We assume the price of both goods equals zero; firms are implicitly supported by advertising, and we assume firm profits are proportional to their market share.

At each integer time, $t \in \mathbb{Z}$, consumers choose to consume the bundle produced by a single firm (we effectively assume attention constraints prevent them from consuming both). Consumer utility is a function of the novelty of the news they consume, the popularity of the news they consume, as well as some random noise term.

There is a novel potential news event at every time $t \in [0..T]$; we will use the variable τ as an index to refer to these news events where $\tau \in [0..T]$. We assume the novelty at time t of news event τ is given by $n_t(\tau)$ and declines exponentially with time. The popularity $p_t(\tau)$ of that news story will depend on the share of people who consume the news story.

While news events occur in real time *t*, news providers only offer bundles of news stories at integer values. At each integer time $t \in [0..T] \cap \mathbb{Z}$, each news provider, *i*, offers a bundle of news stories $Z_i(t)$ where a bundle consists of a set of events covered, and a function $q_t(\tau)$ which specifies the quantity of coverage for each story (alternatively one can think of q quality). The provider is constrained by space so that $\sum q_t(\tau) \leq 1$. Thus choice of coverage of any one news story has a cost in terms of opportunity cost of other stories not covered. At each time *t*, each consumer chooses one news provider, choosing based on the novelty and the popularity of the news stories within the bundles offered. There is a unit mass of consumers. News consumers' preferences are given by the sum of their felicity functions $U(Z_i(t))$ from the consumption of news in each period. We assume discount rate of 1 for simplicity and that felicity is a stationary function that depends only on the novelty and popularity associated with each news event plus $\varepsilon_{t,\tau}^{i}$ weighted by quantity of coverage. Heterogeneity of news consumer is given by $\varepsilon_{t,\tau}^{i}$, an i.i.d. mean zero noise term which follows distribution $F(\varepsilon)$. We can now write utility as

$$V = \sum_{t=0}^{T} U(Z_i(t))dt$$
$$U(Z_i(t)) = \sum_{\tau \in Z(t)} q_t(\tau)\left[u\big(p_t(\tau), n_t(\tau)\big) + \varepsilon_{t,\tau}^{i}\right] \tag{1}$$

News providers face zero costs and charge zero price but they are constrained by space. At any time *t*, firms simultaneously choose which events they wish to report on. Profits are strictly increasing in market share and are additively separable across time periods. We will consider sub-game perfect Nash Equilibrium.

The assumption that there are no producer costs in the model is not completely innocuous. While the actual cost of investigative journalism to cover breaking news has been much discussed, our contention is that for a major news outlet, the marginal cost of hiring a reporter to cover breaking news versus hiring a reporter to write about popular news is small relative to the opportunity cost of not covering a story that would have brought in higher sales. In the data presented here, a one standard deviation increase in the popularity of the news story covered brings in an extra 100,000 issues of newsstand sales alone. The impact of that decision on revenue arguably outweighs the impact of hiring an extra reporter.

From this framework, we can derive some basic results. To better match the cover story data we analyze, and for expositional simplicity, we consider the special case where firms choose to report on only one event in each period. However, the appendix contains derivations for the more general model where firms divide their reporting resources across all possible events of the day.

Analysis of base case

Consider first the simplest model where at any time *t* each firm is limited to reporting on a single event $Z_i(t) = \{z_i(t)\}$ rather than a set. Also, for now, we assume that popularity is an exogenously given function $p_t(\tau)$ differentiable in *t*. Essentially, we assume that firms are small enough in the market for news that their choices have no effect on the popularity of a news story. We consider extensions that endogenize popularity by assuming firms are big enough to the affect popularity of a news story by their investments rather than being simply ‘‘popularity-takers’’ in the appendix.

Proposition 1. *In the Market for News with homogeneous consumers when popularity is exogenous to firm choice, firms pool on the same strategy and report on the event with novelty and popularity that maximizes consumer utility.*

The proof for this and all propositions are found in the appendix.

The proposition suggests that each firm’s choice of story should not depend on competitor’s choice. While this proposition depends upon and follows directly from our assumptions that firm choices have no effect on future popularity and thus no effect on future profits, as well as the separability and independence assumptions of consumer utility, it offers a testable prediction that we can use to test the model. Evidence that firms respond to competitors choices would suggest either that the exogeneity assumption is violated, and that the firms do have ‘‘market power’’ in determining popularity or there is separation due to consumer preferences. This pooling behavior is common to two player uni-dimensional Hotelling models like ours.

A direct implication of Proposition 1 is that we can think of the firm *i*'s problem at a given time *t* to simply report on the event τ that maximizes consumer utility:

$$z_i(t) = \text{argmax} u\big(p_t(\tau), n_t(\tau)\big) \qquad s.t. \tau \leq t \tag{2}$$

Proposition 2. *In the Market for News with exogenous popularity and homogeneous consumers, if consumer preferences for popularity and novelty are first-order homogeneous Cobb-Douglas* $(u(p,n) = p^{\alpha} n^{1-\alpha})$ *, then firms either choose the most novel story (t =* ***s)*** *or choose to report on the story at a point where the relative rate of change between popularity and novelty equals a constant:*

$$\dot{p}\frac{/p}{\dot{n}/n} = -\frac{1-\alpha}{\alpha} \tag{3}$$

If we further assume novelty to be exponential that decays at rate γ − where $n_{t(\tau)} = e^{\gamma(\tau - t)}$ such that $\frac{n}{n} = \gamma$– then we can characterize an interior solution as the point where

$$\frac{\dot{p}}{p} = -\gamma\left(\frac{1-\alpha}{\alpha}\right) \tag{4}$$

In other words, firms either pursue a novel strategy where they choose to cover the newest story available, or if instead they choose to pursue a popular strategy where they choose to report on an older story at the point where the decrease in novelty is offset by the increase in popularity weighted by how much the consumer values popularity relative to novelty. Eq. (4) allows an estimate of the relative weight consumers place on novelty relative to popularity.

Proposition 3. *In the Market for News with Cobb-Douglas utility and exponential novelty decay, if firms choose to report on the most novel story* $(t = \tau)$ *then the rate of increase in popularity at today's date* τ *must be weakly higher than in the case of a non-novel story.*

This proposition follows from standard maximization conditions, that in order for the corner solution to be optimal, consumer utility from consuming the most current news story at time τ must be higher than consuming a story at time $\tau - \varepsilon$ i.e. the first order condition must be weakly positive. When firms report on the most novel news story, the rate of increase in popularity at the current time τ must be higher than when they report on an older story.

Under this framework, if popularity follows a mean reverting process than firms should follow reporting on novel stories but less popular stories with less novel but more popular stories and vice versa. We will look for this oscillation in the data.

There are no firm costs in this model and the firms play a zero sum game, so welfare only depends on consumer utility, and is maximized here because consumer utility is maximized. However, if we consider the extension in the appendix where utility is increased when there is greater investment in coverage of a particular story, then society would do better when there is more diverse coverage of the news.

Empirical strategy

To bring the model to data, we are going to think of each news event as a keyword, see Table 1 for the most popular keywords in our dataset. We will track the popularity of each keyword amongst the general public by tracking the frequency of those words in a set of 64 of the most popular US newspapers from the ProQuest database. We are primarily interested in how the novelty and popularity of these keywords influence the cover story selection of the main weekly newsmagazines: Time, Newsweek and US News and World Reports.

A key concept from the model is the influence of news organizations in the market for attention. The model makes predictions about the behavior of news organizations when they are ''small,'' i.e. their choice of news coverage does not affect the future popularity of that keyword. We can use these predictions to look for market influence in the data.

We take two strategies to analyze the importance of novelty and popularity in cover story selection. The first will employ a semantic analysis of real world news content by examining their keywords. The second will use user ratings of novelty, popularity and similarity using workers recruited through the Amazon Mechanical Turk system.

The data

The news cover stories are from three national news weekly magazines: Time, Newsweek, and U.S. News & World Report.[4] The sample contains 1655 cover stories published from February 2001 to August 2009 (615 from Time, 549 from Newsweek, and 491 from U.S.

[4] In June 2008, the U.S. News & World Report switched its publication frequency from weekly to biweekly.

News & World Report[5]). Each observation of a cover story contains the following information: title, author, keywords, and type (people, geographic or subject). In addition, we obtained Circulation Data from the Audit Bureau of Circulations, and include both weekly magazine subscriptions and newsstand sales. Our analysis will focus on single copy newsstand sales as we did not have sufficient power to detect the impact of cover story selection on future changes in weekly subscribers.

These three magazines are the top three in market share in the news magazine market in the US and together dominate the news weekly market (see Figs. 1a and 1b). Some of our analysis will focus only on Time and Newsweek as they both sell significantly more than US News, especially in the single copy sales market, and since US News exited the newsweekly market in 2008 by switching to a biweekly format. We will take advantage of this shift in market structure to look for market structure effects (see Fig. 2).

To measure the popularity and novelty of each cover story, we take two approaches: (1) an analysis of the keywords employed, (2) coding of article titles using an internet subject pool.

To analyze the keywords, we count the number of articles in daily newspapers that covered the same topic (keyword) in previous weeks. We match each keyword (as indexed by ProQuest) using the ProQuest publication database. ProQuest is an online database that indexes articles from more than 9600 publications. We count the number of times each keyword appears in one of the 64 most popular newspapers indexed by ProQuest, a list of which can be found in the Appendix.[6] Each keyword from a magazine's cover story was matched into the newspaper

---

[5] There was some missing data in the sample. In addition, some issue has more than one cover story. Or had special issues above and beyond the 52 weeks of the year.

[6] We chose those 64 newspapers based on three Inclusion Criteria: (1) Ranked in the top 100 U.S. Newspaper ranking by average weekday readership. (Scarborogh Research, release 2, 2006). (2) Newspaper title contains the name of one of the 15 most populous U.S. states (3) General newspapers. Newspaper title contains at least one of the following: times, daily, news, post, journal, tribune, herald, observer, or press).

database and we counted the number of articles categorized by ProQuest that used that keyword. We calculate weekly popularity for two weeks prior to the delivery date of each issue and up to three weeks after.[7]

The **popularity** of a magazine cover story will be the number of articles that appeared in the ProQuest newspaper database that are classified with the keywords identified for the magazine cover story. The popularity variables we use include:

1. POPW-2 - The number of related newspaper articles found in ProQuest published between 14 days before the magazine delivery date and 7 days before the magazine delivery date.
2. POPW-1 - The number of related newspaper articles found in ProQuest published between 7 days before the magazine delivery date and the day before the magazine delivery date.
3. POPW - The number of related newspaper articles found in ProQuest published between the magazine delivery date, and 7 days after.
4. POPW+1 - The number of related newspaper articles found in ProQuest published between 7 days after the magazine delivery date and 14 days after.
5. POPW+2 - The number of related newspaper articles found in ProQuest published between 14 days after the magazine delivery date and 21 days after.

We measure the novelty in the following two ways:

1. The inverse of the frequency that keyword appears in our dataset (high frequency occurrence means low novelty).
2. The change in popularity in the previous week. Specifically, the popularity score in week T - 1 divided by the popularity score in week T - 2.[8]

[7] Normally the delivery date is 6 days before the printed issue date for all three magazines.

Summary statistics are shown in Table 2. The popularity scores are very similar across weeks, though there is a statistically significant peak in popularity in W - 1 the week before that keyword is chosen, with popularity in the week the cover story is published, W, coming in second.

We validate the keyword analysis by recruiting workers to code the articles based on article titles. Workers were recruited through Amazon's Mechanical Turk rating system, and paid 2 cents each to rate an article or pair of articles.[9] Articles were rated twice to allow us to cross validate ratings. Each worker was shown the article title and the ProQuest keywords associated with that article and then asked the following questions:

(1) On a scale of 1-5 (1 being least and 5 being most), how novel is the following cover story title (in bold) with relevant keywords, if any (listed below the title).

(2) On a scale of 1-5 (1 being least and 5 being most), how newsworthy is the following cover story title (in bold) with relevant keywords, if any (listed below the title).

(3) On a scale of 1-5 (1 being least and 5 being most), how interesting is the following cover story title (in bold) with relevant keywords, if any (listed below the title).

(4) Please compare the similarity between the following cover stories titles (in bold) with relevant keywords, if any (listed below the title), on a scale from 1 to 5 (with 1 being least and 5 being most).

Different workers answered the novelty, newsworthiness, interestingness and similarity questions for each article or in the case of similarity, each article pair. Table 3 regresses the survey responses on our keyword based measures of novelty and popularity. While assessments

---

[8] We considered many other variants of changes in popularity scores, they largely yielded similar results.
[9] The use of MTurk to recruit experimental subjects has been growing in economics research. See for example, Horton et al. (2011) who cross validated experimental results in the Harvard lab with results obtained from MTurk using standard economic games.

of novelty and popularity years after the fact are certainly prone to measurement error, Table 3 shows that these measures had at least some explanatory power, and such measures are very plausibly exogenous for our IV estimations.

Results

Linear regression analyses can be found in Tables 4-7. Table 4 takes the keyword as the unit of analysis. Each observation is a keyword, and we regress the single copy newsstand sales of the magazines associated with a cover story featuring that keyword in a given week, on the popularity of that keyword in surrounding weeks. Column 2 adds in measures for the novelty of that keyword based on frequency of that keyword in the dataset, and based on the change of popularity for that word from the prior week. Year, and month controls to account for cyclicality of the news cycle were added for column 3 and 4. In terms of the theoretical framework, each observation can be thought of as a news event selected for the cover story. Most weeks have multiple keywords, so standard errors are clustered at the issue level. It is notable that all specifications, even the ones without fixed effects, have a surprisingly good fit with adjusted R-squared ranging from 0.64 to 0.69.

The popularity of these keywords in other media prove to be very good predictors of sales success; a one standard deviation increase in popularity increases weekly newsstand sales by 20%. However, attention seems to be limited to one to two weeks. Stories that were popular the week of and the week prior increased sales, but stories that were popular two weeks prior to publication, were associated with lower sales today. Similarly, ongoing stories that remain popular in future weeks also are associated with lower sales.

The patterns in popularity are also reflected in the impact of our measures for novelty. Note the frequency measure of novelty ranges from 0 (appears in every single issue of the magazine) to 1 (appears only once). A shift in novelty from 0 to 1 effectively doubles the newsstand sales of the magazine. Although for context, it is important to note that single copy sales account for only about 5% of total circulation, and while we looked at the impact of these variables on changes in subscription numbers, the imprecise timing of our measures made it difficult to say anything conclusive.

Column 4 adds story level variables for the average novelty and average popularity of all the keywords associated with a given issue of a magazine. We find that the story level variables actually have most of the explanatory power, and therefore, we conduct the remainder of the analysis in the paper where one observation equals a single magazine issue.

Table 5 presents the same single issue sales figures regressed now on average novelty and average popularity of all the keywords associated with each issue's cover story, so that an observation is now at the issue level. Columns 1 and 2 look at the effect of your competitors choices on sales, both in terms of whether they choose to report on the same keywords (any matches and average matches), and on the impact of the popularity of what your competitor is covering. To facilitate comparison, Columns 1 and 2 focus only on Time and Newsweek so that we have a well-defined single competitor. Even early in our dataset in 2001, Time and Newsweek had much higher circulation and by the end of our dataset in 2009, US News effectively exited the market. Coverage by your competitor of an especially popular event helps own sales, even after controlling for cyclical time trends, and similarity of coverage. This suggests that when both magazines cover the same popular story, then both see an increase in sales, but when your competitor is covering a popular story that perhaps the public is tired of, it also helps sales to cover something different.

Columns 3 and 4 attempt to examine the effect of market structure on our analysis by splitting data before and after 2006. The year 2006 represents a shift for all three news magazines, when both single copy sales and circulation after a decade of steady revenues, begins to decline, most notably for US News, as consumers switched over to online news. Consistent with this trend, the impact of popularity on sales decreases by half in 2006 (the coefficient on PopW-1 falls from 709 to 352), and the life span of a story becomes even shorter. Before 2006, a story that was popular the week after publication (PopW) would still generate increased sales, whereas after 2006, stories that are still popular the week after 2006 are associated with significantly less customer interest. Similarly, we find that the beneficial effects of coverage of a novel story only existed before 2006. Cover stories of novel stories after 2006 hurt newsstand sales.

Dynamics of novelty and popularity

Table 6 now considers trends in novelty and popularity over time. Specifically we see whether novelty and popularity of this week's cover story is positively or negatively correlated with the novelty and popularity of the story covered last week. We also consider how the novelty and popularity for one magazine is related to the novelty and popularity of coverage by its competitors.

The theory predicts that we should see oscillation between novelty and popularity. A novel story should be followed by a popular story, and vice versa. We can think of the editor's task as selecting from a menu of different events in the news, each with their own novelty and popularity score, and deciding which to put on the cover.

One concern for these regressions is the endogeneity problem from using novelty and popularity one week to predict novelty and popularity the following week, because the novelty and popularity of keywords are highly correlated across weeks when the same keywords are reported on. We stick to the frequency measure of novelty here to avoid some of the mechanical relationship between novelty and popularity. We will also use our MTurk worker assessments of story's novelty and popularity as instruments to circumvent endogeneity concerns.

Here we find little support for our main hypothesis, but endogeneity clearly plays a role. Instead we find evidence for positive serial correlation in coverage of popular coverage from week to week. From Column 1, I am less likely to cover a novel story if I covered a popular story last week. Instead, from Column 4, I am more likely to cover a popular story if I covered a popular story, although the correlation is small. The coefficient on the previous weeks popularity is 0.11 in Column 4. Perfect positive serial correlation would imply that coefficient should be 1.0. Columns 2 and 5 attempt to see how editors respond both to what their competitors (in this case either Time or Newsweek) are doing in the same week, or what they did in the previous week. Here we find that the serial correlation in popularity is half due to what each magazine did itself in the previous week but half due to editors echoing the choices of their competitors when it comes to popular stories. Finally, Columns 3 and 6 try to account for the endogeneity by using the survey measures of novelty and popularity collected from MTurk as first stage instruments for our keyword based measure. However, given the weakness of our instruments, these estimates are suspect, but they do suggest caution when interpreting the other estimates in this Table.

Finally Table 7 considers what determines whether firms cover the same story as their competitor. We are essentially asking whether a magazine is choosing to imitate the reporting of their competitor, or to innovate, by reporting on a novel story of their own. Here we use the Change in Popularity Measure for Novelty to avoid the mechanical association between having a low frequency keyword and a low probability of matching that keyword to your competitor. While nothing predicted whether there was a keyword match (perhaps because this occurred 90% of the time), the average number of keyword matches was strongly predicted by the novelty of the news story. This is consistent with the model's prediction that in a competitive environment, firms wind up pooling on the same story that maximizes consumer utility. This relationship breaks down after 2006. However, it is hard to attribute this to any one of the many shifts to the market that occurred during this time.

Other specifications

Other specifications that were considered but are not yet reported here include regressions for change in paid subscription instead of newsstand sales; probit/logit regressions for indicator dependent variables; sub-samples of only keywords that referred to people or geographic location; other measures of novelty; sub-samples that excluded the 20 most popular keywords that accounted for approximately 1/3 of the keywords in the dataset; specifications that included keyword fixed effects; specifications in logs rather than levels. All showed similar patterns to the results reported here.[10]

[10] Tables for such specifications available upon request.

Conclusions

We have introduced a theoretical model of how news providers choose to cover different news stories based on how novel or popular those stories are amongst the general public. We used this model to construct a dataset that measures the novelty and popularity of the cover stories of the major news weeklies between the years 2001 and 2009, based on the keywords that describe those stories. We validated our keyword based measures of novelty and popularity by having Mturk workers code the cover stories. This time period is also unique as it encompasses a period of major change in the market structure of the news industry as major newspapers and magazines begin to exit the market given the rise of online news.

We find that our measures of novelty and popularity do a surprisingly good job in predicting news stand sales; our regression models explain 65 to 70% of the variation. We also find that consumers care about novelty and popularity, and that they consider the coverage of both major news providers when making purchase decisions. We find that consumers do like coverage of popular stories, but that popularity only helps for one to two weeks. We find evidence that the average lifespan of popular stories have declined since 2006 to under one week, and that novelty has lost its effectiveness at selling magazines. The data is less clear in explaining how editors choose their stories, but we find evidence that magazines tend to follow popular stories with another popular story, and that before 2006, news magazines were likely to pool on novel stories, but that this pattern disappears after 2006.

We consider this analysis exploratory and hope it inspires future research. The model provides a systematic method to think about content analysis in the news market, and we construct and validate one way to measure the attributes of novelty and popularity in weekly news magazines.

Novelty and popularity have natural analogues in a broader industrial organization context. Producing a novel story can be thought of as innovation, while producing a popular story can be thought of technological transfer or network effects. The news media is an industry with consumption externalities and externalities from innovation. One consumer's consumption of a news story makes it more popular and therefore increases its value to other news consumers. Similarly, once one news provider breaks a new story, others can imitate the news provider who broke the story without doing as much work. By studying the behavior and strategies of news content providers who make these innovation decisions daily, we have access to much higher frequency data than would be available in other industries. Industries with similar externalities such as software or fashion, might follow similar patterns, and so the findings here can offer insight into innovation more generally.

Table 1. Most popular keywords.

Keyword	Freq.	Percent
American history	32	3.40
Bush, George W	63	6.27
Candidates	50	5.31
History	35	3.72
Hospitals	13	1.38
Iraq	45	4.78
Leadership	13	1.38
Obama, Barack	36	3.72
Personal profiles	86	9.14
Political campaigns	35	3.72
Politics	68	7.23
Presidential elections	57	6.06
Presidents	65	6.91
Ratings & rankings	17	1.81
Terrorism	64	6.80
US	113	12.01
United States	149	13.71
Total	941	100.00

Keywords were collected using Proquest's keywords for the cover stories of the three main news weeklies.

Table 2. Summary statistics.

	Time	**Newsweek**	**US News**
Single copy	154266.5	134288.9	40305.43
Sales	(135415.6)	(101048)	(13387.81)
Novelty by	0.2916	0.2448	0.2917
Frequency	(.3495)	(.3346)	(.3552)
Novelty by	0.1643	0.1552	0.0896
ΔPopularity	(.6999)	(.6683)	(.5702)
PopW−2	83.6854	91.0362	113.0749
	(196.8287)	(190.8573)	(238.7587)
PopW−1	91.915	96.6825	118.1569
	(204.2947)	(190.9574)	(240.3179)
PopW	92.3909	94.4823	110.2085
	(208.5375)	(189.2293)	(225.9574)
PopW+1	89.5438	90.6137	108.0736
	(209.1717)	(183.7618)	(220.6022)
PopW+2	86.9227	88.59	110.2072
	(208.5702)	(181.8956)	(223.2582)
Any matches	0.8524	0.8524	
	(0.3549)	(0.3549)	
Average	0.3708	0.3708	
Matches	(0.3022)	(0.3022)	
Observations	4568	5564	7144

Standard deviations in parentheses.

Novelty for each cover story was measured both by how frequent it showed up in the data and by the change in popularity of the keywords associated with that story. Popularity was measured by the frequency those keywords appeared in Proquests database of daily newspapers in the preceding and proceeding weeks subscripted as W-tW-t. Similarity was assessed by comparing the number of keyword matches between Time magazine and its main competitor Newsweek.

Table 3. Validating the novelty and popularity measure.

	(1) Novelty by frequency	(2) Novelty by ΔPopularity	(3) Any match keywords	(4) Avg matches keywords
Novelty (1–5 scale)	−0.0131*** (0.00173)	0.00666* (0.00351)		
Interestingness (1–5 scale)	0.0110*** (0.00149)	−0.000287 (0.00303)		
Newsworthiness (1–5 scale)	−0.00448*** (0.00153)	0.00435 (0.00311)		
Similarity (1–5 scale)			0.0322*** (0.0106)	0.0809*** (0.00877)
Constant	0.301*** (0.00841)	0.0937*** (0.0171)	0.822*** (0.0196)	0.265*** (0.0162)
Observations	16,589	16,589	1230	1230
Adjusted R-squared	0.007	0.000	0.007	0.064

Here we compare our keyword-count based measures of novelty and similarity on survey measures collected from online survey takers. Each article was rated twice to allow for cross validation, to assess each article's novelty, interestingness, newsworthiness, and similarity. We regress our keyword-count based measures of novelty and similarity on survey takers assessment.

Standard errors in parentheses.

** $p < 0.05$.

*** $p < 0.01$.

* $p < 0.1$.

Table 4. Single copy sales on popularity and novelty.

	(1) Single copy sales	**(2) Single copy sales**	**(3) Single copy sales**	**(4) Single copy sales**
PopW−2	−214.6***	−169.2***	−177.9***	20.13
	(14.55)	(16.01)	(15.29)	(17.51)
PopW−1	225.2***	175.6***	187.2***	−21.82
	(19.69)	(20.99)	(19.95)	(22.96)
PopW	163.1***	159.1***	135.5***	−0.0871
	(22.76)	(22.73)	(21.60)	(25.15)
PopW+1	−71.59***	−63.58***	−68.33***	−1.518
	(22.57)	(22.57)	(21.43)	(24.61)
PopW+2	−83.26***	−81.78***	−67.69***	2.460
	(15.00)	(14.99)	(14.34)	(16.57)
Novelty by		−1646	−4549**	−2126
Frequency		(1955)	(1870)	(1829)
Novelty by		7620***	7051***	3860***
ΔPopularity		(1133)	(1075)	(1054)
Average				−603.2***
PopW−2				(28.47)
Average				597.6***
PopW−1				(38.13)
Average				418.0***
PopW				(45.59)
Average				−140.6***
PopW+1				(46.92)
Average				−222.8***
PopW+2				(31.94)
Year and month controls	No	No	Yes	Yes
Magazine controls	Yes	Yes	Yes	Yes
Observations	16,161	16,161	16,161	16,161
Adjusted R-squared	0.643	0.644	0.682	0.698

In this table regresses the number of single copy sales (e.g. sales at newsstands) on our measures of novelty and popularity. Each observation is at the keyword level, standard errors are clustered by issue. We include measures of the popularity of the cover story keyword from two *weeks before to two weeks after the cover story was published. We first consider a specification without novelty to because of co-linearity concerns. Adding time fixed effects in Column 3 does not change estimates. Column 4 adds the average popularity of all off the keywords used in the cover of that issue.*

Standard errors in parentheses.

∗$p < 0.1$.

*** $p < 0.01$.

** $p < 0.05$.

Table 5. Single copy sales on competitor's novelty and popularity.

	(1) **Single copy sales**	(2) **Single copy sales**	(3) **Single copy sales before 2006**	(4) **Single copy sales after 2006**
PopW−2	−676.4∗∗∗	−545.4∗∗∗	−891.9∗∗∗	−70.29
	(159.2)	(187.2)	(132.0)	(130.2)
PopW−1	515.4∗∗	357.4	709.9∗∗∗	352.5∗∗∗
	(202.3)	(229.7)	(179.6)	(129.5)
PopW	174.4	206.6	233.0∗	−255.2∗∗∗
	(140.8)	(141.7)	(131.3)	(84.56)
Novelty by		−7,312	−15,864	4,310
Frequency		(19,369)	(17,851)	(11,015)
Novelty by		15,343	26,074∗∗∗	−16,028∗∗
ΔPopularity		(11,686)	(10,034)	(7,426)
Any Matching	−5,925	−1,425	−3,715	−1,452
Keywords	(12,467)	(12,924)	(11,842)	(6,518)
Average # of	−3,801	−15,355	−1,377	8,592
Matching keywords	(17,319)	(19,119)	(15,265)	(8,676)
Competitor's	−843.7∗∗∗	−664.2∗∗∗		
PopW−2	(159.1)	(187.0)		
Competitor's	877.2∗∗∗	706.2∗∗∗		
PopW−1	(154.0)	(181.2)		
Comp novelty		−3,460		
Frequency		(19,428)		
Comp novelty		21,036∗		
ΔPopularity		(11,636)		
Month/year Controls	Yes	Yes	Yes	Yes
Magazine Controls	Yes	Yes	Yes	Yes
Observations	791	791	874	434
Adjusted R-squared	0.713	0.714	0.691	0.850

In this table regresses single copy sales of a magazine on the popularity and novelty of the keywords of its cover story and the popularity and novelty of the cover story of its competitor. Here an observation is at the issue level, and we focus only on Time and Newsweek so that *"competitor" is well defined. We also include a measure of similarity of cover stories by counting the number of matching keywords.*

Table 6. Trends in novelty and popularity.

	(1) Novelty (by frequency)	(2) Novelty (by frequency)	(5) Novelty (by frequency)	(3) Pop_{W-1}	(4) Pop_{W-1}	(6) Pop_{W-1}
Previous issue's novelty (by frequency)	0.000823	−0.00421	0.773**	−9.618	−2.055	54.57
	(0.0316)	(0.0427)	(0.351)	(15.48)	(18.22)	(197.0)
Previous issue's PopW−1	−0.000137**	−1.05e−05	0.00169	0.204***	0.112***	−0.824
	(6.14e−05)	(8.89e−05)	(0.00123)	(0.0301)	(0.0379)	(0.691)
Competitor's novelty (by frequency)		0.122***			−18.30	
		(0.0426)			(18.19)	
Competitor's PopW−1		−0.000185*			0.237***	
		(9.65e−05)			(0.0412)	
Competitor's previous issue's novelty (freq)		−0.0681			−33.49*	
		(0.0425)			(18.16)	
Competitor's previous issue's PopW−1		−0.000356***			0.130***	
		(8.80e−05)			(0.0376)	
Magazine controls	Yes	Yes	Yes	Yes	Yes	Yes
Month/year controls	Yes	Yes	No	Yes	Yes	No
Instrumental variables	No	No	Yes	No	No	Yes
Observations	1191	638	1144	1191	638	1144
Adjusted R-squared	0.714	0.723		0.519	0.575	

Here we consider how novelty and popularity evolve over time. Columns (1) and (2) examine how the current issue's novelty, depends on the novelty and popularity of the cover story from the previous issue, and on the novelty and popularity of the cover stories published by its competitor. Columns (3) and (4) does the same for the current issue's popularity. Columns (5) and (6) repeats the same using an IV analysis, using MTurk survey data as an instrument.

Standard errors in parentheses.

*** $p < 0.01$.

** $p < 0.05$.

* $p < 0.1$.

Table 7. Similarity with competitor.

	(1)	**(2)**	**(3)**	**(4)**
	Any keyword match indicator	**Average number of keyword matches**	**Average number of keyword matches**	**Average number of keyword matches**
PopW−1	0.000218	0.000174	0.000232	1.45e−05
	(0.000488)	(0.000361)	(0.000441)	(0.000623)
PopW	0.000120	0.000252	0.000244	0.000311
	(0.000486)	(0.000360)	(0.000437)	(0.000627)
Novelty by	0.0145	0.0850□□□	0.104□□□	0.00419
ΔPopularity	(0.0294)	(0.0218)	(0.0241)	(0.0475)
Magazine controls	Yes	Yes	Yes	Yes
Month/year controls	Yes	Yes	Yes	Yes
Observations	923	923	602	321
Adjusted R-squared	0.135	0.097	0.147	0.074

In this table regresses the similarity of the cover stories between Time and Newsweek on the popularity of those stories and the novelty of those stories.

Standard errors in parentheses.

*** $p < 0.01$.

A

	TIME	Newsweek	U.S. News & World Report
Publication Frequency	Weekly	Weekly	Weekly (change in Frequency after June 2008)
Year of First Issue	March 3, 1923	February 17, 1933	1933 (U.S. News) 1946 (World Report) 1948 (merger)
City of Publication	New York London (Time Europe) Hong Kong (Time Asia) Sydney (The South Pacific Edition)	New York	Washington D.C.
Average Number of Subscriptions (2000-2009)	3.7 million	2.9 million	1.9 million
Average Weekly Single Copy Sales (2000-2009)	120,773	135,295	39,719

Figure 1a. Profile of news weeklies.

Sources: time.com, newsweek.com, usnews.com, Audit Bureau of Circulation.

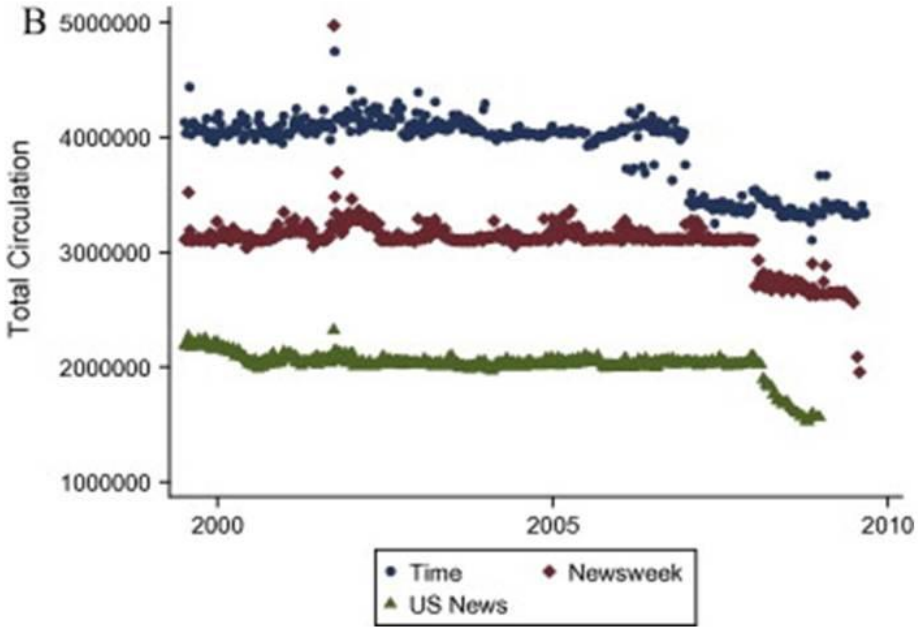

Figure 1b. Total paid sales (single copy plus circulation).

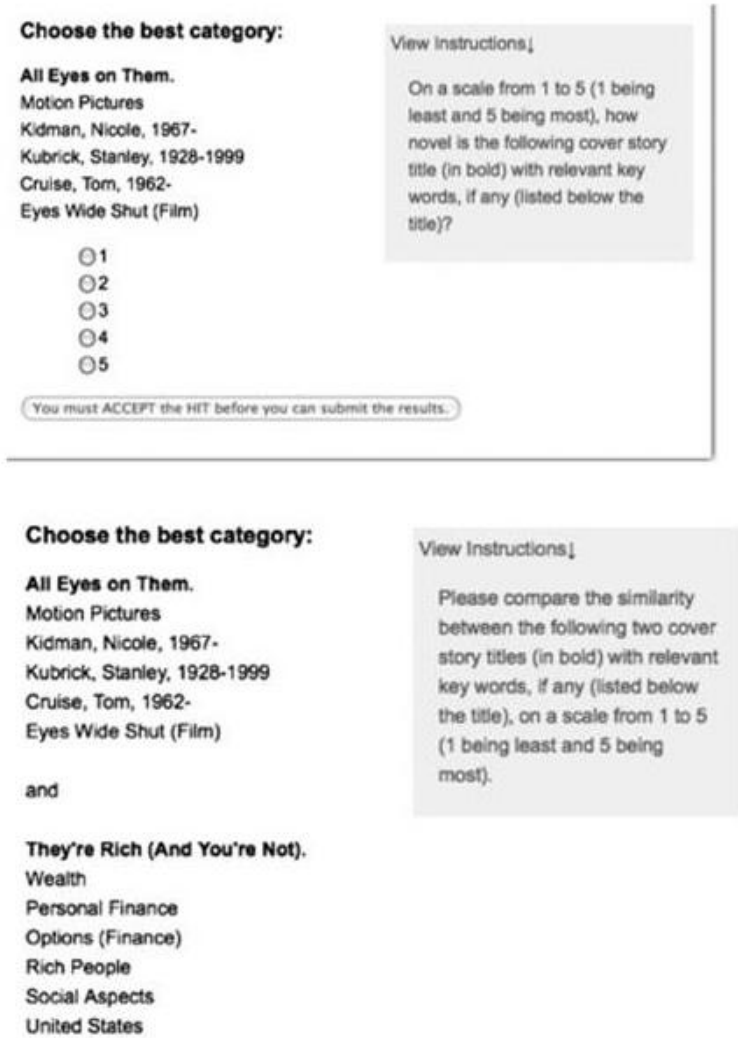

Figure 2. Example MTurk survey questions. Sample questions given to online survey participants to assess novelty and popularity of stories.

Tundra Times; Anchorage, AK	The Daily Record; Baltimore, MD
Arkansas Times; Little Rock, AR	Maine Times; Portland, ME
Arizona Capitol Times; Phoenix, AZ	Detroit News; Detroit, MI
Phoenix New Times; Phoenix , AZ	Michigan Chronicle; Detroit, MI
California Journal; Sacramento, CA	Michigan Citizen; Highland Park, MI
California Voice; Oakland, CA	Daily Record and the Kansas City Daily News-Press; MO
News from Native California; Berkeley, CA	St. Louis Daily Record / St. Louis Countian; St. Louis, MO
Oakland Post; Oakland, CA	The Riverfront Times; St. Louis, MO
Sacramento Observer; Sacramento, CA	The Charlotte Post; Charlotte, NC
San Francisco Business Times; San Francisco, CA	Daily Record; Rochester, NY
San Francisco Metro Reporter; San Francisco, CA	Inner City Press; Bronx , NY
Sun Reporter; San Francisco, CA	New York Amsterdam News; New York, N.Y.
Washington City Paper; Washington, D.C.	New York Beacon; New York, N.Y.
Washington Informer; Washington, D.C.	New York State Conservationist; Albany, NY
Broward - Palm Beach New Times; Fort	New York Times; New York, NY
Miami New Times; Miami, FL	Syracuse New Times; Syracuse, NY
Miami Times; Miami, FL	Wall Street Journal; New York, NY
South Florida Times; Coral Springs, FL	The Athens News; Athens, OH
The Jacksonville Free Press; Jacksonville, FL	Philadelphia Tribune; Philadelphia, PA
Columbus Times; Columbus, GA	Philadelphia Weekly; Philadelphia, PA
Georgia Trend; Norcross, GA	Spartanburg Herald - Journal; Spartanburg, SC
The Atlanta Journal - Constitution; Atlanta, GA	The State Journal; Charleston, SC
The Atlanta Tribune; Roswell, GA	The Tennessee Tribune; Nashville, TN
Chicago Defender; Chicago, IL	Dallas Observer; Dallas, TX
Chicago Independent Bulletin; Chicago, IL	Houston Press; Houston, TX
Chicago Review; Chicago , IL	The Texas Observer; Austin, TX
Windy City Times; Chicago, IL	The Virginia Quarterly Review; Charlottesville, VA
The Saturday Evening Post; Indianapolis, IN	USA Today; McLean, VA
The Liberty Press; Wichita, KS	North American Post; Seattle, WA
Boston Globe; Boston, MA	Seattle Times; Seattle, WA
The Beacon Hill Times; Boston, MA	Rhinelander Daily News; Rhinelander, WI
The Massachusetts Review; Amherst, MA	The Daily Reporter; Milwaukee, WI

Figure A1. Newspaper database.

References

Amaldoss, W., Jain, S., 2010. Reference groups and product line decisions: an experimental investigation of limited editions and product proliferation. Manage. Sci.

Anderson, S., Coate, S., 2005. Market provision of broadcasting: a welfare analysis. Rev. Econ. Stud.

Bernheim, B.D., 1994. A theory of conformity. J. Polit. Econ., 841-877

Beshears, J., et al., 2011. The Effect of Providing Peer Information on Retirement Savings Decisions. No. w17345, National Bureau of Economic Research.

Bikhchandani, S., Hirshleifer, D., Welch, I., 1992. A theory of fads, fashion, custom, and cultural change as informational cascades. J. Polit. Econ., 992-1026.

Brewer, M.B., 1991. The social self: on being the same and different at the same time. Pers. Soc. Psychol. Bull. 17 (5), 475-482.

Chiang, C.F., Knight, B., 2011. Media bias and influence: evidence from newspaper endorsements. Rev. Econ. Stud., rdq037.

Duflo, E., Saez, E., 2003. The role of information and social interactions in retirement plan decisions: evidence from a randomized experiment. Quart. J. Econ., 815-842.

Falkinger, J., 2007. Attention economies. J. Econ. Theory 127, 266-294.

Fan, Y., 2013. Ownership consolidation and product characteristics: a study of the US daily newspaper market. Am. Econ. Rev. 103 (5), 1598-1628.

Franceschelli, I., 2011. When the Ink is Gone: The Transition from Print to Online Editions. Working Paper.

Gal-Or, Esther, Geylani, Tansev, Yildirim, Tuba Pinarc, 2012. The impact of advertising on media bias. J. Mark. Res. 49 (1), 92-99.

Galvis, Angela Fonseca, Snyder Jr., James M., Song, B.K., 2012. Newspaper Market Structure and Behavior: Partisan Coverage of Political Scandals in the US from 1870 to 1910. Unpublished Manuscript.

Gentzkow, Shapiro, Jesse M., 2006. What Drives Media Slant? Evidence from U.S. Daily Newspapers. NBER Working Papers 12707, National Bureau of Economic Research, Inc.

Gentzkow, M., Shapiro, J.M., 2008. Competition and truth in the market for news. J. Econ. Perspect.

Gentzkow, M., Shapiro, J.M., Sinkinson, M., 2011. The effect of newspaper entry and exit on electoral politics. Am. Econ. Rev. 101 (7), 29803018.

George, L., 2007. What's fit to print: the effect of ownership concentration on product variety in daily newspaper markets. Inform. Econ. Policy 19 (3), 285-303.

Gomez Rodriguez, M., Leskovec, J., Krause, A., 2010. Inferring networks of diffusion and influence. In: Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. AMC, pp. 1019-1028 (July).

Heath, C., Ho, B., Berger,J., 2006. Focal points in coordinated divergence.J. Econ. Psychol. 27 (5), 635-647.

Hong, S., 2012. Online news onTwitter: Newspapers' social media adoption and their online readership. Inform. Econ. Policy 24(1), 69-74.

Horton, Rand, Zeckhauser, 2011. The online laboratory: conducting experiments in a real labor market. Exp. Econ.

Karni, E., Schmeidler, D., 1990. Fixed preferences and changing tastes. Am. Econ. Rev.

Kwak, H., Lee, C., Park, H., Moon, S., 2010. What is Twitter, a social network or a news media? In: Proceedings of the 19th International Conference on World Wide Web. ACM, pp.

591-600 (April). Leskovec,Jure., Backstrom, Lars., Kleinberg,Jon, 2009. Meme-trackingand the dynamics of the news cycle. Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. ACM.

Mullainathan, Schleifer, 2005. The market for news. Am. Econ. Rev. Pesendorfer, W., 1995. Design innovations and fashion cycle. Am. Econ. Rev.

Prior, Markus., 2013. Media and political polarization. Ann. Rev. Polit. Sci. 16, 101-127.

Simon, D.H., Kadiyali, V., 2007. The effect of a magazine's free digital content on its print circulation: cannibalization or complementarity? Inform. Econ. Policy 19 (3), 344-361.

Wu, F., Huberman, B.A., 2008. Popularity, novelty and attention. In: Proceedings of the 9th ACM Conference on Electronic Commerce (Chicago, Il, USA, July 08-12, 2008). EC '08. ACM, New York, NY, pp. 240-245. 10.1145/1386790.1386828

Appendix A. Proofs

Proof of Proposition 1. It will be helpful to define the market share, *s*(*t*), of firm *i* defined as the fraction of consumers who choose news provider *i* at time *t*. If τ is the news story covered by firm *i* and τ' is the news story covered by firm *i*'s competitor, where G is the c.d.f. of the difference of the noise parameter, ε, then

$$s(t) = \Pr\left[u\big(p(\tau), n(\tau)\big) + \varepsilon^{i}_{Zt} > u\big(p(\acute{\tau}), n(\acute{\tau})\big) + \varepsilon^{i}_{Z\acute{t}}\right]$$

$$s(t) = \Pr\left[u\big(p(\tau), n(\tau)\big) + u\big(p(\acute{\tau}), n(\acute{\tau})\big) > \varepsilon^{i}_{Z\acute{t}} - \varepsilon^{i}_{Zt}\right] \quad (5)$$

$$s(t) = \mathrm{G}\left(u\big(p(\tau), n(\tau)\big) + u\big(p(\acute{\tau}), n(\acute{\tau})\big)\right)$$

Since the c.d.f. is an increasing function, the firm's market share, and hence profits are increasing in the utility its news coverage provides for consumers. At any time, s, backward induction allows us to consider only the Nash Equilibrium of that stage game, and furthermore, given that firms in Nash Equilibrium treat the behavior of others as a constant, and the fact that consumers are ex ante homogeneous, we can think of firms as choosing the news coverage that maximizes consumer utility (see Fig. A1).

Proof of Proposition 2. Taking first order conditions, the firm chooses the event τ to satisfy

$$\frac{\partial u}{\partial \tau} = \frac{\partial u}{\partial p}\frac{\partial p}{\partial \tau} + \frac{\partial u}{\partial n}\frac{\partial n}{\partial \tau} + \lambda(t - \tau) = 0 \quad (6)$$

If we let consumer utility be Cobb-Douglas between the popularity and the novelty of the news event consumed

$$u(p, n) = p^{\alpha} n^{1-\alpha} \quad (7)$$

then we can show that the firm either chooses a corner solution by choosing the most novel story available, where $\tau = t$, , or the firm selects an interior solution $(\lambda = 0)$ where the relative rate of change between popularity and novelty equals a negative constant:

$$\frac{\dot{p}\ /p}{\dot{n}\ /n} = -\ \frac{1-\ \alpha}{\alpha} \tag{8}$$

Proof of Proposition 3. Standard maximization theory tells us that two conditions that are necessary for a corner maximum (i.e. a novelty strategy), (1) that the value function is increasing at the corner and (2) the second order conditions are not satisfied at the point where the first order condition equals zero.

Note in the case of Cobb-Douglas utility with exponential novelty decay, he first order condition is given by $\frac{\dot{p}}{p}\ (\tau) \geq \frac{(1+\alpha)y}{\alpha}$ while the second order condition is given by

$$\frac{\partial^2 \Delta u}{\partial t^2} = \alpha \left[\frac{\partial^2 p}{\delta t^2}\frac{1}{p(t)} - \frac{\partial p}{\delta t}\frac{1}{p(t)^2}\right] + (1-\alpha)\left[\frac{\partial^2 n}{\delta t^2}\frac{1}{n(t)} - \frac{\partial n}{\delta t}\frac{1}{n(t)^2}\right]\ \square$$

Appendix B. Extensions to theory

The theory presented has focused on the simplest case where firms report only on a single story, and have no effect on the future popularity of the stories they cover.

B.1. Analysis with endogenous popularity and small or short lived firms

Now, instead of assuming that popularity is exogenous, instead, we let popularity depend on the amount of coverage in previous periods. To that end, we assume that in every period, firms choose a set of events to cover and a quantity of coverage for each event such that the total quantity fits into the space constraints. Each event has some inherent interestingness factor ρ_τ, such that the popularity of any story τ at time *t* is given by the interestingness factor times the sum of all prior coverage on that event:

$$\mathrm{P}_t(\tau) = \rho_\tau \sum_{i=0}^{I} \sum_{j=0}^{t} q_j^i(\tau) \tag{9}$$

To simplify the analysis, assume for now that firms are either short-lived or myopic or small relative to the industry, so that they do not consider the impact of their behavior on future investments. We will relax this assumption as well in the following section.

Again, firms maximize profits by maximizing readership and their readers' utility from their choice of news coverage.

$$\Pi \cong U = \max_{q(-)} \int_{\tau=0}^{t} q(\tau)u\big(p(t)\big)d\tau \quad \text{s.t.} \sum q(\tau) \preccurlyeq 1 \tag{10}$$

Maximizing the above for each $q(t)$, and using the fact that the utility is separable by events, and that novelty is independent of quantity, we can get the following first order condition where λ is the Lagrange multiplier which represents the shadow cost of space.

$$\frac{dU}{dq(\tau)} = u + q(\tau)\frac{\partial u}{\partial p}\,\frac{2p}{\partial q(\tau)} - \lambda = 0 \tag{11}$$

From the definition of popularity, we get $\frac{\partial \mathrm{p}}{\partial q(\tau)} = \rho_\tau$ and assuming again that utility is Cobb-Douglas, we can simplify the above to get:

$$q(\tau) = \frac{\lambda - u}{\alpha\rho}\,p(\tau) \tag{12}$$

B.2 Possible extensions

In the previous sub-section, we assumed simultaneous action and myopic play. We could also introduce sequential news coverage, with a Stackleberg leader and follower at each time. The model can also be generalized with monopolistic and oligopolistic competition, with endogenous popularity and with the firms each producing a bundle (vector) of news stories each period each with different costs. We are still exploring the model's implications such as those on welfare. Thus far, the model has focused on symmetric firms, but by adding in asymmetry, we can draw a link to the literature on horizontal and vertical differentiation.