Content and Competition in Local Newspaper Markets

Elie Singer

Honors Thesis

Advisor: Professor Sevgi Yuksel

Department of Economics

University of California, Santa Barbara

Abstract

I propose a methodology to isolate the effect of competition on media content using local newspaper closures as an exogenous change in competition to closured newspapers' competitors. I define five topical metrics and construct a specialized categorization scheme to measure newspaper content over time. By comparing content across the periods before and after a newspapers' competitor closes, I hold all factors that may affect content constant save the level of competition in the market and chronological time. Following the theoretical model proposed by Perego and Yuksel (2020), I hypothesize that decreasing competition should incentivize general content and disincentivize specialized content. I test my hypothesis on a case-study of 13 local newspapers in California over the period from 2000 to 2020. I am unable to differentiate the effect of decreasing competition from other time-related factors, and therefore cannot make a definitive conclusion from the limited data. However, my methodology and categorization scheme may be applied a more extensive dataset further test whether the effect of competition on local newspaper content can be observed and isolated from other factors.

Introduction

Media content impacts social welfare by shaping citizens' beliefs, knowledge, and behavior in important political and societal dimensions. A growing body of research links media content to behavioral outcomes including voter turnout, knowledge about political parties, clearance rates, and sentencing lengths (Matrorocco 2021, Lim et al. 2014, Ash and Poyker 2021, Dellavigna and Kaplan 2007). Many of these behavioral outcomes directly affect social welfare and are of great importance to policy makers.

Perego and Yuksel (2020) theorize that competition between media providers amplifies social disagreement, which in turn may decrease social welfare. Because it is difficult to isolate competition from other factors that affect content, little empirical evidence exists prove or disprove their model. I offer a method to test the theoretical model proposed by Perego and Yuksel in the specific context of local newspapers. Following Agirdas (2015), I employ local newspaper closures as an exogenous change in competition to the closed newspaper's competitors. By comparing my content metric across the periods before and after a newspaper's competitor closes, I hold all factors that may affect content constant save the level of competition in the market and chronological time.

The following assumptions motivate local newspaper closures as an exogenous change in competition to a closed newspaper's competitors. (1) Local newspapers compete for a set readership, and interest in local news does not change over short periods of time. (2) Local newspapers are frequently the only source of information on local politics, community events, and public affairs (Lindgren et al. 2019). It follows that local newspapers do not face competition from media providers outside of their locality, and the exit of a newspaper from a local market decreases the level of competition for the remaining firms. (3) Closures are caused by factors that affect all local newspapers, including generational demographic changes and the rising popularity of social media platforms (Nonkes 2020). This implies the closure of a newspaper's competitor is unrelated to specific characteristics of the newspaper and only affects its competitive environment.

To measure content, I define five topics separated into general (*crime*, *education*, and *politics*) and specialized (*business* and *nonlocal*) groups. I calculate the ratio of topical headlines P_{topic} as the number of headlines published in each time period pertaining to a certain topic, divided by the total number of headlines published during that time period. I hypothesize that decreasing competition should positively affect the ratio of general topics and negatively affect the ratio of specialized topics. I then regress a two-period time indicator against P_{topic} to measure the change in P_{topic} across the pre-closure and post-closure periods for each source.

My original design was to collect data from a control group of sources operating in markets that did not experience a closure and use both groups to estimate the change in P_{topic} for all newspapers across chronological time. I would then normalize the measured values of P_{topic} for each source in my treatment group (of newspaper that did experience a change in competition) with respect to this estimate before comparing the pre-closure and post-closure periods. I hoped to achieve a similar estimate given sufficient overlap between pre-closure and post-closure periods in the data. However, the case-study data are separated into timeframes with non-overlapping pre-closure and post-closure periods. Time relative to closure dates is almost perfectly correlated with chronological time, so I am unable to construct an estimate of P_{topic} independent of closures. I cannot isolate the effect of competition on content from other time-based factors, and instead compare the results of simple linear regressions using only relative time as the independent variable across two non-overlapping timeframes. If competition caused observable trends in my data, I would expect to see the same trends across both timeframes. While this result would not confirm the causality of the change in competition, it would indicate that the change in competition may reasonably have some effect on content for the period 2000 to 2020.

I find that the data indicate some evidence of short-term trends in P_{topic} across time, but these trends are inconsistent between the two timeframes studied. I conclude that there is no evidence competition influenced newspaper content as measured in my data. Given the limitations of my data and methodology, I take this result as an indication that if competition does affect local newspaper content, the affect is subtle and difficult to measure. However, I cannot generalize my conclusions without more data and a sufficient set of control sources.

The remainder of my paper is organized as follows. Section I describes the related literature on media content and the contributions my analysis makes to the field. Section II describes the data used in my case-study. Section III outlines my methodology, including my metric of newspaper content, my headline categorization scheme, and my regression models. Section IV details the results of my categorization scheme and time comparisons. Finally, Section V describes the trends in the data and regression results.

I. Related Literature

My core contribution is to expand upon previous analysis of competition and media content using an application of natural language processing to categorize local newspaper headlines. A large body of research exists on the effect of competition on media, and there are many models optimizing newspaper headline classification. However, to my knowledge there are no models specifically designed to categorize local news headlines and using natural language processing to measure newspaper content by the topics I propose has been undertaken elsewhere.

My underlying theoretical framework comes from Perego and Yuksel (2021), who provide a detailed model for information specialization among information producers in a competitive environment. Their model describes how under certain competitive assumptions, producers selling information about a policy with an uncertain outcome are incentivized to tailor content to specific agents in order to differentiate their product from competitors and maximize profit. Under less competitive circumstances, producers are incentivized to produce information that is appealing to a larger and more diverse audience, which they argue may increase social welfare. My analysis applies their theory to local newspaper markets. I measure information specialization as a negative change in coverage in *crime*, *education* and *politics*, and a positive change in *business* and *nonlocal* stories (as defined in Section III).

Several recent studies have documented the effect of consolidation in the television broadcasting industry on content. Many have used the exogenous shock of Sinclair (a large, national broadcasting company) acquisitions of local television networks to measure how television content differs between locally and nationally owned stations. Mastrorocco and Ornaghi (2020) linked decreased coverage of local crime news (as a result of Sinclair acquisitions) to a negative change in violent crime clearance rates. Martin and McCrain (2019) showed that TV stations acquired by Sinclair reported fewer local politics and more national politics. In a similar study, Miho (2020) demonstrated that Sinclair ownership of local television stations increased Republican voting in the 2016 election. These analyses employed the same basic framework I propose to use on local newspapers, examining differences in content in the periods before and after a major event (in this case an acquisition) through a difference in differences model. While these studies did not all directly employ content analysis, they all relied on text-analysis results from previous research (example Hedding et al. 2019) in their assumption that Sinclair acquisition increases the conservative tone and share of national news of a local station. In a more direct use of content analysis, Blankenship and Vargo (2021) conducted a detailed study of six television stations over a four-year period to identify significant changes in coverage of local events.

Employing a different framework, Simonov et al. (2021) used the quasi-random assignment of cable news station numbers to measure how differences in content between Fox News network shows on the COVID-19 pandemic affected social distancing behavior of network viewers. Using the same instrument (originally from Martin and Yurukoglu 2017), Ash and Poyker (2019) linked increased viewership of conservative media to increased sentencing lengths among elected judges. Both studies used difference in differences models to isolate the effect of conservative media on viewer behavior.

In the study of local newspapers, Hayes and Lawless (2018) showed that a decline in coverage of congressional elections in local news resulted in less participation in congressional elections. Their study employed a content analysis of newspapers in the lead-up to the 2010 and 2014 elections. In the analysis most relevant to what I propose, Agirdas (2015) examined the change in political content of local newspapers in the period before and after a competitor closed. Agirdas limited the study to newspapers with a pre-existing political bias and used coverage of unemployment in periods where the party opposing the paper's political slant had control of the presidency as a measure of political news. The study categorized articles based on key work counts and used a difference in differences model to find a causal relationship. Other relevant studies that employ similar content analysis to what I propose include Hayes and Lawless (2015), Larcinese and Puglisi (2011), Getznow and Shapiro (2010), and Nimark and Pitschner (2019).

II. Data

I collect headlines from 13 newspapers operating in seven California localities that experienced a change in competition from 2000 to 2020. I collect raw text data from the media archive NewsBank. The newspapers I study overwhelmingly focus on local news and have a predominantly local readership. For each source, I collect all headlines from a sample of four to eight days per month over a roughly six-year period. The period is centered around the date the source's competitor closed, resulting in two three-year periods of greater and lesser competition respectively. The number of days and articles sampled varies by source and ranges from 204 days, 2,212 articles (*Coalinga Record*) to 700 days, 49,730 articles (*Sonoma Index Tribune*). On three occasions, I sample from the same newspaper across two non-overlapping periods corresponding to independent closures of difference competitors. I label these occasions with the newspaper's name followed by the letters "A" and "B" to denote the separate closure events. When possible, I collect only headlines published on Sunday to eliminate potential between-day variability in topical coverage. However, some of the sources were not available on Sunday in the NewsBank database, either because the newspaper did not publish on Sundays or because the NewsBank did not make articles published on Sundays available at the time I retrieved the data. When I could not collect articles published on Sundays, I took whichever days were available. In most cases whether I could collect articles from Sunday corresponds to whether the newspaper published every day. Newspapers from which I collected articles only published on Sunday were mostly daily papers and those that I could not collect articles published on Sunday were mostly non-daily papers. The differences in the number of days and headlines collected, whether I study a newspaper across two separate timeframes, and what day of the week I collect articles from all introduce variability between sources that may affect content. Further, my sources come from different localities, meaning there is variation in the size of the readership, the number of competitors, and other aspects of the local media market and community. All these dimensions reasonably affect content. Because I am concerned only with the *change* in content across periods of varying competition, my concern is whether these aspects affect a newspaper's reaction to a competitor closing. Newspapers that experienced multiple changes in competition over a short time period may behave differently than those that experienced only one change. In other words, the effect of the second change in competition on content may differ from the effect of the first. Daily newspaper may also respond differently to changes in competition than non-daily newspapers. This may depend on whether the newspaper that closed published daily. I assume that the variation in topical coverage between days of the week remains constant over time, but this may also not be true. Additionally, differences between individual localities and local media markets may affect how local newspapers respond to changes in competition. My study design assumes a large enough number of sources to account for these differences. If competition truly affects content and my metric of content is accurate, then I should see an overall affect across all sources regardless of between-source variation, given a large enough sample size. My case study data do not have enough sources for this to be true.

Table 1 lists the number of days and sources, closure market, size of the locality, and whether I collected headlines from for each source. The Market ID column gives and integer key denoting in which closure-market the newspaper operated. The sources are separated in to seven "closure-markets," each corresponding to the closure of a different newspaper. Figure 1 shows the days that I collected headlines for each source, colored by closure-market grouping. All newspapers colored the same way operate in the same market. The black dots mark the change in competition.

Table 1: Source Information

Newspaper Name	No. Days	No. Articles	Closure Market ID	Market Size	Sunday
sonoma_index_tribune	700	49730	2	0	0
the_press_democrat	272	18988			1
pacifica_tribune	324	13234	3	0	0
san_mateo_county_times	262	27833			1
chico_enterprise_record	336	11819	4	0	1
coalinga_record	204	2212	5	0	1
fresno_bee_A	269	13995			1
kingsburg_recorder_A	478	6912			0
selma_enterprise_A	562	7240			0
fresno_bee_B	306	9879	8	0	1
kingsburg_recorder_B	381	3098			0
selma_enterprise_B	692	5035			0
sacramento_bee	304	18072	7	1	1
mercury_news	366	42272	12	1	1
milpitas_post	551	9595			0
san_jose_examiner	498	8239			1

Figure 1: Sources by Closure Market
(*Black dots indicate change in competition)*

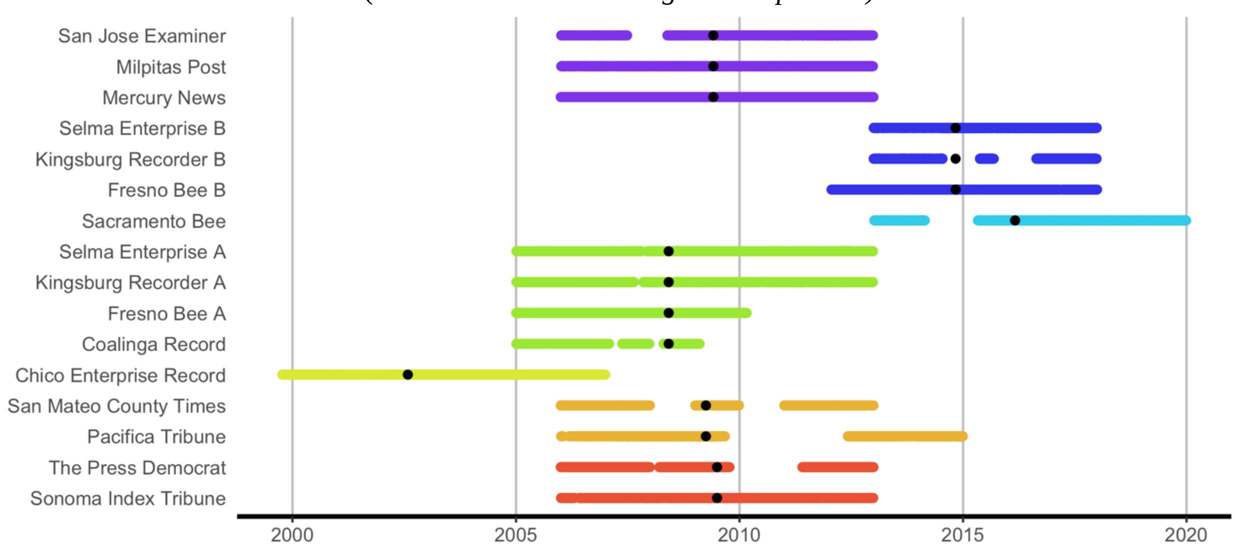

Figure 2 shows the same plot of sources over chronological time but is colored by pre-closure and post-closure periods. The gaps in the colored lines indicate periods where I could not collect headlines due to their inaccessibility from the NewsBank archive. Many of the sources had substantial time gaps where I could not collect headlines. This lowers my accuracy in measuring their content across the three-year timeframe. In particular, *Pacifica Tribune, San Mateo County Times, The Press Democrat, Coalinga Record*, and *Kingsburg Recorder B* all had substantial time gaps directly after the change in competition occurred. Figure 2 clearly shows the natural separation of the sources into two distinct groups; sources that experienced a change in competition around 2009 (timeframe A), and sources that experienced a change in competition around 2015 (timeframe B). From Figure 2 we see that there was almost no overlap between pre-closure and post-closure periods across these two timeframes. The lack of overlap means that I cannot separate chronological time and time relative to closure in my data.

Figure 2: Sources by Pre-Closure and Post-Closure Periods
(*Black dots indicate change in competition)*

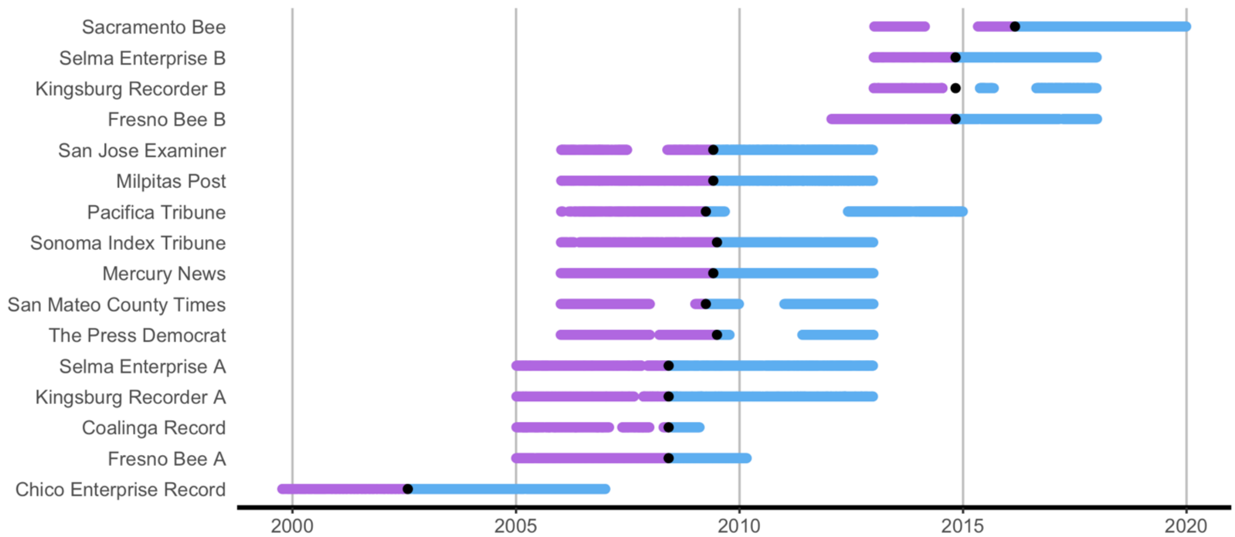

III. Methods

Using local newspapers in California as a representative case study, I outline a methodology to isolate and study the effect of competition on media content. I propose five topic definitions and calculate the number of headlines pertaining to a topic in each time period divided by the total number of headlines in that period (P_{topic}) as metric for newspaper content. Using techniques from natural language processing, I build a unique categorization scheme to measure P_{topic} over time for each source. I then compare the measured value of P_{topic} between the pre-closure and post-closure periods. I discuss a method to account for chronological time trends unrelated to changes in competition, but cannot separate chronological time from time relative to closures my case study data.

Measuring Content

I define P_{topic} as the number of headlines pertaining to a given topic divided by the total number of headlines, as shown in Equation 1. I then construct a unique categorization scheme calculate P_{topic} for each source and time period of interest. My categorization scheme omits human judgement and applies the same logic to all headlines across sources time to ensure unbiased results.

Equation 1.

$$P_{topic} = \frac{\text{No. of Topical Headlines}}{\text{Total No. of Headlines}}$$

Topic Definitions: I define five topics separated into general and specialized groups. My definitions and groupings are motivated by their separability, relative ease of categorization, and connections to measurable behavioral outcomes. I define three general topics; *crime*, *politics*, and *education*, and two specialized topics; *business* and *nonlocal*. Definitions for these topics are listed below.

<u>General topics</u>

Education: any stories covering local schools (elementary-high) or local colleges aside from major college sports. Includes local school accomplishments, funding, high school sports, and other local education issues.

Crime: any stories covering local crime.

Politics: any stories covering local politics, elections, or government. Includes coverage of local races, bills, local government votes, and any opinions, editorials or news stories related to local government and politics. Does not include national or world politics.

<u>Specialized Topics</u>

Nonlocal: any stories covering national or international news aside from national sports and business. Includes national politics, world news, and US news.

Business: any stories covering the economy (non-political), major businesses, local businesses, the stock market, or housing sales.

I do not include major college or national sports in the *nonlocal* and *education* topics. High school sports are included in *education* because they pertain to local schools and bring attention to local education in the same way as other news on local education would. While national and major college sports are of interest to many people living outside the locality of the national or college sports teams, this is generally not the case with high school sports.

Behavioral Outcomes: I choose each topic to correspond with a measurable behavioral outcome. While I do not explore these relationships in my analysis, my motivation is to lay the foundation for future work to study how revealed changes in content (if present) affects these behavioral outcomes.

Education: Increased coverage of local education presumably raises awareness of local high schools and community colleges. This may affect political pressure to fund local schools, as well as residents' tendency to send children to private schools.

Crime: Changes in crime reporting has been shown to affect criminal sentencing lengths and clearance rates in the context of local television (Matrorocco 2021, Lim et al. 2014, Ash and Poyker 2021, Dellavigna and Kaplan 2007). It is reasonable that the same relationship may be present in the context of local newspapers.

Politics: A large body of research demonstrates the importance of local newspapers in raising awareness for local political issues, helping candidates convey their message, communicating the preferences of taxpayers to local government, and limiting the negative externalities caused by corporate and governmental corruption (Nonkes 2020, Nespor 2020, Barnett 2014, Hoffman 2010). However, to my knowledge this research assumes that coverage of local politics remains constant over time. If a change in local political issues is observed, it is of great interest to study how this may affect the behaviors previously linked to local political coverage.

Nonlocal: Agirdas (2015) showed that local newspapers with an existing political bias may alter the quantity of politically biased content in response to changes in competition. As politically biased news typically relates to national politics, this case study exemplifies a change *nonlocal* coverage. Other research connects politically biased news to political preferences and behavioral outcomes with important social impact (Eberl 2019, DellaVigna and Kaplan 2018, Prior 2013).

Justification of Topic Groupings: My motivation in topic groupings is to reflect what a typical resident would desire from their local newspaper. Past research has demonstrated that reading news is a leisure activity (Nonkes 2020). I assume that residents will choose to read their local newspaper over other sources of information both in order gain information on their locality, and because of the community-building orientation of local newspapers. Crime, local politics, and local education all convey important local information and should be desirable to the typical reader. I assume the typical reader does not desire national and world news from their local newspaper because they would have otherwise chosen a different source of information. I further support this grouping by the existence of openly politically biased newspapers and the existence of political interest groups with clearly expressed preferences toward news that agrees with their bias (Agirdas 2015). I further assume that the typical resident is not interested in business coverage. I justify this assumption through the prevalence of numerous media sources that intentionally market themselves as business oriented.

Categorizing Headlines

My categorization scheme has three major components: a training data set, a keyword filtering scheme, and a Naïve Bayes classifier. All three components contribute the resulting magnitude of P_{topic}. An important feature of a content metric is its replicability and the absence of human-based judgement bias (Argirdas 2015). Many headlines are ambiguous, and human-based categorization requires making many on-the-spot judgements. It is well established that people making quick judgement decisions in this context are likely to introduce bias from unrelated factors, and the outcomes of their decisions are not consistent over time (Kahneman 2011). By applying the same machine-based categorization scheme to every headline, I ensure that my categorization is unbiased and consistent across all time periods and sources. My categorization scheme is also replicable and could be applied to new data to expand my analysis or test results on other data.

While the first two components (the training set and keyword filtering scheme) introduce human-based judgement bias, this bias only affects the sensitivity of my model to "ambiguous" headlines. I define the ambiguity of a headline imprecisely as the agreeability between interpreters to categorize the headline in the same way. While I do not calculate this metric, ambiguity could be measured as the probability that two or more randomly selected people would categorize an article in the same way. Future research may test formally measure ambiguity to more robustly describe a model's sensitivity. A key assumption I make is that the level of ambiguity (i.e., the proportion of ambiguous headlines present) remains constant over time. I expect that any trends present in a model that is less sensitive to ambiguity should also be present in a model that is more sensitive to ambiguity within a certain range. The first two components of my categorization scheme are designed to minimize the overall false-positive rate[1] and ensure that my model captures as many less ambiguous headlines as possible. I am less concerned with more ambiguous headlines for two reasons. (1) Ambiguous headlines are more difficult to accurately categorize. Given my time constraints I was unable to develop a model more sensitive to ambiguous headlines without dramatically increasing the number of false-positive classifications. (2) I assume that less ambiguous headlines will provoke a greater behavioral response among readers. Words that a model less sensitive to ambiguity may miss will likely be less provocative to a human reader. Returning to my motivation in defining topics related to measurable behavioral outcomes, it is of greater interest to study less ambiguous headlines first in this context.

Generating the Training Set: Local newspaper headlines differ from other textual bodies in the frequent occurrence of locality-specific vocabulary. Identifying locality-specific vocabulary proved crucial in distinguishing between local and non-local stories, particularly in the *politics*, *education*, and *nonlocal* topics. The most difficult categorization tasks were (i) differentiating between local and national political stories, (ii) differentiating between local crime and war or immigration stories, and (iii) finding local high school stories. I required training data containing a sufficient locality-specific vocabulary to accurately distinguish between local and non-local stories. The existing body of classified news headlines did not possess the locality-specific vocabulary I required. For this reason, I created my own training set of 6,923 headlines. Because categorizing manually would be impractical for a sample this large, I developed a key-word-search method to identify likely candidates headlines for each topic. I then manually checked the stories to confirm they were categorized correctly.

Keyword Filter: I apply keyword filtering using an expanded version of the keyword set used to construct my training data. Because I use the frequency of each word in the entire corpus as part of my vectorization method, filtering as many headlines that clearly do not pertain to any topics before

[1] I define false positives as headlines that are categorized to a given topic but do not reasonably pertain to that topic.

vectorizing greatly improves the accuracy of my model. The main limitation in keyword filtering is that I immediately discount any headlines that do not match my keyword set. While my keyword set is extensive, I presumably miss some headlines that should reasonably be categorized to one of the topics.
Naïve Bayes Classifier: After reducing the raw dataset to a smaller one consisting only of articles with a keyword match, I apply standard cleaning and tokenizing steps to the raw text. I vectorize with a term frequency—inverse document frequency (TF–IDF) vectorizer. I compare various classification methods including basic logistic regression, SVM, Naïve Bayes, and CNN. The CNN proposed by Kim (2014) performed the best according to my test metrics, but because of time limitations and the size of my dataset, I chose a Naïve Bayes classifier with a multinomial distribution. The performance of this model according to my test metrics was marginally lower than Kim's CNN, but given time limitations in the model design process, my implementation of Naïve Bayes was faster. Future research would compare categorization results between models across the entire dataset and run similar post-categorization regressions using other classifiers. I use 80% of my training data to fit the model and the remaining 20% as a test dataset. Due to computational limitations, I split the raw data into 10 subsets and run the classifier on each subset. I vectorize each subset independently, meaning my vectorization scheme differs slightly between subsets. This is because the term frequency and inverse document frequency of tokens differs across subsets. While the form of the Naïve Bayes classifier remains consistent, the test data are vectorized differently across subset-trials and the precision scores therefore vary.
Test Metrics: The precision and recall scores were nearly identical in all trials, so I only report the precision. I calculate three precision scores: one for the training-fit data (the 80% of my training data used to fit the model), one for the training-test (the remaining 20% of training data *not* used to fit the model), and one for an independent test set of 1000 randomly sampled headlines. The number of headlines in each subset after filtering ranged from 8,969 to 14,659 with a median of 12,150. The additional 6,923 training observations were appended to each subset before vectorization. The number of features varied from 109,741 to 171,767 across subset-trials with a median of 130,678 features. The average precision across all 10 models was 0.904 ($sd = 0.049$) for the training-fit data, 0.835 ($sd = 0.029$) for the training-test data, and 0.952 ($sd = 0.016$) for the independent test data. While the standard deviations indicate some variation between subsets, this variation was not large and within an acceptable range. The lowest precision value on the independent test set (which best measures the actual accuracy) was 0.919, which is high enough for the purpose of this analysis.

Time Comparison of P_{topic}

I hypothesize that decreasing competition caused by the closure of a competitor incentivizes the remaining newspapers to increase coverage of general topics and decrease coverage of specialized topics. If my hypothesis is true, then ratio of topical headlines P_{topic} should increase with decreasing competition for the *crime*, *education* and *politics* topics, and decrease with decreasing competition for the *business* and *nonlocal* topics.
Using the categorization scheme and topic definition described, I calculate the value of P_{topic} (as defined in Equation 1) per month-long period for every source. The number of months varies per source and averages around 72 (see Table 1 for more details). I define the two-period indicator I_{comp} to measure the difference in P_{topic} across the pre-closure and post-closure periods. I_{comp} is defined for every source-month pair as zero if the given month was prior to the closure-date (corresponding to the source) and one if the given month was after the closure date. See Equation 2.

Equation 2.

$$I_{comp}(\text{month},\ \text{source}) = \begin{cases} 0 & \text{if month } < \text{source}_{\,closuredate} \\ 1 & \text{if month } > \text{source}_{\,closuredate} \end{cases}$$

Note that $source_{closuredate}$ in Equation 2 is the closure date of the source's competitor (i.e., the date that the change in competition occurred). For a given source, I_{comp} is zero for all months in the pre-closure period and one for all months in the post closure period.
According to my hypothesis, the date of the change in competition occurred should not matter. A similar change in content should exist between the pre-closure and post-closure periods across newspapers whose competitors closed at different points in chronological time. However, there are many factors unrelated to competition that change across chronological time and may affect competition. In my original design, I planned to collect data from many sources with overlapping pre-closure and post-closure periods as well as control sources from markets that did not experience a closure. In this way I could account for chronological time trends affecting all sources by normalizing the values of P_{topic} for each source-month pair with respect to the average value of P_{topic} across all sources. In case study data, I am unable to estimate this value due to the lack of control sources and overlapping timeframes. I cannot construct a chronological time metric that is independent from relative time to closure.
As detailed in Section II, my data is naturally separated into two timeframes; sources that experienced a change in competition around 2009 (timeframe A), and sources that experienced a change in competition around 2015 (timeframe B). I study the effect of I_{comp} on P_{topic} across both timeframes separately and compare the results. Because there is almost no overlap between timeframes, a significant effect from I_{comp} on P_{topic} in both timeframes does not necessarily mean that competition affects content. However, consistent trends in P_{topic} across both timeframes would suggest that competition may play a role. The lack thereof would merit the opposite conclusion.
Regression Model: Equation 3 outlines the regression model applied to each timeframe. I regress the independent variable I_{comp} against the dependent variable P_{topic}. The magnitude of the coefficient β_1 gives the change in P_{topic} between the pre-closure and post-closure periods (i.e., the change in the ratio of topical headlines across this time). The sign of β_1 gives the direction of the change between periods of greater and lesser competition. In the absence of unrelated factors, a negative coefficient would indicate a positive relationship between P_{topic} and the level of competition and a positive coefficient would indicate a negative relationship between P_{topic} and the level of competition. As previously mentioned, the magnitude and sign β_1 do not necessarily represent the effect of competition but indicate if the effect may be observable in the data. Following my hypothesis, I expect a positive effect of I_{comp} for the *crime*, *education* and *politics* topics and a negative effect for the *business* and *nonlocal* topics.

Equation 3.

$$P_{topic} = \beta_0 + \beta_1 \bullet I_{comp}$$

IV. Results

Table 2 summarizes the change in the average value of P_{topic} between the pre-closure and post-closure periods for each source and topic. The distribution of the change in the average value of P_{topic} for *crime* is right skewed in the positive direction. The mean and median across all sources is near zero, but five sources saw a positive change greater than 0.5% and two sources saw a positive change greater than 1%. *Sacramento Bee* is an outlier in the negative direction and was the only source with a substantial negative change. The distribution for *education* is also nonsymmetric. While over half of the sources saw a change in the range of -3% to +2%, three sources saw a negative change greater than -7% in magnitude. The distribution for *politics* was slightly right skewed, but generally symmetric about the mean of +0.3%. There is relatively minor variation between sources for *politics,* and no clear outliers. The distribution for *business* has a large peak around zero and is symmetric save one outlier, *San Mateo County Times*, with a positive change of +3.5%. There was almost no change in the average value of P_{topic} for *nonlocal* in any of the sources. This reflects the extremely low values of P_{topic} for *nonlocal* headlines and large number of months where P_{topic} was zero. As shown in Table 3, the average value in P_{topic} for *nonlocal* did not exceed 0.02% across any timeframe.

In three sources (*Kingsburg Recorder A*, *Pacifica Tribune,* and *Selma Enterprise A*) the difference in the average value of P_{topic} deviated substantially from the mean across all sources in both the *education* and *politics* topics. All three of these sources see large negative changes for *education*, however the sign of the change in *Pacifica Tribune* was opposite to the other for *politics*. This pattern may indicate these newspapers had some common factor that caused them to behave differently than the others.

Table 2: Changes in P_{topic} (%) Between Pre-closure and Post-closure Periods

	General			*Specialized*		
Source	Crime	Education	Politics	Business	Nonlocal	Market ID
sonoma_index_tribune	-0.046	-2.765	-0.177	-0.467	-0.037	2
the_press_democrat	-0.008	2.811	-0.372	-1.183	-0.022	
pacifica_tribune	-0.136	-7.236	-2.335	-0.104	-0.015	3
san_mateo_county_times	0.460	0.958	0.377	3.583	-0.061	
chico_enterprise_record	0.265	0.027	-0.728	0.840	-0.067	4
coalinga_record	0.000	1.415	0.940	0.000	0.000	5
fresno_bee_A	-0.104	0.505	-0.030	1.305	0.050	
kingsburg_recorder_A	-0.130	-8.218	2.456	-0.009	0.000	
selma_enterprise_A	-0.124	-7.623	2.994	-0.048	0.000	
fresno_bee_B	0.136	3.136	0.790	-1.867	0.020	8
kingsburg_recorder_B	1.157	-2.918	0.035	0.000	0.042	
selma_enterprise_B	0.937	-2.838	-0.495	0.000	0.000	
sacramento_bee	-1.453	-1.223	0.268	0.246	0.012	7
mercury_news	0.109	-0.338	0.201	0.108	0.005	12
milpitas_post	0.696	-0.759	0.623	1.033	0.005	
san_jose_examiner	0.637	-1.514	-0.031	0.708	0.001	

Table 3 lists the mean and standard deviation of P_{topic}, grouped by topic during the pre-closure and post-closure periods, across all sources in the given timeframe. There is substantial variation in P_{topic} across all topics save *nonlocal* where the means are trivially close to zero. In general, the variation in P_{topic} is greater during timeframe B than timeframe A.

Table 3: Test Statistics for P_{topic}(%) By Topic, Timeframe, and Period

Topic	Timeframe	Period	Average P_{topic} *(Standard Deviation)*
Crime	A	pre-closure	0.22 *(0.54)*
		post-closure	0.43 *(0.81)*
	B	pre-closure	0.72 *(1.27)*
		post-closure	0.90 *(1.17)*
Education	A	pre-closure	8.25 *(8.08)*
		post-closure	7.57 *(5.21)*
	B	pre-closure	10.36 *(7.98)*
		post-closure	9.79 *(7.10)*
Politics	A	pre-closure	1.96 *(2.13)*
		post-closure	2.80 *(3.24)*
	B	pre-closure	1.53 *(1.47)*
		post-closure	1.79 *(2.12)*
Business	A	pre-closure	1.11 *(1.31)*
		post-closure	1.57 *(2.03)*
	B	pre-closure	1.12 *(1.52)*
		post-closure	0.49 *(1.49)*
Nonlocal	A	pre-closure	0.02 *(0.09)*
		post-closure	0.01 *(0.09)*
	B	pre-closure	0.00 *(0.00)*
		post-closure	0.02 *(0.11)*

Figure 3 plots the average value of P_{topic} for each month across chronological time. Plot I shows all sources, Plot II shows only sources in timeframe A, and plot III shows only sources in timeframe B. I exclude the *nonlocal* topic because the values of P_{topic} are trivially small. Comparing Figure 3 and Table 3 shows that the between-month variability in the average value of P_{topic} generally aligns with the variation in P_{topic} across sources. Both values are typically higher in timeframe B.

Figure 3.

Plot I:
Topical Headline % By Month, Chronological Time

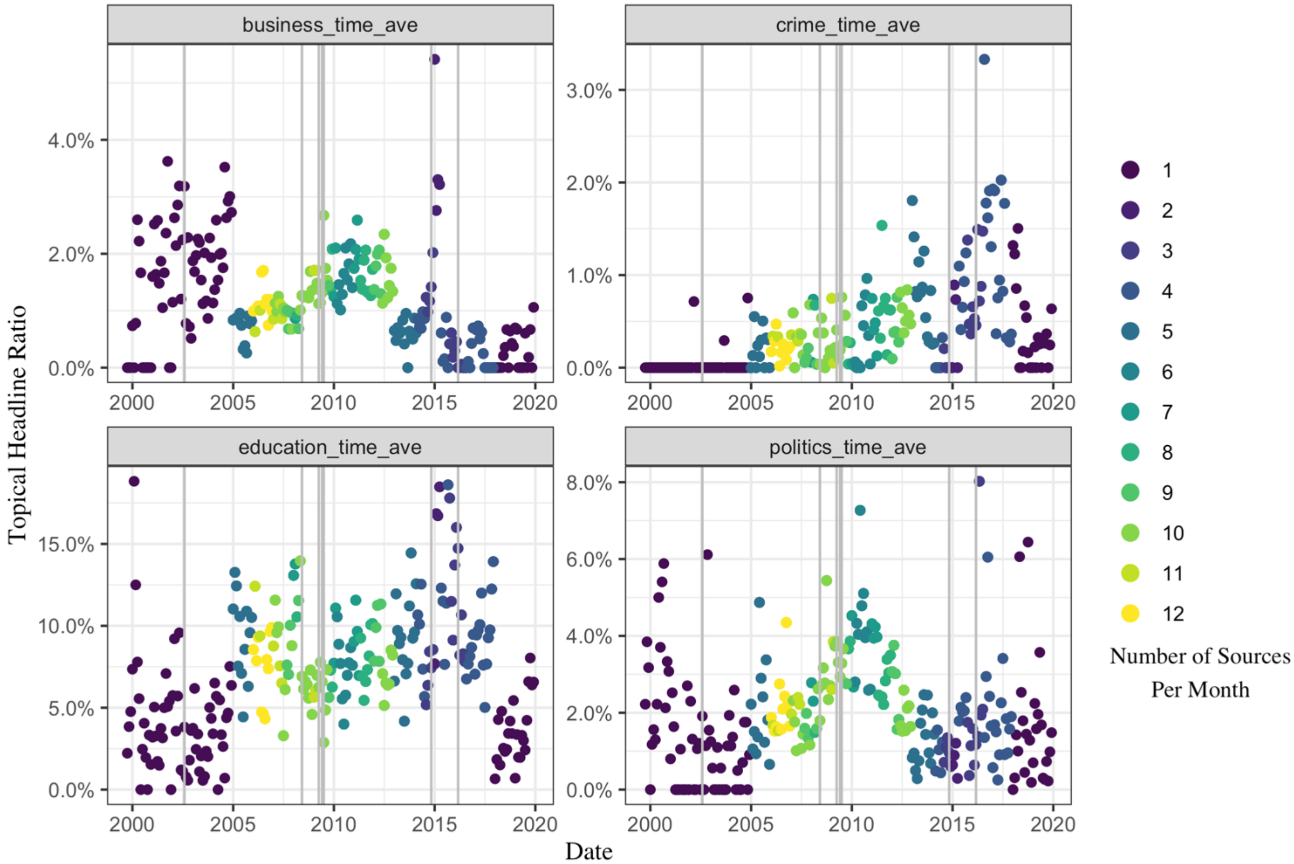

Plot II:

Topical Headline % By Month, Chronological Time (Period A)

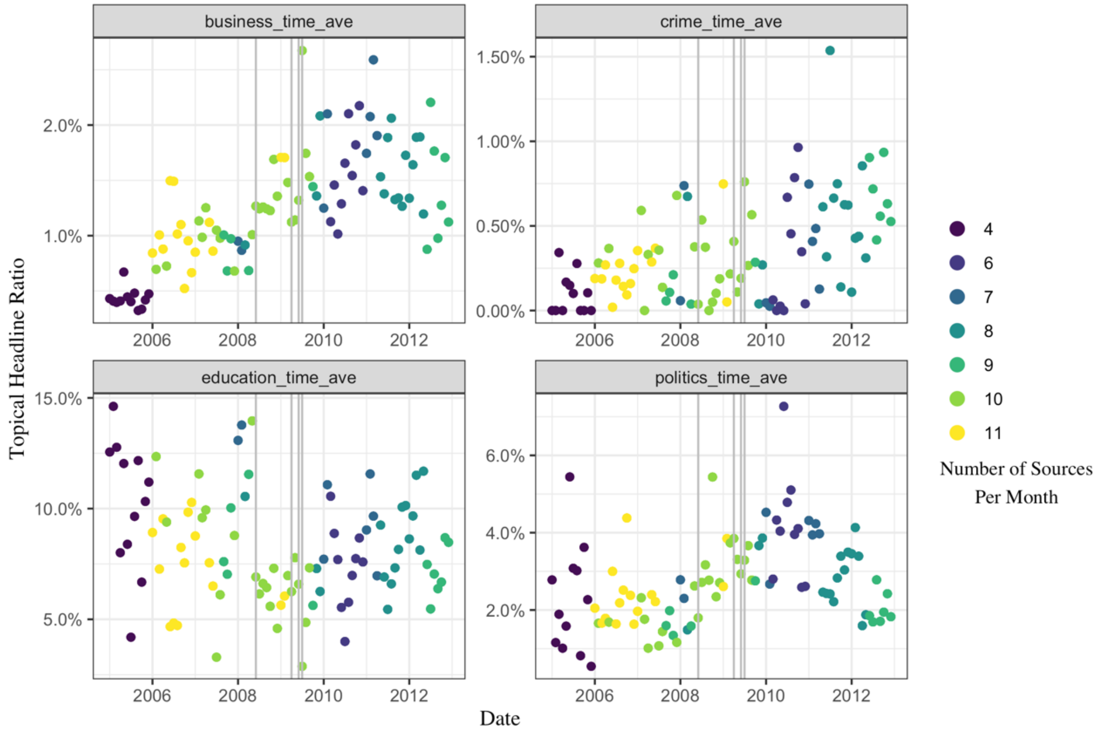

Plot III:

Topical Headline % By Month, Chronological Time (Period B)

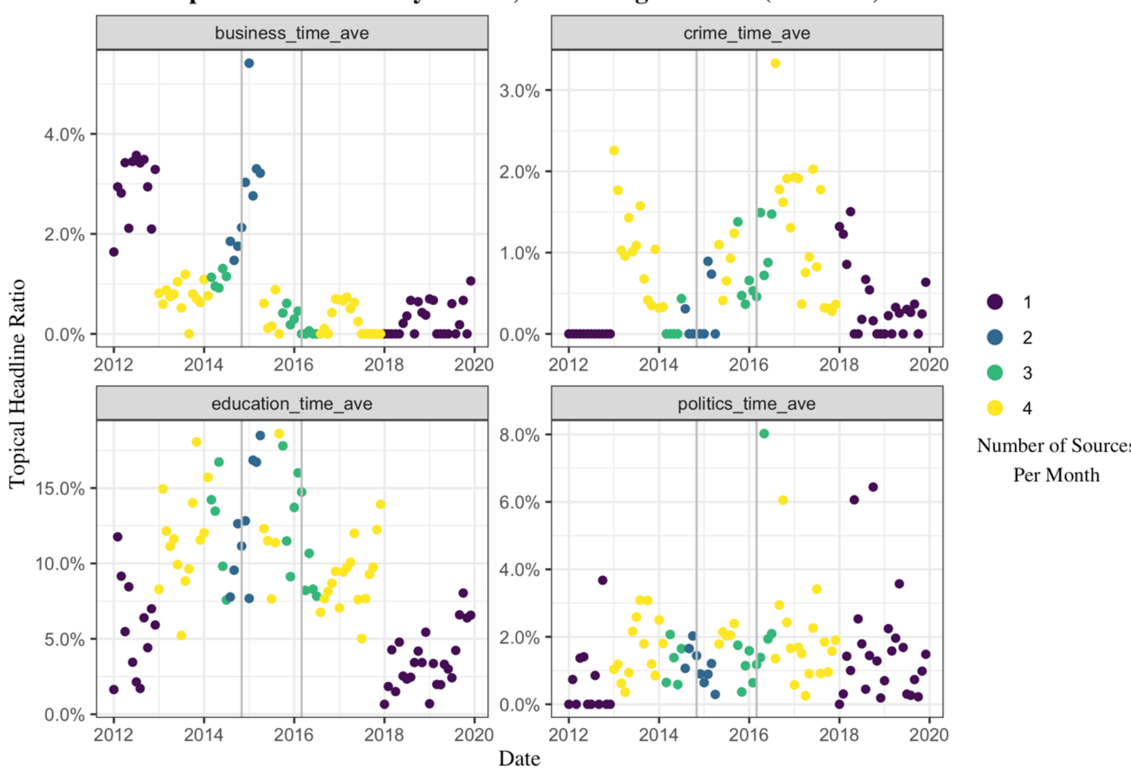

Figure 4 plots the measured value of P_{topic} for each source and month against months relative to closure. Each point shows the value of P_{topic} on the vertical axis for a given source-month pair. Time (in months) is scaled so that negative values correspond to the pre-closure period and positive values correspond to the post-closure period. The black lines indicate the month of the closure. I exclude the *nonlocal* topic because the values of P_{topic} are trivially small.

Figure 4.

Topical Headline % By Month, Relative to Closure Date

Black Line Indicates Date of Closure

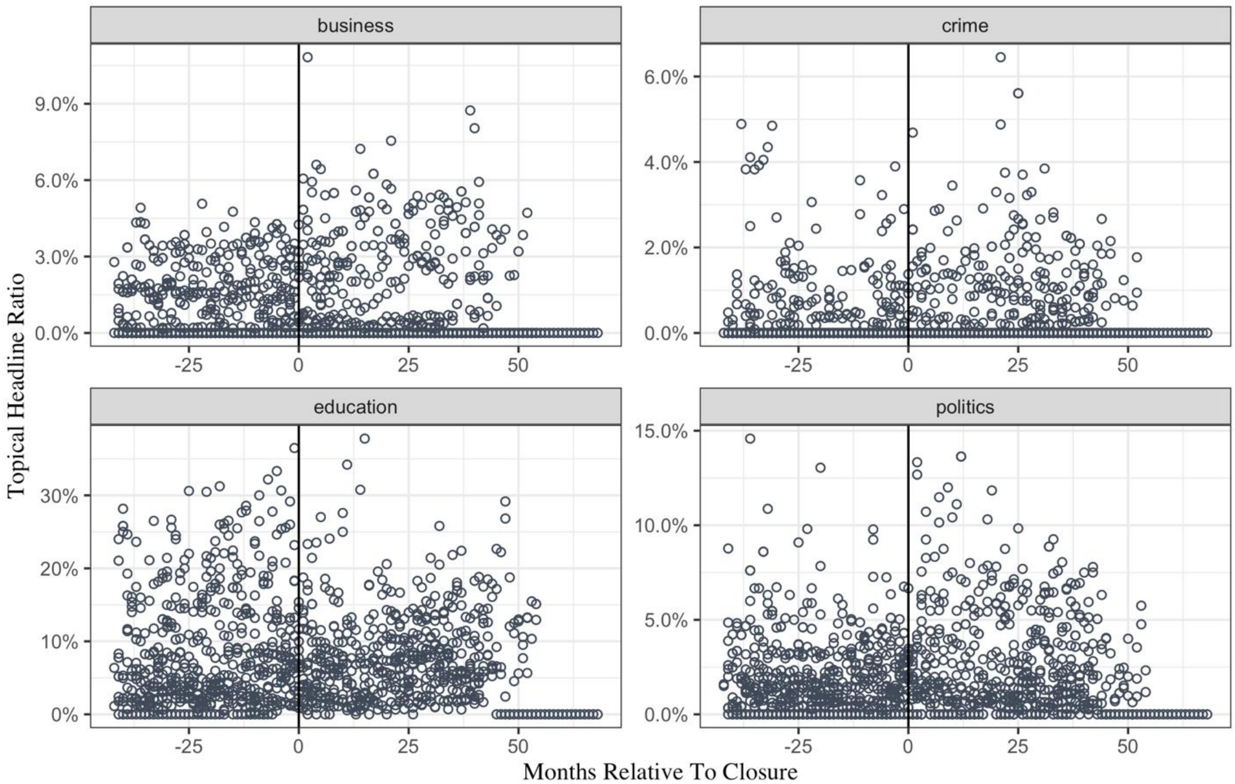

Comparison of Regression Results for Timeframes A and B.

The results of the two-period regression models applied across timeframes A and B indicate that while short-term time trends in P_{topic} may exist for a subset of the topics, the directionality and magnitude of the relationship between the time-indicator I_{comp} and P_{topic} are not consistent across any topics. No causal relationship can be determined between changes in competition and P_{topic} with the available data, and the observable trends in timeframe A are likely not related to the change in competition.

Table 4: Regression Results for Two-Period Change in I_{comp}, Timeframe A

Dependent Var.	Independent Var.	Coefficient	Std. Error	Adj. R^2	MSE
Crime	Intercept	0.002***	<0.001	0.020	0.001
(General)	I_{comp}	0.002***	0.001		
Education	Intercept	0.083***	0.004	0.001	0.009
(General)	I_{comp}	NS	0.005		
Politics	Intercept	0.020***	0.001	0.061	0.031
(General)	I_{comp}	0.013***	0.002		
Business	Intercept	0.011***	0.001	0.017	0.004
(Specialized)	I_{comp}	0.005***	0.001		
Nonlocal	Intercept	<0.001***	<0.001	-0.000	<0.001
(Specialized)	I_{comp}	NS	<0.001		

Table 5: Regression Results for Two-Period Change in I_{comp}, Timeframe B

Dependent Var.	Independent Var.	Coefficient	Std. Error	Adj. R^2	MSE
Crime	Intercept	0.007***	0.001	0.001	<0.001
(General)	I_{comp}	NS	0.002		
Education	Intercept	0.104***	0.008	-0.003	0.002
(General)	I_{comp}	NS	0.010		
Politics	Intercept	0.015***	0.002	0.001	<0.001
(General)	I_{comp}	NS	0.002		
Business	Intercept	0.011***	0.002	0.038	0.002
(Specialized)	I_{comp}	-0.006**	0.002		
Nonlocal	Intercept	NS	<0.001	0.004	<0.001
(Specialized)	I_{comp}	NS	<0.001		

Crime: The two-period time indicator I_{comp} had a significant positive affect on the ratio of *crime* stories published during timeframe A, but was not significant during timeframe B. The ratio of crime headlines increased by 0.2% across the two periods measured in timeframe A. The magnitude of the affect was relatively large, resulting in a doubling of the average value of P_{topic} across the two periods ($\mu_{\text{pre-closure}}$= 0.2%, $\mu_{\text{post-closure}}$ = 0.4%). The direction of the affect agrees with my hypothesis, which predicts decreasing competition should increase P_{topic}. However, Figure 3 shows that the data does not fit a clear linear trend. While the number of months during the post-closure period in which P_{topic} is greater than 0.5% increases substantially, the variability of P_{topic} increases from with the mean ($\sigma_{\text{pre-closure}}$ = 0.5%, $\sigma_{\text{post-closure}}$ = 0.8%). While the time-indicator is significant in the *crime* regression and there is some positive trend in the data, I cannot determine a clear linear relationship between P_{topic} and the time-indicator from the data in timeframe A. I_{comp} is not significant at all during timeframe B. Figure 3 shows that the data in timeframe B shows no increasing trend across the pre-closure and post-closure periods. Ignoring the dark-colored values (denoting points where only one source was used to calculate the average) the average value P_{topic} first seems to decrease sharply in 2013 and then rises steadily from 2014 to 2018. However, there is too much variation in the data during both periods to conclude any trend. While the mean value of P_{topic} increases from $\mu_{\text{pre-closure}}$= 0.7% to $\mu_{\text{post-closure}}$ = 0.9%, the standard deviations during both periods are greater in magnitude than the corresponding means ($\sigma_{\text{pre-closure}}$ = 1.3%, $\sigma_{\text{post-closure}}$ = 1.2%). Looking at the two periods together, the data suggests that whatever trend may exist in timeframe A only reflects short-term time variability unrelated to the change in competition. If competition truly caused the change in P_{topic} during timeframe A, I would expect a similar negative effect during timeframe B. As this is not the case, I conclude that competition did not have a clear measurable effect on P_{topic} for the *crime* topic.

Education: The two-period time indicator was not significant for either period in the *education* regression. If any trend is to be seen, the average value of P_{topic} first decreases from around 2006 to 2009, and then increases again from 2009 onwards. The average value of P_{topic} only decreases by 0.7% across the two periods. However, the high variability of P_{topic} in both the pre-closure ($\sigma_{\text{pre-closure}}$ = 5.2%) and post-closure ($\sigma_{\text{post-closure}}$ = 8.1%) periods prevent any clear relationship. With the given data, I cannot conclude that any relationship exists between P_{topic} and the time indicator during timeframe A.

Politics: The time-indicator had a significant positive affect on the ratio of *politics* stories published during timeframe A, but was not significant during timeframe B. While the regression coefficient of I_{comp} indicates a 1.3% increase in P_{topic} across the two periods, Figure 3 shows the data do not fit a linear trend. If any trend is to be seen, the average value of P_{topic} follows a parabolic curve—increasing from 2006 to 2010 and then decreasing from 2010 to 2012. There is less variability in the data ($\mu_{\text{pre-closure}}$= 2.0%, $\sigma_{\text{pre-closure}}$ = 2.1%, $\mu_{\text{post-closure}}$ = 3.2%, $\sigma_{\text{post-closure}}$ = 2.7%) than other topics, and I conclude that there may be some short-term time trend present. However, the trend is clearly not linear across timeframe A. The time-indicator was not significant during timeframe B. The average value of P_{topic} changed by only 0.3% ($\mu_{\text{pre-closure}}$= 1.5%, $\mu_{\text{post-closure}}$ = 1.8%) and the data show no positive trend. Looking at both time periods together, I conclude that P_{topic} may follow a short-term time trend in timeframe A, but the trend is likely not caused by a change in competition.

Business: P_{topic} somewhat follows a linear trend for *business* over timeframe A. The time-indicator was significant in the business regression, causing a 0.5% increase in P_{topic} across the two periods. This change is substantial relative to the period averages ($\mu_{\text{pre-closure}}$= 1.1%, $\mu_{\text{post-closure}}$ = 1.6). While the variation in P_{topic} is also substantial ($\sigma_{\text{pre-closure}}$ = 1.3%, $\sigma_{\text{post-closure}}$ = 2.0%), Figure 3 shows a clear increase between the pre-closure and post-closure periods and somewhat of a linear trend. The direction of the trend disagrees with my hypothesis that decreasing competition should cause a decrease in specialized content. While the time-indicator is significant during timeframe B, the direction of the effect is negative, and the magnitude is quite small (only 0.05%). Looking only at the months calculated with 3 or more sources (colored yellow and), Figure 3 shows a vaguely decreasing trend in the data. However, the trend not obvious. Comparing both timeframes, I conclude that while P_{topic} may follow a short-term time trend in timeframe A, there is no such trend in timeframe B. No conclusive relationship between I_{comp} and P_{topic} exists in the data.

V. Discussion

After comparing the effect of I_{comp} on P_{topic} between timeframes A and B, I deduce that (1) short term time-based trends likely exist for some of the topics in my studied timeframe, and (2) these trends are likely unrelated to changes in competition. Table 2 shows that (3) there are significant differences in the change in P_{topic} between the pre-closure and post-closure periods across sources. There is some indication that sources operating in the same closure-market may respond similarly, but this result is far from robust. I further conclude that (4) there is enough variability in P_{topic} between sources and across time that my estimate on the overall trend in P_{topic} is nontrivial only during periods where I collect data from more than ~5 sources. This implies that my sample size in timeframe B may be too small to observe any trends if present. As previously mentioned, I am unable to separate unrelated chronological time factors from competition. Even so, the data disagree with my hypothesis that competition affects content. I conclude that if competition has any effect on content, it is too subtle to accurately measure with my small sample size and high between-source variability.

While my comparison of the effect of I_{comp} on P_{topic} between timeframes A and B helps deduce whether the trends seen in timeframe A may be caused by competition, the data in timeframe B are far less than timeframe A. Timeframe A consists of eleven sources from four markets while timeframe B consists of only four sources from two markets. As shown in Figure 3, the variation in the average values P_{topic} between consecutive months greatly increases with fewer sources. Averages taken with greater than ~8 sources follow far better trends and have substantially less between-month variation than averages taken with a lesser number of sources. The amount of between-month variation greatly increases on averages taken with less than ~5 sources. The fact that timeframe A has almost three times the number of sources and that timeframe B is below the ~5 source threshold both suggest the data from timeframe B cannot reliably reveal trends in the average value of P_{topic} across time. Comparing timeframes A and B therefore has significant limitations. While I generally infer that the trends in timeframe A are not robust enough to justify the lack of trends in timeframe B, I would require a greater sample of sources in timeframe B to make a definite conclusion.

Trends Across Closure Markets: My model assumes that while differences across closure markets may affect the magnitude in the change of P_{topic} between the pre-closure and post-closure periods, the presence and directionality of the change should remain consistent. The data show consistency among sources in the same market for some topics but not all. As mentioned in the Section IV, the *Kingsburg Recorder A* and *Selma Enterprise A* follow similar trends in the *education* and *politics* topics. Both see a negative change in *education* and positive change in *politics* large in magnitude relative to the other sources (-8.2% and 7.6% respectively compared to the mean of -1.6% for *education*; +2.5% and +3.0% respectively compared to a mean of 0.3% for *politics*). While I cannot conclude these changes are outliers due to the small sample size, they are certainly on the far tails of the distribution across all sources. Both sources see a similar change in *crime* (-0.13% and -0.12% respectively) and *business* (-0.01% and -0.05% respectively). The magnitudes of these changes lie near the center of their respective distributions across all sources. The closeness of these results suggests that there may exist a similarity between the *Kingsburg Recorder A* and *Selma Enterprise A.* However, the other two sources from the same closure-market (*Coalinga Record* and *Fresno Bee A,* market 5) do not follow the same pattern. The difference in the average value of P_{topic} changes in the opposite direction for *education,* and the *Fresno Bee* has almost no change (-0.03%) for *politics*. While *Kingsburg Recorder* and *Selma Enterprise A* see a small negative change for *business, Fresno Bee A* has a substantial positive change (1.3%). All this suggests that whatever factors may have caused the similarities between the *Kingsburg Recorder A* and *Selma Enterprise A* are not present across the entire closure-market. Another important point is that these two newspapers saw similar changes in both *crime* and *education* during timeframe B (i.e., *Kingsburg Recorder B* and *Selma Enterprise B*). Both sources saw a negative change around -3% for *education* and a positive change around +1% for *crime*. The magnitudes of these changes are not as substantial relative to the other sources as in timeframe A, but still deviate significantly from the center of their respective distributions across all topics. The similarity between the two newspapers across both timeframes further suggests there may be a common factor affecting their behavior.

The directions of the changes in the average value of P_{topic} are consistent across the sources in market 2 for all topics save *education*. A similar result holds across the sources in closure market 12; the directions of the changes are consistent across all topics save *politics* where *Mercury News* and *Milpitas Post* both saw a positive change and *San Jose Examiner* saw a negative change. The are no clear consistencies between the sources grouped by timeframes A and B.

There are also no clear trends among the newspapers that experienced multiple changes in competition (*Fresno Bee*, *Kingsburg Recorder* and *Selma Enterprise*) aside from the similarity between *Kingsburg Recorder* and *Selma Enterprise* already mentioned. The decline in magnitude between periods A and B for *education* (from approximately -8% to approximately -3%) supports my inclination that a second change in competition may cause a lesser affect. However, this phenomenon does not appear among other topics and sources.

I presume that market size may affect how newspapers respond to changes in competition. The direction of the change in the average of P_{topic} between the pre-closure and post-closure periods aggress between the *Sacramento Bee* (closure market 7) and the sources in closure market 12 for the *education* topic. However, the results do not agree for the other topics, indicating there is no such relationship present in the data. But because there are only two distinct closure-markets in the large market size group, I cannot properly address this question.

There are no discernible trends when the data is grouped by the day of the week published (i.e., the Sunday indicator), save the connection between *Kingsburg Record* and *Selma Enterprise* already mentioned.

Short Term Chronological Time Trends: I find some evidence of short-term chronological time trends in the average value of P_{topic} across all sources. As mentioned in Section IV, *business* and *crime* increase from 2005 to 2012. *Politics* first increases from 2005 to 2010 and then decreases from 2010 to 2015. There are many phenomena that could explain these trends, and I do test any specific explanations in my analysis. I will mention two major events that occurred during this period—the 2008 election and the great recession. Because I include national politics in the *nonlocal* topic, the great recession offers a better explanation for the trends in *business* and *politics.* Increasing coverage of business and local governmental issues would be consistent the increased importance of government and financial issues during that period.

Improvements in Categorization: While I justify my chosen model through relatively high precision on selected test datasets, my categorization scheme could be improved to capture less ambiguous headlines. Such a model would presumably produce higher values of P_{topic} across all topics. While I compare the results of several models on my selected test datasets, these datasets were not large enough to assess whether models of higher sensitivity show different trends in P_{topic} across time. Comparing the results from different categorization schemes would test the robustness of the trends observed with my current categorization scheme and may possibly show different trends.

The values of P_{topic} using the current categorization scheme were extremely low. While I asses that most newspaper headlines do *not* pertain to any of my topics, I believe my current scheme gives a lower bound within the reasonable range of P_{topic}. It follows that the lack of observable trends in my data may be a sensitivity issue; differences in topical coverage may only be seen when including less ambiguous headlines. This is particularly pertinent to the *nonlocal* topic, where the average value of P_{topic} is near-zero. Improving the categorization scheme is an important next step to the analysis.

Conclusion

I propose a methodology to study the effect of competition on content. Using the case study of 13 newspapers operating in California localities from 2000 to 2020, I use newspaper closures instrument to capture the effect of competition on content. I define the metric P_{topic} to measure content across time and design a categorization scheme for local newspaper headlines using techniques from natural language processing. My categorization scheme addresses human-judgement bias and sample-size limitations and improves upon simple key-word searches employed in previous studies. I hypothesis that P_{topic} should increase with decreasing competition for the three general topics (*crime, education*, and *politics*), and decrease with decreasing competition for the two specialized topics (*business* and *nonlocal*). I choose topics corresponding to measurable behavioral outcomes, thereby laying the foundations for future research to study the effect of local newspaper content on behavior. A key limitation in my analysis is that I am unable to isolative the effect of competition from unrelated factors that may affect content over time in the case study data. Because the data is separated into distinct groups with no overlap between pre-closure and post-closure periods and lacks control sources, I am unable to construct an estimate of the trend in P_{topic} across all headlines.

I find that while the data shows some evidence of trends in *business, crime*, and *politics*, the trends are short-term, not necessarily linear, and are not consistent across the two timeframes. This result indicates that the observed trends are likely not caused by the change in competition. However, the small sample size in timeframe B makes trends in P_{topic} from this period less reliable than in timeframe A. I only find observable trends in the average value of P_{topic} when roughly five or more sources are used to calculate the average, indicating there is enough variation across sources and between months that trends across time may not be visible with fewer than five sources. This result confirms that data from many sources are needed to capture differences in content across time.

I further examine if the market a newspaper operates in, the market size, or the days I collect articles from affect how P_{topic} changes across time. While I find similarities in the behavior of newspapers with shared characteristics, I cannot conclude with my data that any of these factors have a significant affect. The change in P_{topic} between the pre-closure and post-closure periods varies substantially across sources and there are no clear trends when sources are grouped by market, market size, or the day of the week I collected from. My sample size is too small to make robust comparisons across groups given the high variability in the change in P_{topic} between sources.

While I do not collect a large or diversified enough dataset in my case study to make any conclusions about how a change competition affects local newspaper content, the method I propose may still be applied to a more complete dataset to study this relationship. My categorization scheme offers a consistent and unbiased method of measuring content across different sources and time periods. I improve upon previous methods of content measurement in this way. My analysis details a first step to study the relationship between competition and media content in the context of local newspaper. Future research may apply this methodology to an expanded data set to further address this important question.

Sources

Agirdas, C. (2015). What drives media bias? New evidence from recent newspaper closures. *Journal of Media economics*, *28*(3), 123-141.

Ash, E., & Poyker, M. (2019). Conservative news media and criminal justice: Evidence from exposure to fox news channel. *Columbia Business School Research Paper*.

Barnett, S., & Townend, J. (2015). Plurality, policy and the local: can hyperlocals fill the gap?. *Journalism Practice*, *9*(3), 332-349.

Blankenship, J. C., & Vargo, C. J. (2021). The Effect of Corporate Media Ownership on the Depth of Local Coverage and Issue Agendas: A Computational Case-study of Six Sinclair TV Station Websites. *Electronic News,* 19312431211043483.

DellaVigna, S., & Kaplan, E. (2008). The political impact of media bias. *Information and Public Choice*, 79.

Eberl, J. M. (2019). Lying press: Three levels of perceived media bias and their relationship with political preferences. *Communications*, *44*(1), 5-32.

Gentzkow, M., & Shapiro, J. M. (2010). What drives media slant? Evidence from US daily newspapers. *Econometrica*, *78*(1), 35-71.

Hayes, D., & Lawless, J. L. (2015). As local news goes, so goes citizen engagement: Media, knowledge, and participation in US House Elections. *The Journal of Politics*, *77*(2), 447-462.

Hayes, D., & Lawless, J. L. (2018). The decline of local news and its effects: New evidence from longitudinal data. *The Journal of Politics*, *80*(1), 332-336.

Hedding, K. J., Miller, K. C., Abdenour, J., & Blankenship, J. C. (2019). The Sinclair effect: Comparing ownership influences on bias in local TV news content. *Journal of Broadcasting & Electronic Media*, *63*(3), 474-493.

Hoffman, Lindsay H., and William P. Eveland Jr. "Assessing causality in the relationship between community attachment and local news media use." *Mass Communication and Society* 13, no. 2 (2010): 174-195.

Janowitz, M. (1967). The community press in an urban setting: The social elements of urbanism. Chicago: University of Chicago Press. (Original work published 1952)

Kahneman, D. (2011). *Thinking, fast and slow. Macmillan.*

Kim, Y. (2014). Convolutional Neural Networks for Sentence Classification. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*. 10.3115/v1/D14-1181.

Larcinese, V., Puglisi, R., & Snyder Jr, J. M. (2011). Partisan bias in economic news: Evidence on the agenda-setting behavior of US newspapers. *Journal of public Economics*, *95*(9-10), 1178-1189.

Lim, C. S., Snyder Jr, J. M., & Strömberg, D. (2015). The judge, the politician, and the press: newspaper coverage and criminal sentencing across electoral systems. *American Economic Journal: Applied Economics*, *7*(4), 103-35.

Lindgren, A., Jolly, B. Sabatini, C. & Wong, C. (2019). Good news, bad news: A snapshot of conditions at small-market newspapers in Canada. National NewsMedia Council. https://localnewsresearchproject.ca/2019/05/01/good-news- bad-news-a-snapshot-of-conditions-at-small-market-newspapers-in-canada/

Martin, G. J., & McCrain, J. (2019). Local news and national politics. *American Political Science Review*, *113*(2), 372-384.

Martin, G. J., & Yurukoglu, A. (2017). Bias in cable news: Persuasion and polarization. *American Economic Review*, *107*(9), 2565-99.

Mastrorocco, N., & Ornaghi, A. (2020). Who watches the watchmen? Local news and police behavior in the United States. University of Warwick, Department of Economics.

Miho, A. (2020). Small screen, big echo? Estimating the political persuasion of local television news bias using Sinclair Broadcast Group as a natural experiment.

Nespor, C. D. (2020). *The Effect of Newspaper Closure on Local Media Ecology*. West Virginia University.

Nimark, K. P., & Pitschner, S. (2019). News media and delegated information choice. *Journal of Economic Theory*, *181*, 160-196.

Nonkes, M. (2020). When a small market newspaper closes: A literature review. Emerging Library & Information Perspectives, 3, 157–183. https://doi.org/10.5206/elip.v3i1.8596

Perego, J., & Yuksel, S. (2018). Media competition and social disagreement. *Work*.

Prior, M. (2013). Media and political polarization. *Annual Review of Political Science*, *16*, 101-127.

Simonov, A., Sacher, S. K., Dubé, J. P. H., & Biswas, S. (2020). *The persuasive effect of fox news: non-compliance with social distancing during the covid-19 pandemic* (No. w27237). National Bureau of Economic Research