Supplementary Appendix: “Judge Lynch” in the Court of Public Opinion: Publicity and the De-legitimation of Lynching

November 13, 2018

A Supplementary Analyses

Figure A1: Lynching coverage as a function of rail network centrality — Dyads (1880–1900)

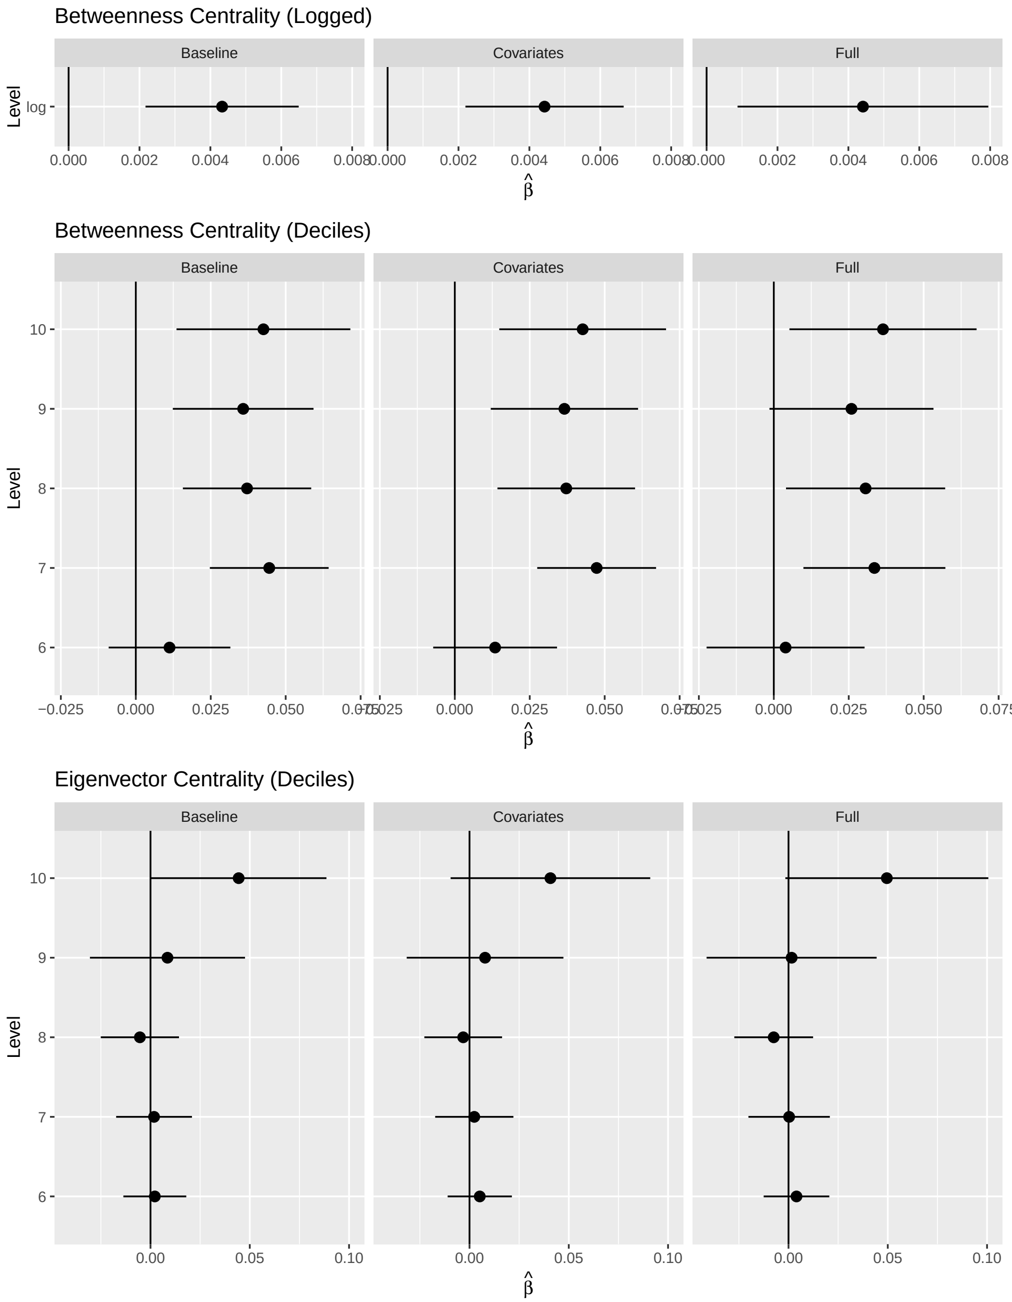

This figure shows the effects of betweenness centrality (logged and in deciles) and eigenvector centrality (in deciles) on probability of coverage using data from 1880 to 1900. *Baseline* models include lynching county, publication county, and year fixed effects. N is 6,827,543, across 2433 lynchings in 1105 counties and 4554 newspapers in 998 counties. *Covariate* models add logged population, logged urban population, logged agricultural and manufacturing output, percent black, percent urban for both lynching and publication counties. *Full* models further add dummies for degree centrality (rail lines connected to a county and its direct neighbors). N is 6,093,405, across 2283 lynchings in 1000 counties and 4330 newspapers in 947 counties. All models cluster standard errors by lynching and publication county.

Table A1: Effects of Rail Centrality and Travel Time on Coverage Rate

	(1)	(2)
Log Betweenness Centrality	0.004***	0.004*
	(0.001)	(0.002)
Eigenvector Centrality (6th decile)	0.002	0.004
	(0.008)	(0.008)
Eigenvector Centrality (7th decile)	0.0002	0.0003
	(0.010)	(0.010)
Eigenvector Centrality (8th decile)	−0.007	−0.007
	(0.010)	(0.010)
Eigenvector Centrality (9th decile)	0.005	0.001
	(0.020)	(0.022)
Eigenvector Centrality (10th decile)	0.039	0.049
	(0.023)	(0.026)
Log Non-Rail Travel Time (days)	−0.005*	−0.006**
	(0.002)	(0.002)
Log Rail Travel Time (days)	−0.017***	−0.016***
	(0.002)	(0.002)
County FE	X	X
Newspaper FE	X	X
Year FE	X	X
Covariates		X
Local Rail Network		X
N	6,827,543	6,051,956
Adjusted R^2	0.134	0.134

*p < .05; **p < .01; ***p < .001

Estimates obtained using OLS, with standard errors clustered by lynching and publication county.

Table A2: Effects of Rail Centrality and Travel Time on Coverage Rate

	(1)	(2)
Betweenness Centrality (6th decile)	0.009	0.004
	(0.010)	(0.013)
Betweenness Centrality (7th decile)	0.042***	0.033**
	(0.010)	(0.012)
Betweenness Centrality (8th decile)	0.033**	0.029*
	(0.011)	(0.014)
Betweenness Centrality (9th decile)	0.033**	0.025
	(0.012)	(0.014)
Betweenness Centrality (10th decile)	0.040**	0.036*
	(0.015)	(0.016)
Eigenvector Centrality (6th decile)	0.001	0.004
	(0.008)	(0.008)
Eigenvector Centrality (7th decile)	0.001	0.0003
	(0.010)	(0.010)
Eigenvector Centrality (8th decile)	−0.007	−0.008
	(0.010)	(0.010)
Eigenvector Centrality (9th decile)	0.006	0.0001
	(0.019)	(0.021)
Eigenvector Centrality (10th decile)	0.027	0.036
	(0.023)	(0.027)
Log Non-Rail Travel Time (days)	−0.005*	−0.006**
	(0.002)	(0.002)
Log Rail Travel Time (days)	−0.017***	−0.016***
	(0.002)	(0.002)
County FE	X	X
Newspaper FE	X	X
Year FE	X	X
Covariates		X
Local Rail Network		X
N	6,827,543	6,051,956
Adjusted R^2	0.134	0.134

*p < .05; **p < .01; ***p < .001
Estimates obtained using OLS, with standard errors clustered by lynching and publication county.

Figure A2: Anti-Lynching Discourse Over Time

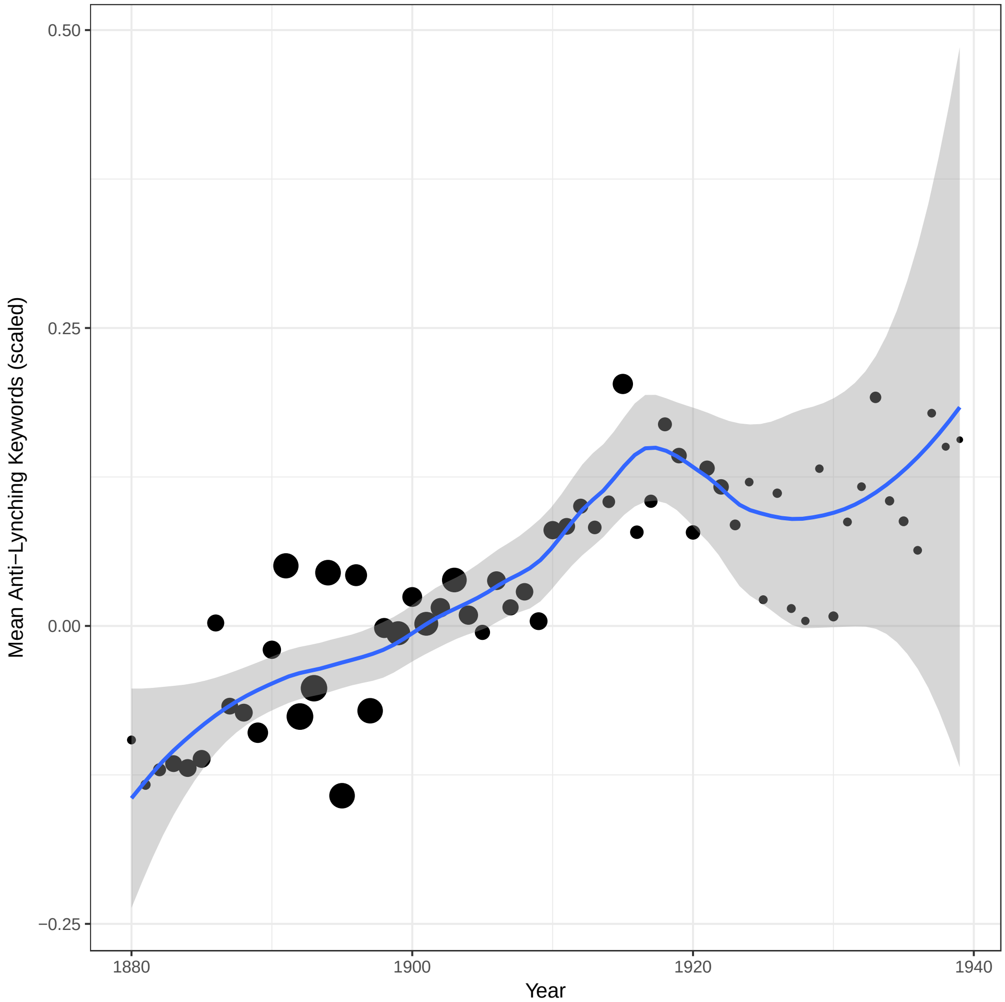

This figure shows the mean residual anti-lynching versus pro-lynching keywords (normalized) in newspaper coverage of lynchings after adjusting for the total number of keywords by year.

A.1 Distance

Increases in access to communication networks offsets distance. Absent technology, distant places should be less likely to cover local lynching events. An examination of the relationship between distance to the lynching reveals that geographic remoteness did offer obscurity to lynch mobs. As per Figure A3, while papers within 50 miles of a lynching addressed lynching in more than 36 percent of their issues in the week following the event, coverage dropped off precipitously: to 30 percent of issues at 100 miles distance, 22 percent at 500 miles, and 20 percent at 1000 miles distance. The same pattern holds even when including year, publication, and lynching-county fixed effects (Figure A4). Given that these estimates ignore the presence of transportation and communication technologies that might undo the effects of distance, it is certain that they understate the effects of distance (see Pred 1973).

Figure A3: Lynching coverage across distance (unadjusted

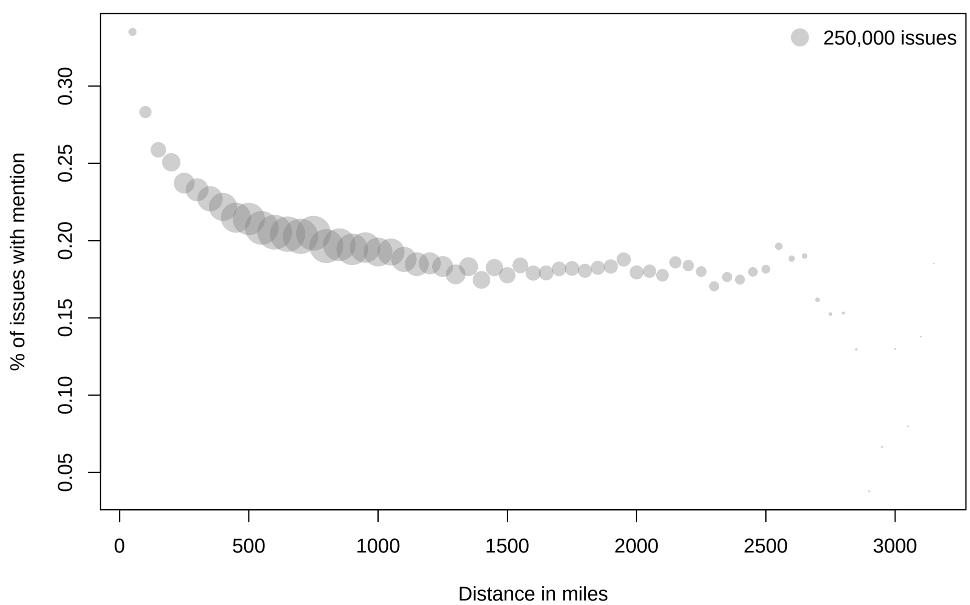

This figure shows the raw rate of lynching coverage across 50-mile distance bins. This is based on 10,212,152 dyad observations, corresponding to 3297 lynching events and 6001 newspapers between 1880 and 1911.

Figure A4: Lynching coverage across distance (adjusted

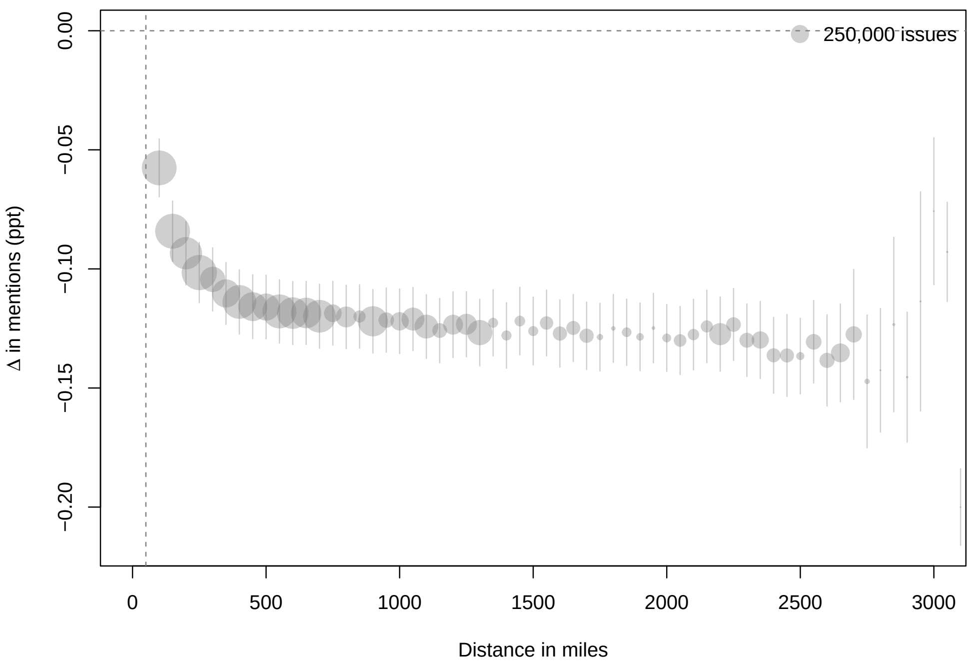

This figure shows the adjusted rate of lynching coverage across 50-mile distance bins (compared to the nearest one). The adjusted model includes year, lynching county, and publication fixed effects. This is based on 10,212,152 dyad observations, corresponding to 3297 lynching events across 1251 counties and 6001 newspapers across 1155 counties between 1880 and 1911. Standard errors are clustered by publication and lynching counties.

A.2 Reduction in Lynchings or Change in Violence

One key question about the effects of media exposure on lynching is whether the reduction in the incidence of lynching is due to actual reductions in violence or in the displacement of violence into less spectacular forms. To address this problem, I take advantage of data collected by Hagen et al. (2013). These researchers collected data on lynch mob formations in Mississippi, Georgia, and North Carolina between 1882 and 1930. While Beck et al. (2016) collect data on mob formations for a wider range of Southern states, this data is not publicly available. I merge this data to a panel of counties from these three state containing census covariates and access to circulation for the years 1880 to 1900. This makes it possible to estimate the effects of changing media exposure on the incidence of lynching through changes in the formation of lynching mobs versus direct effects on the success of lynching. To do this, I estimate the effect of access to out-of-state circulation on any attempted lynchings and any successful lynchings, with year and county fixed effects, and three sets of covariates (none, economic and demographic variables, and economic and demographic variables and local railroad network). I then use the `mediation` packaged in `R` to estimate the direct effect of access to circulation on successful lynchings and the indirect effect through lynching attempts.

Across the three different specifications, 68, 69, and 72 percent of the effect of increasing access to circulation on the incidence of lynching is indirectly through reducing mob formations. But the direct effect of access to circulation on the incidence of lynching is -0.014, -0.015, and -0.023 ($p = 0.05, 0.06, 0.06$), even when mobs form. This implies that a fifty percent increase in access to circulation yields a $0.5, 0.6, 0.9$ percentage point decrease in the probability of a successful lynching. While reductions in mob formations might be due, in part, to a displacement of public violence into more private repertoires, reductions in mob successes indicates that increasing access did actually lead to fewer deaths. This is because thwarting mobs usually involved police protection for the intended target. While it is possible that survivors of thwarted lynchings were later convicted and executed, there is little evidence that legal executions and lynchings were substitutes (Beck et al. 1989; Keil and Vito 2009)

B Robustness

In this section, I discuss and report robustness checks to the analyses reported in the main body of the paper. These sections follow the same order as presented in the paper.

B.1 Networks

In the main body of the paper, Figure 1 and Table 1 show the effects of increasing rail network centrality and decreasing railroad travel time on coverage of lynching using dyadic data. These data use all reported lynchings between 1880 and 1910, all newspaper archives, and define coverage as occurring within a window of 7 days of a lynching. Similarly, these estimates use fixed effects to partial out time-invariant attributes of lynching-counties, newspapers, and years. For both the effects of centrality and of travel time, I repeat the main analyses across different subsets of the data defined by the following attributes:

- **Lynching Sample** I consider the five different samples of lynching events discussed in section D.1. These vary across the quality of the source (academic vs. not) and the set of states considered (all vs. former slave states).

- **Time Period** In my analysis, I look at within-county variation in how lynchings were reported. To boost the number of counties with multiple lynchings, the main results uses events from 1880 to 1910. However, the railroad data I use do not change between 1900 and 1910. While, as I discuss in the paper, the changes in this time are less substantial than in the preceding two decades, I consider the robustness of my results to using samples that include lynchings between 1880 and 1900 or 1880 and 1910.

- **Archive** My data on coverage comes from four different digitized newspaper archives. One concern might be that the results are due to particular attributes of one of these archives. I consider the robustness of my results to keeping all four archives and dropping each archive in turn (five different configurations in total).

- **Window** My definition of coverage uses the presence of keywords within a window of time. In the main reported specifications, I examine coverage within a 7 day window. But as discussed below in Section D.3.2, this is a somewhat arbitrary choice. I repeat the analysis across windows of 3 to 11 days.

In total, considering all possible samples with these attributes gives 450 different subsets of the data. For each of these 450 possible samples, I estimate the effect of each network measure and rail travel time using three different model specifications. While all models include lynching-county, newspaper, and year fixed effects and cluster standard errors by lynching-county and publication-county, the set of covariates changes.

- The only covariate included is distance between the lynching county and newspaper county.

- Adding time-varying controls for logged agricultural output, logged manufacturing output, logged total population, and logged urban population for both the lynching-county and publication-county (8 variables in total).

- Further adding dummies for the degree centrality (number of direct rail links a county has with other counties) for the lynching county and the counties neighboring the lynching county (excluding links with the lynching county).

In total, this gives 1350 different possible models for each independent variable that I consider. In the following sections, I report, in broad brushstrokes, notable patterns in the robustness of these models. Finally, I repeat the analyses using the samples from the main body of the paper using the random effects within estimator (Bell and Jones 2015) in place of fixed effects. This approach makes use of the same unbiased within estimator as fixed effects but preserves efficiency. I consider random effects for only one sample, as these models are substantially more costly with respect to computational resources.

B.1.1 Dyads (Centrality)

Betweenness (Logged) The effects of logged betweenness centrality do not differ across lynching sample, newspaper archive, or time period. However, I find that while the magnitudes are the same across coverage windows (Figure B2), standard errors are larger with shorter windows, (Figure B1). Similarly, while effect magnitudes are similar across different sets of covariates (Figure B4), standard errors are larger when including local railroad degree centrality (Figure B3). This is likely because including the dummies for degree centrality cuts down on variance in betweenness, increasing standard errors.

If this is the case, using random instead of fixed effects should resolve this issue. Figure B5 shows that when using random effects, the effects of logged betweenness centrality are highly significant.

Figure B1: Effect of Network Centrality on Coverage: t statistic by coverage window

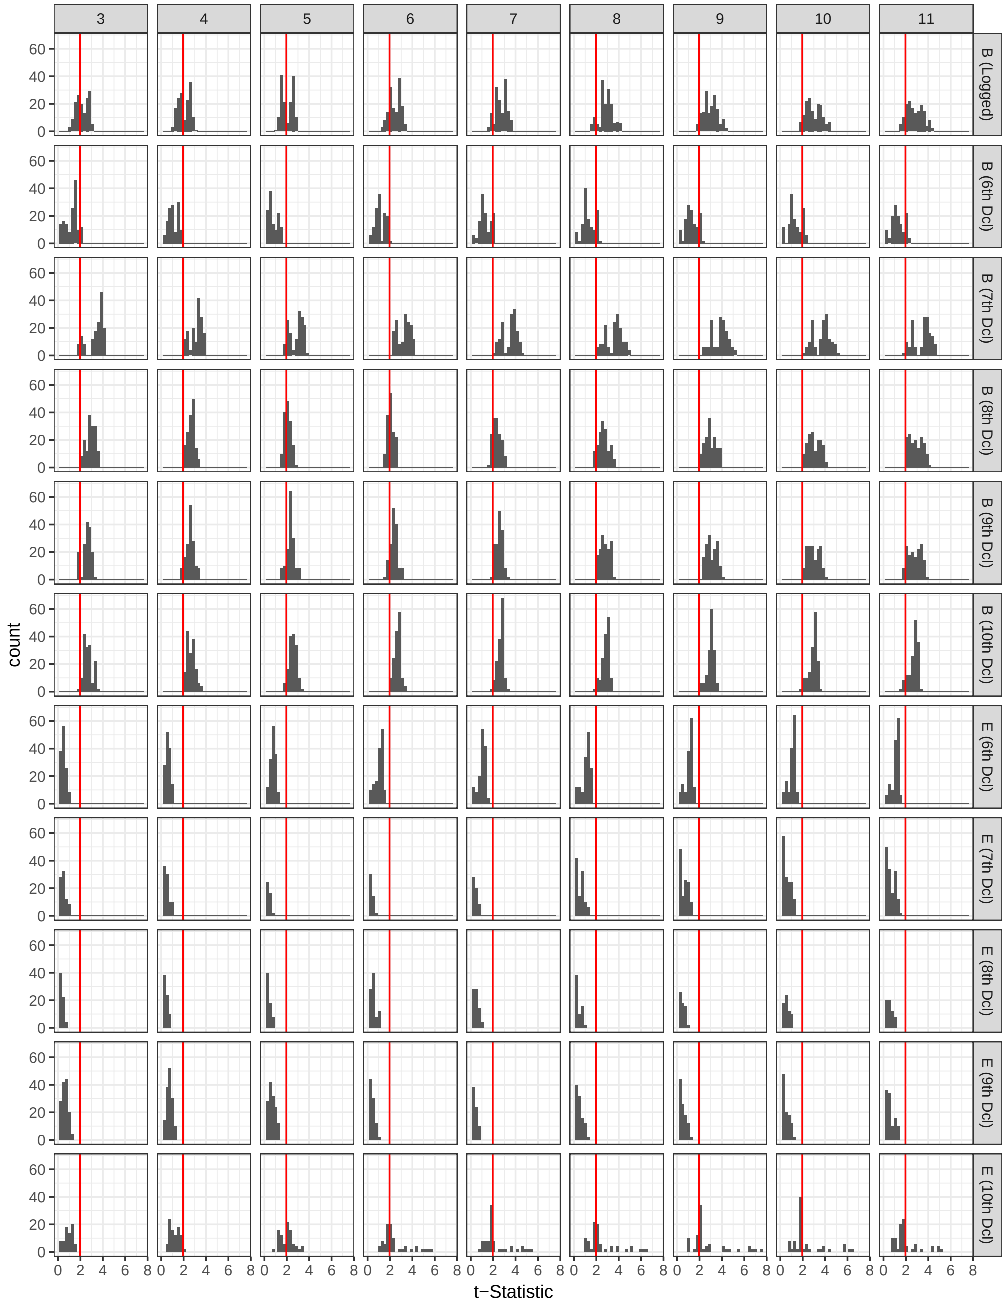

This summarizes the t statistics for the effect of betweenness and eigenvector centrality on coverage rates across 450 different samples and 3 different model specifications by coverage window used, using dyad data.

Figure B2: Effect of Logged Betweenness Centrality on Coverage: coefficients by coverage window

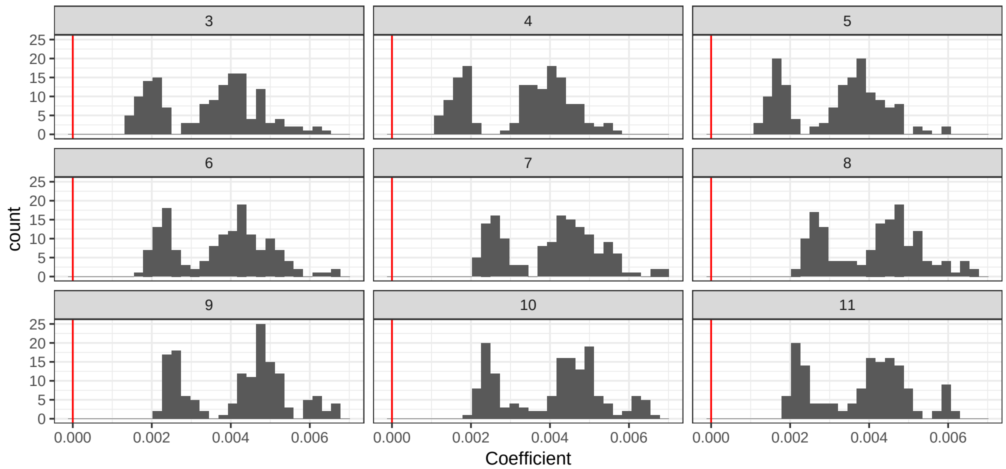

This summarizes the coefficient for the effect of logged betweenness centrality on coverage rates across 450 different samples and 3 different model specifications by coverage window used, using dyad data.

Figure B3: Effect of Network Centrality on Coverage: t statistic by covariates

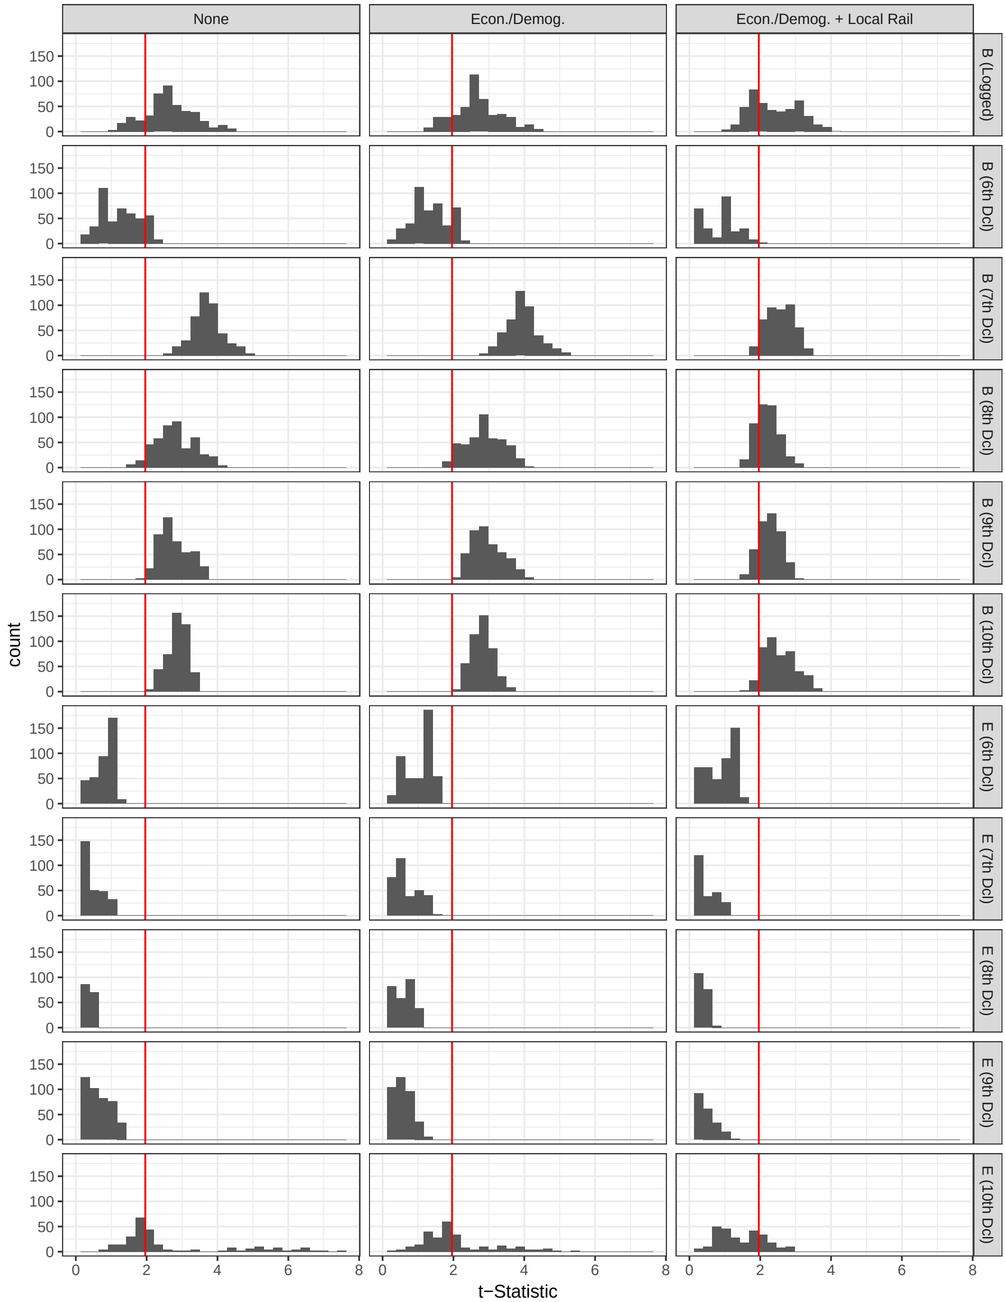

This summarizes the t statistics for the effect of betweenness and eigenvector centrality on coverage rates across 450 different samples and 3 different model specifications by covariates used, using dyad data.

Figure B4: Effect of Logged Betweenness Centrality on Coverage: coefficients by covariates

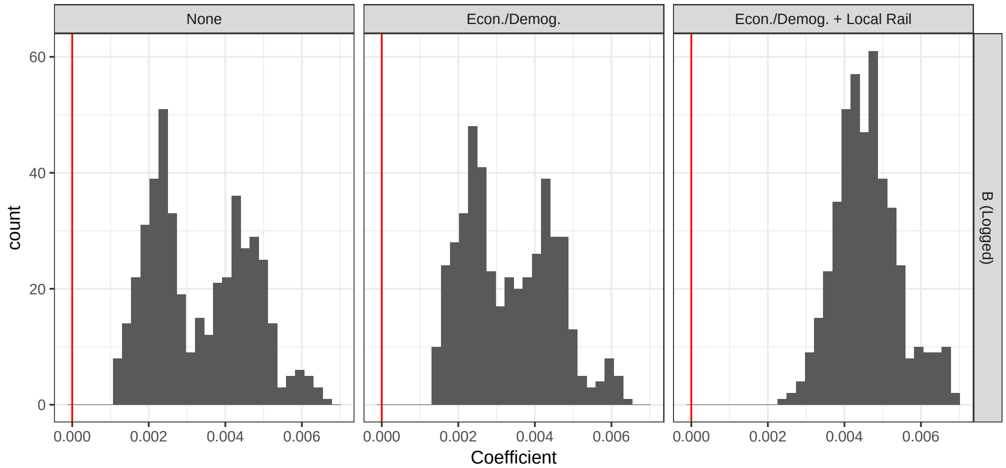

This summarizes the coefficients for the effect of logged betweenness centrality on coverage rates across 450 different samples and 3 different model specifications by covariates used, using dyad data.

Betweenness (Deciles) The effects of betweenness centrality deciles are consistent across lynching samples, archives, windows, and time periods. Moving into the 7th through 10th deciles of betweenness centrality show substantial increases in rate of coverage. While, as with logged betweenness, the inclusion of railroad degree increases the standard errors (see Figure B3), the use of random effects again clarifies that this is primarily due to the loss of variance imposed by fixed effects. When using Random Effects, these effects are highly significant (Figure B5).

Eigenvector (Deciles) The effects of eigenvector centrality are less robust. Figure B3 shows that the t statistics of Eigenvector centrality often fall below 2. But at the same time, the effects seem to vary widely in their significance (and, though not shown, in their point estimates). This appears to be due to the fact that lynchings rarely occur in places in the highest eigenvector centrality decile. In fact, when looking only at Southern lynchings, this term drops out of the analysis entirely. Consequently, the effects of eigenvector centrality are less certain. They are almost always positive, but beyond that conclusions are hard to draw.

Figure B5: Lynching coverage as a function of network centrality — Dyads (1880–1900), Random Effects

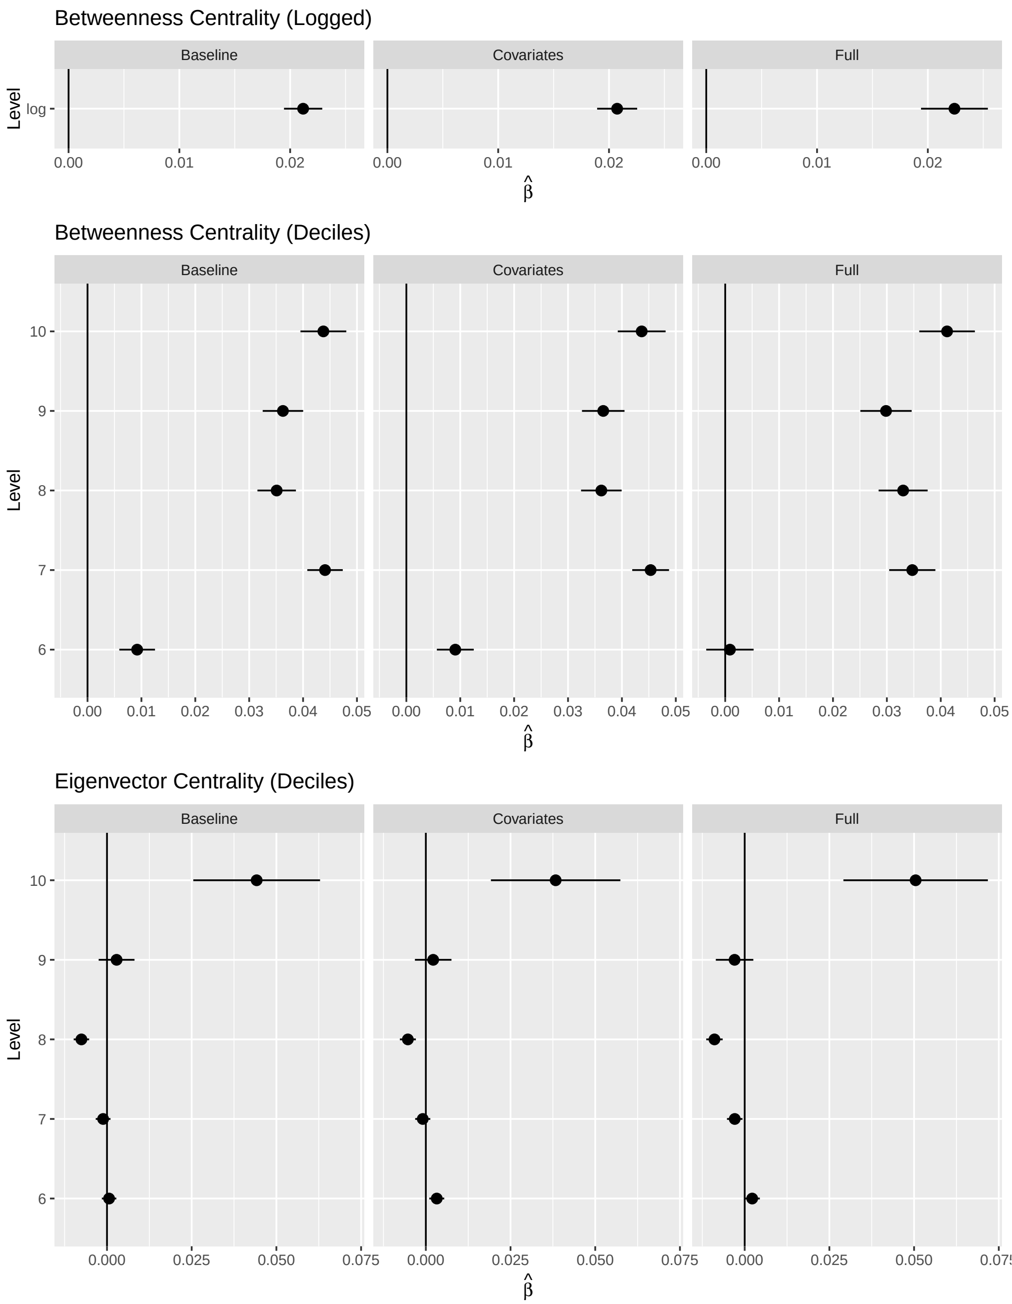

This figure shows the effects of betweenness and eigenvector centrality on probability of coverage using data from 1880 to 1900 and a coverage window of 7 days, estimated using the random effects within estimator. *Baseline* models include lynching county, publication county, and year random effects and controls for logged non-rail travel time. N is 6,827,543, across 2433 lynchings in 1105 counties and 4554 newspapers in 998 counties. *Covariate* models add logged population, logged urban population, logged agricultural and manufacturing output, percent black, percent urban for both lynching and publication counties. *Full* models further add random effects for degree centrality (rail lines connected to a county and its direct neighbors). N is 6,093,405, across 2283 lynchings in 1000 counties and 4330 newspapers in 947 counties.

B.1.2 Dyads (Time)

There are no differences in the effects of railroad travel time across windows, archives, and time periods. Again, the inclusion of dummies for local rail network increases standard errors, (Figure B6, Left Panel) but does not appear to affect point estimates. Restricting the sample to only Southern lynchings increases standard errors and also weakens the effect (Figures B7 and B8). But the use of covariates and of Southern-only lynchings only affects overall significance when conditioning on distance-time. When conditioning on non-rail travel time (i.e., including waterways), the effect of rail travel time is consistently significant. This is more of an issue for Southern states, because in much of the South water transport (particularly on the Mississippi) on rivers was important prior to the expansion of railroads.

When using random effects within estimators (Bell and Jones 2015), the effects of rail travel time are highly significant and do not depend on the set of included covariates (Figure B9).

Figure B6: Effect of Rail Travel Time on Coverage: t statistic by covariates

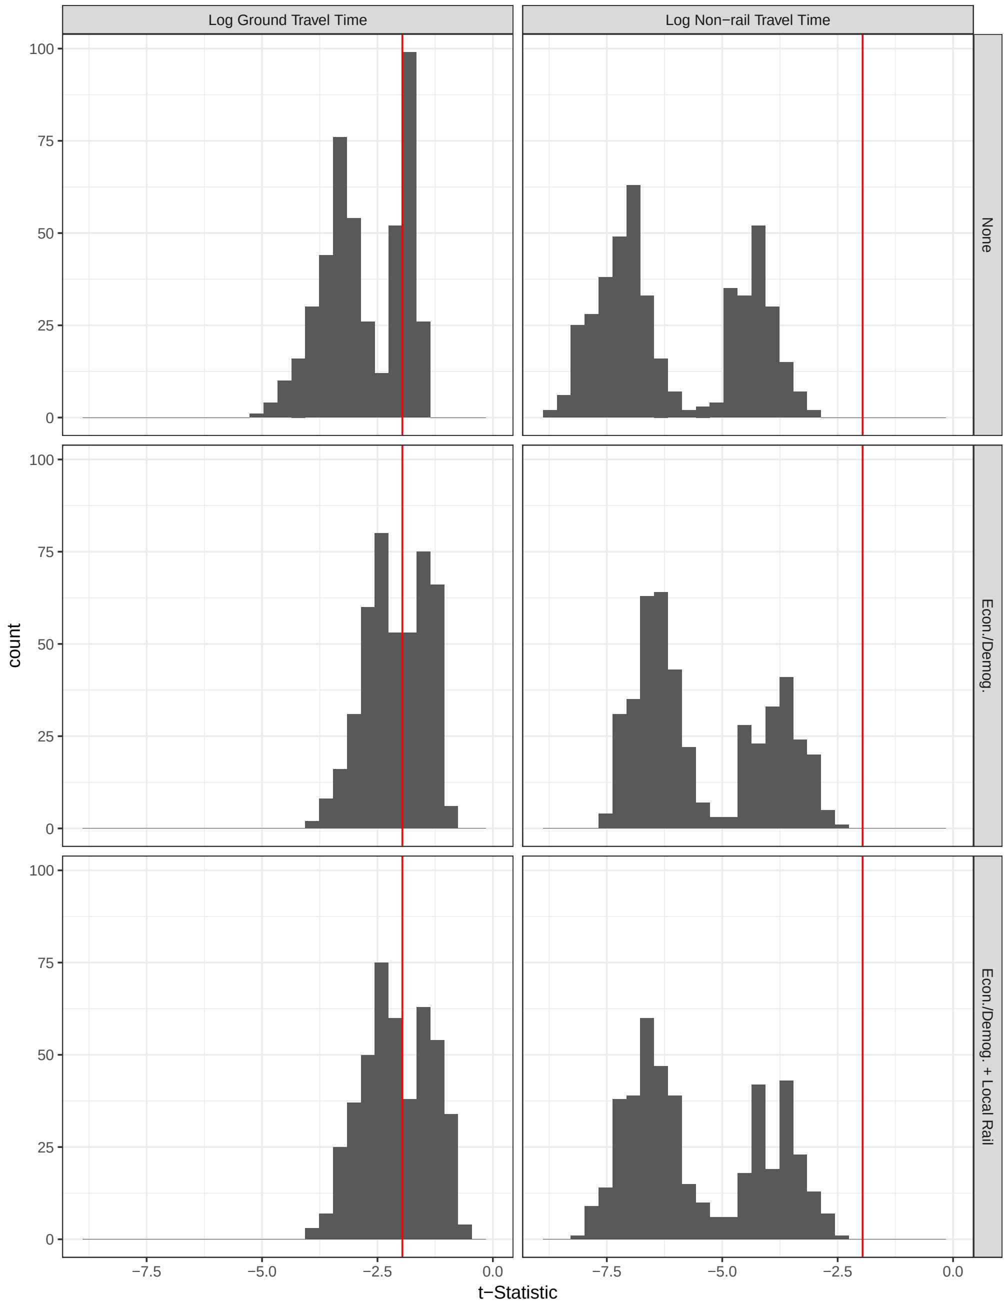

This summarizes the t statistic for the effect of logged railroad travel time on coverage rates across 450 different samples and 6 different model specifications by covariates used, using dyad data.

Figure B7: Effect of Rail Travel Time on Coverage: t statistic by lynching sample

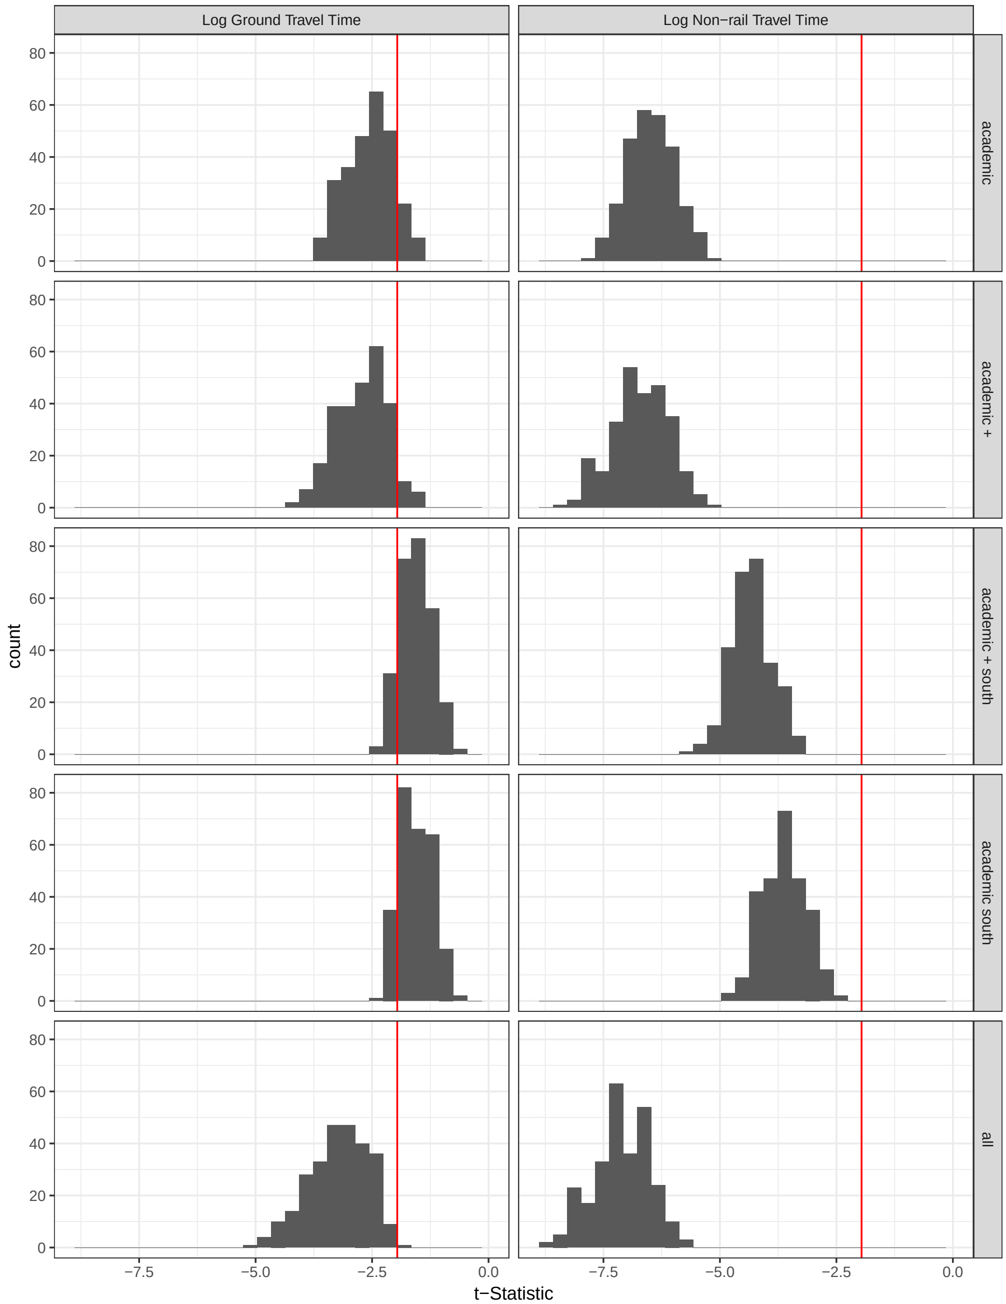

This summarizes the t statistic for the effect of logged railroad travel time on coverage rates across 450 different samples and 6 different model specifications by lynching sample, using dyad data.

Figure B8: Effect of Rail Travel Time on Coverage: coefficients by lynching sample

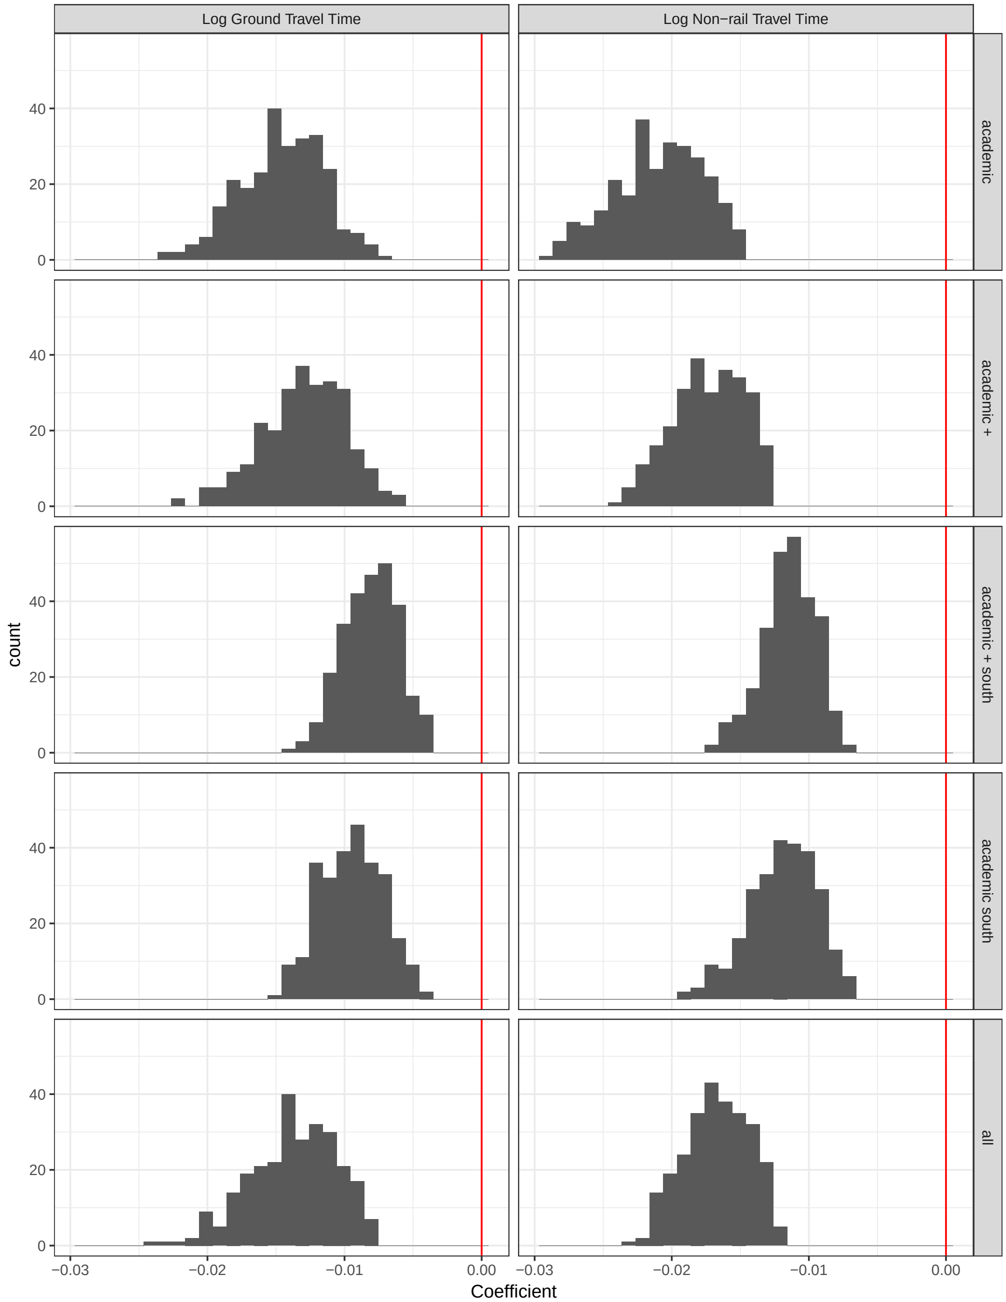

This summarizes the coefficient for the effect of logged railroad travel time on coverage rates across 450 different samples and 6 different model specifications by lynching sample, using dyad data.

Figure B9: Lynching coverage as a function of railroad travel time — Dyads (1880–1900), Random Effects

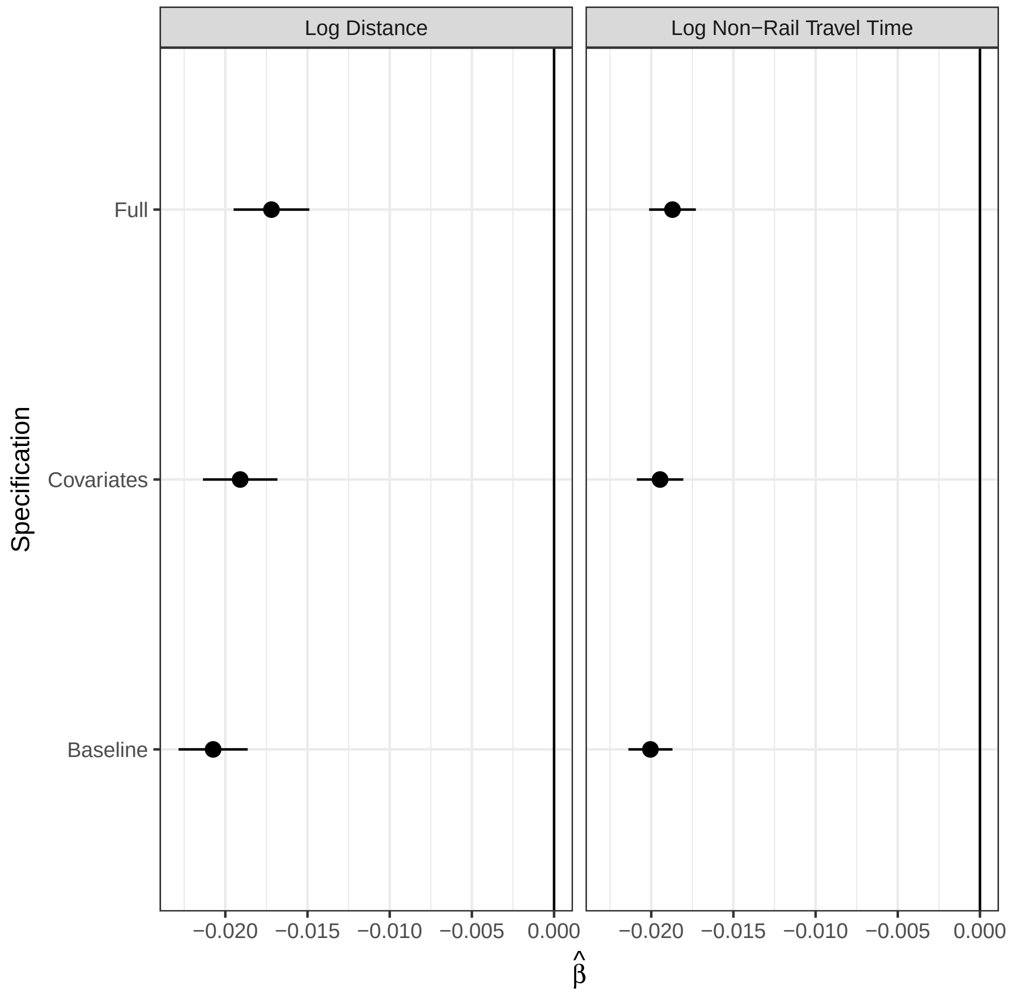

This figure shows the effects of logged railroad travel time on probability of coverage using data from 1880 to 1900 and a coverage window of 7 days, estimated using the random effects within estimator. *Baseline* models include lynching county, publication county, and year random effects and controls for logged non-rail travel time. N is 6,827,543, across 2433 lynchings in 1105 counties and 4554 newspapers in 998 counties. *Covariate* models add logged population, logged urban population, logged agricultural and manufacturing output, percent black, percent urban for both lynching and publication counties. *Full* models further add random effects for degree centrality (rail lines connected to a county and its direct neighbors). N is 6,093,405, across 2283 lynchings in 1000 counties and 4330 newspapers in 947 counties.

B.1.3 Events

One objection to my use of dyadic data is that the use of clustering or random effects for lynching counties and newspaper publication counties insufficiently accounts for non-independence of the errors. As a result, I collapse the lynching-event and newspaper-issue dyads to a panel of lynching events where coverage is measured as the fraction of realized over potential coverage.[1] For the main sample between 1880 and 1910, this yields an N of 3231 across 1239 unique counties. In equation 1, i and t represent a lynching county and year, respectively. α_{year} and α_i are year and lynching-county fixed effects.

$$\begin{aligned} Y_{it} = &\alpha_{year} + \alpha_{county-i} + \boldsymbol{\delta} * \mathbf{Centrality_{it}} + \\ &\boldsymbol{\gamma} * \mathbf{X_{it}} + \varepsilon_{county-i} \end{aligned} \tag{1}$$

When replicating the analysis using the same sample as used in the main paper, the results are similar when using this event (rather than dyad) data. While this ignores potentially useful information (such as distance to publications and publication fixed effects that capture transcription quality or coverage propensity), the results are broadly similar. Estimating equation 1 for betweenness and eigenvector centrality, Figure B10 shows that increases in both kinds of centrality lead to coverage in a greater fraction of issues. These patterns hold up even when including economic and demographic controls and local railroad construction. It is reassuring that magnitudes of these effects are similar to the dyad results: a 50 percent increase in betweenness yields a 0.1 percentage point increase in coverage; moving to the top decile in betweenness leads to a 4.0 (SE 1.9) percentage point increase, and a similar change in eigenvector centrality yields a 5.6 (SE 3.1) percentage point increase.[2]

[1] I consider coverage in the fraction of issues, fraction of unique newspapers with any coverage, and fraction of unique publication counties with any coverage.

[2] These results are similar when measuring coverage as fraction of newspapers (Figure B11) or fraction of publication counties (Figure B12).

Figure B10: Lynching coverage as a function of rail network centrality — Lynching Events

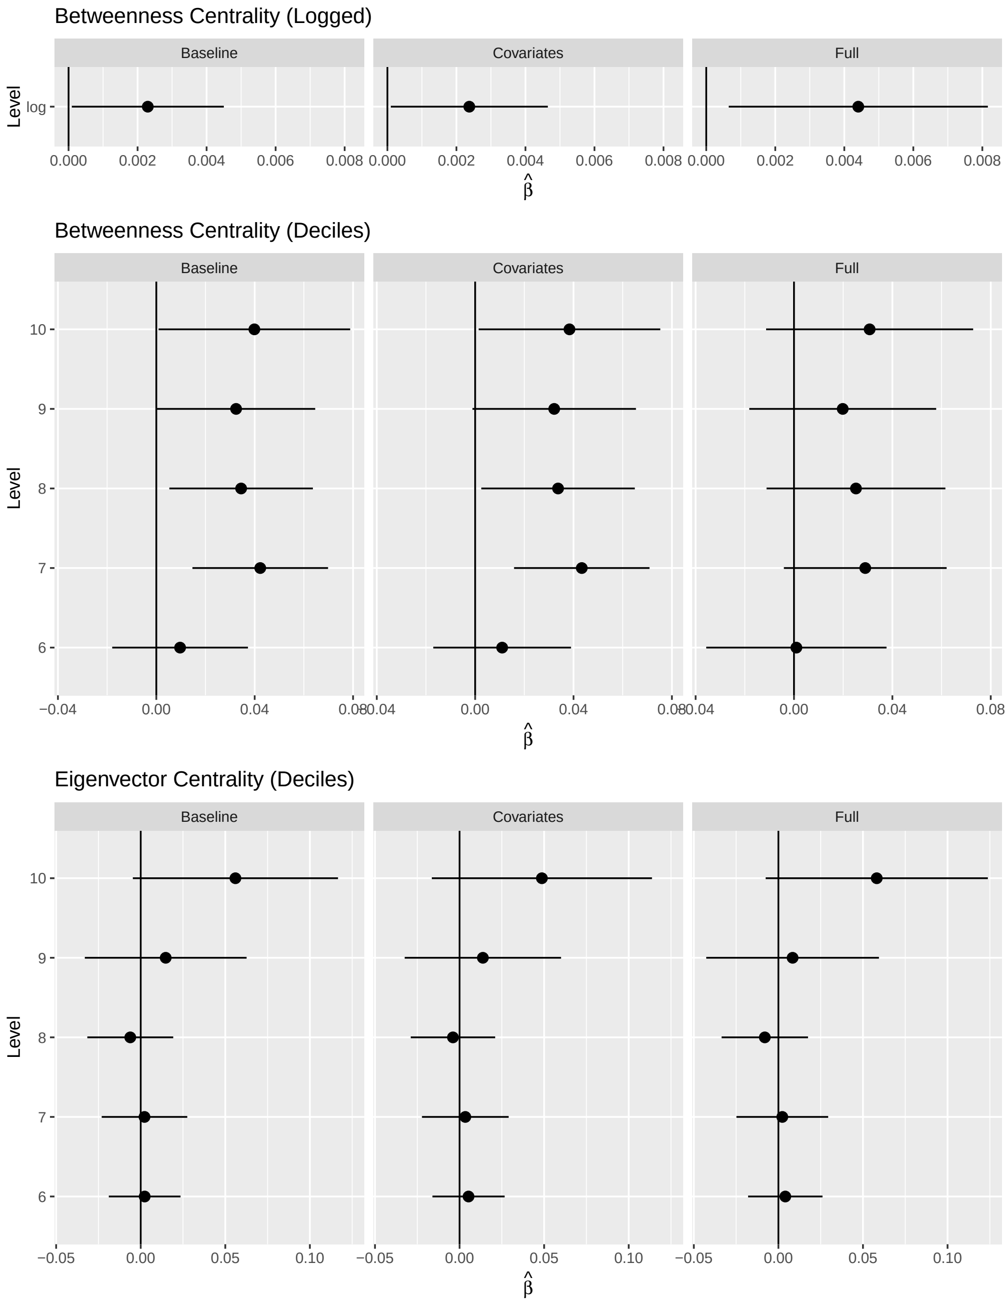

This figure shows the effects of betweenness centrality (logged and in deciles) and eigenvector centrality (in deciles) on the fraction of newspaper issues covering lynching using data from 1880 to 1910. *Baseline* models include lynching county and year fixed effects. N is 3231 lynchings in 1239 counties. *Covariate* models add logged population, logged urban population, logged agricultural and manufacturing output, percent black, percent urban for lynching counties. *Full* models further add dummies for degree centrality (rail lines connected to a county and its direct neighbors). N is 3044 in 1125 counties. All models cluster standard errors by lynching county.

Figure B11: Lynching coverage (papers) as a function of rail network centrality — Lynching Events

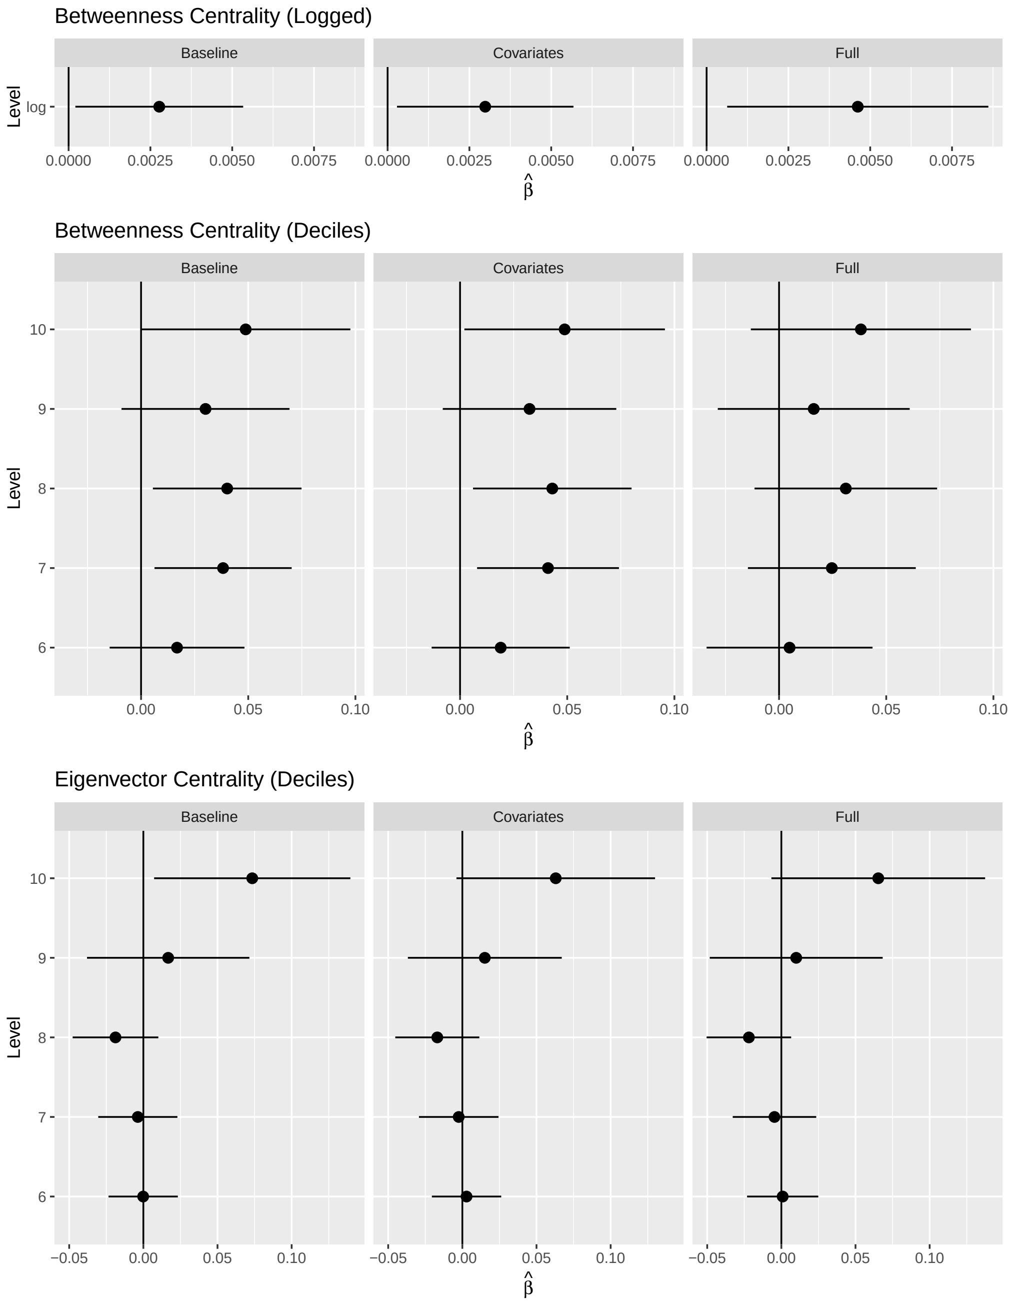

This figure shows the effects of betweenness centrality (logged and in deciles) and eigenvector centrality (in deciles) on the fraction of newspapers covering lynching using data from 1880 to 1910. *Baseline* models include lynching county and year fixed effects. N is 3231 lynchings in 1239 counties. *Covariate* models add logged population, logged urban population, logged agricultural and manufacturing output, percent black, percent urban for lynching counties. *Full* models further add dummies for degree centrality (rail lines connected to a county and its direct neighbors). N is 3044 in 1125 counties. All models cluster standard errors by lynching county.

Figure B12: Lynching coverage (publication counties) as a function of rail network centrality — Lynching Events

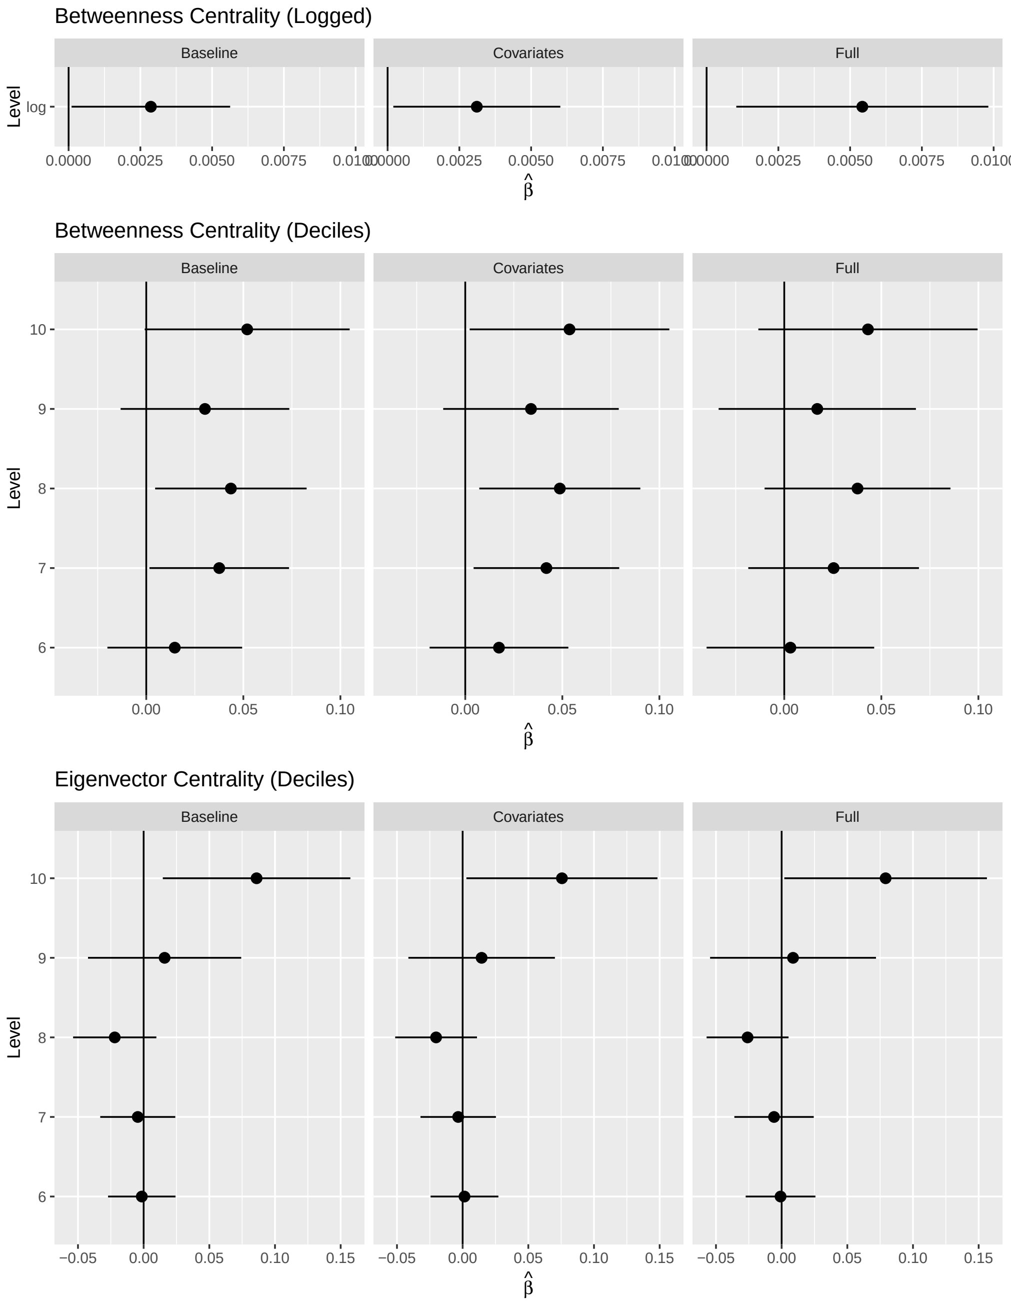

This figure shows the effects of betweenness centrality (logged and in deciles) and eigenvector centrality (in deciles) on the fraction of possible counties with coverage of lynching using data from 1880 to 1910. *Baseline* models include lynching county and year fixed effects. N is 3231 lynchings in 1239 counties. *Covariate* models add logged population, logged urban population, logged agricultural and manufacturing output, percent black, percent urban for lynching counties. *Full* models further add dummies for degree centrality (rail lines connected to a county and its direct neighbors). N is 3044 in 1125 counties. All models cluster standard errors by lynching county.

Like with the dyadic data, I also replicate these analyses across 450 different samples and 3 different sets of covariates. The effects of network centrality do not differ substantially across lynching sample, time period, or text archives.

Overall, as with the dyadic data, betweenness results are robust but eigenvector centrality results are not. Effects of betweenness are less significant when using full set of covariates, but (not shown) magnitudes of effect are the same (Figure B13. While median t statistics are significant for logged and decile betweenness, not all specifications and samples reach full significance. This is, in part, driven by smaller windows (see Figure B14). However, none of the specifications show effects in the wrong direction, and the effect magnitudes are similar. Adding covariates and using smaller windows increase in standard errors, but effects magnitudes appear to stay the same. This is encouraging.

If the concern is primarily about loss of power due to inclusion of additional covariates, we should expect that this issue is ameliorated by the use of random effects within estimators (Bell and Jones 2015).[3] Figures B15, B16, and B17 show that increases in betweenness centrality (both logged and in deciles) significantly increased coverage of lynching. The only exception is that the effect of logged betweenness is not significant at $p < 0.05$ when including controls for local railroad construction. This is not concerning as the point estimates do not change much.

The effects of eigenvector centrality are less consistent. While they are broadly positive, the standard errors and point estimates vary quite widely. Thus, while it appears that higher levels of eigenvector centrality are consistently related to higher levels of coverage, the size of this effect and its uncertainty make conclusions hard to draw.

[3]Random effects models standardize all parameters except dummy variables to help aid convergence. Thus, coefficient estimates are on different scales than fixed effects results.

Figure B13: Effect of Network Centrality on Coverage: t Statistics by Covariates, Event Data

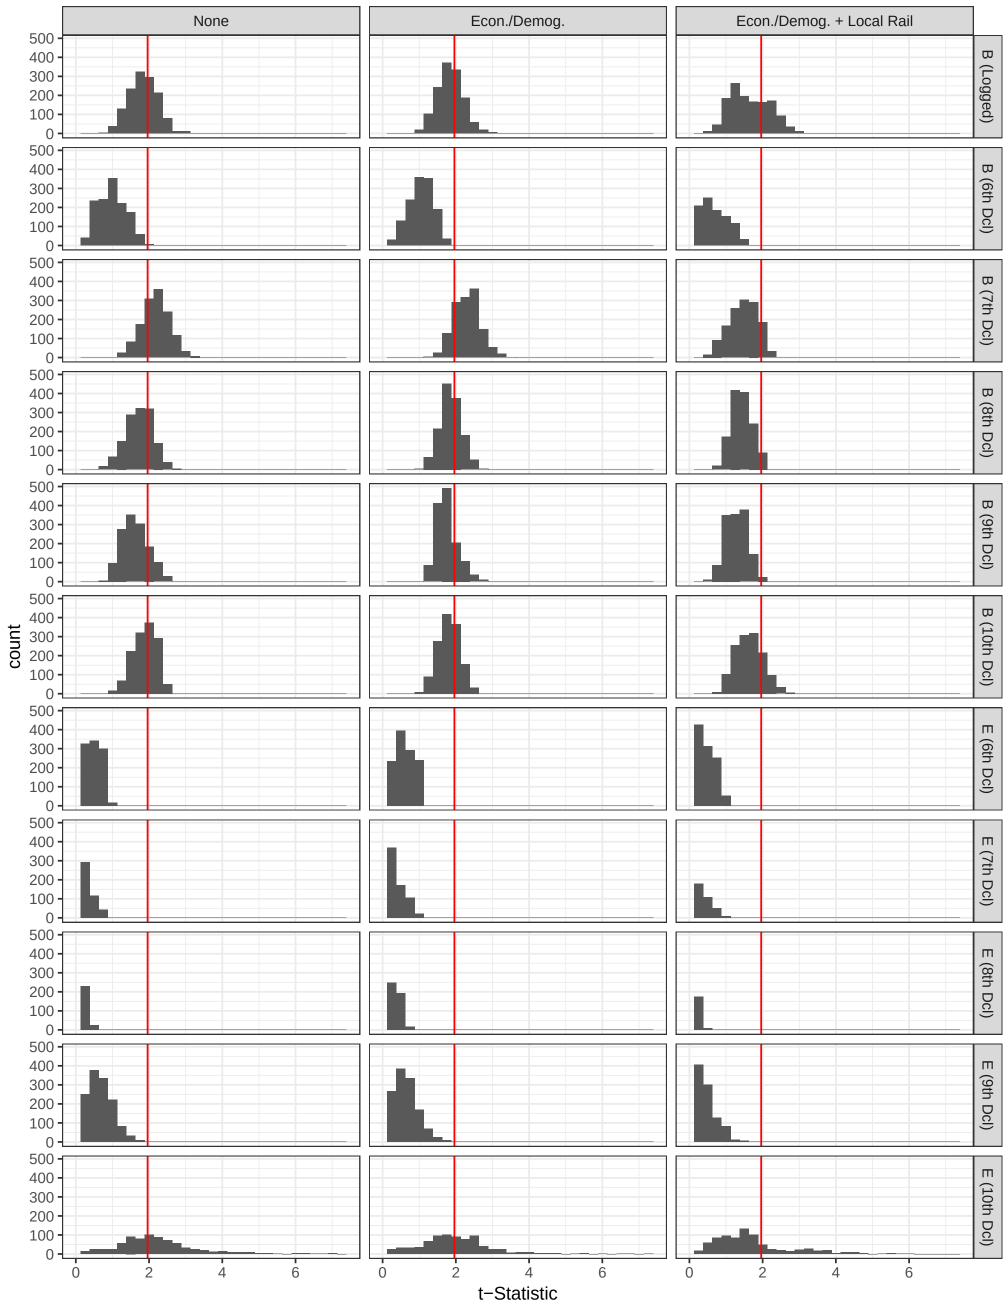

This figures summarizes the t statistics for the effects of betweenness (logged and in deciles) and eigenvector (in deciles) centrality within rail networks on fraction of coverage, across 450 different samples, 3 different sets of covariates, and 3 different measures of coverage. These are reported across different sets of covariates.

Figure B14: Effect of Network Centrality on Coverage: t Statistics by Coverage Window, Event Data

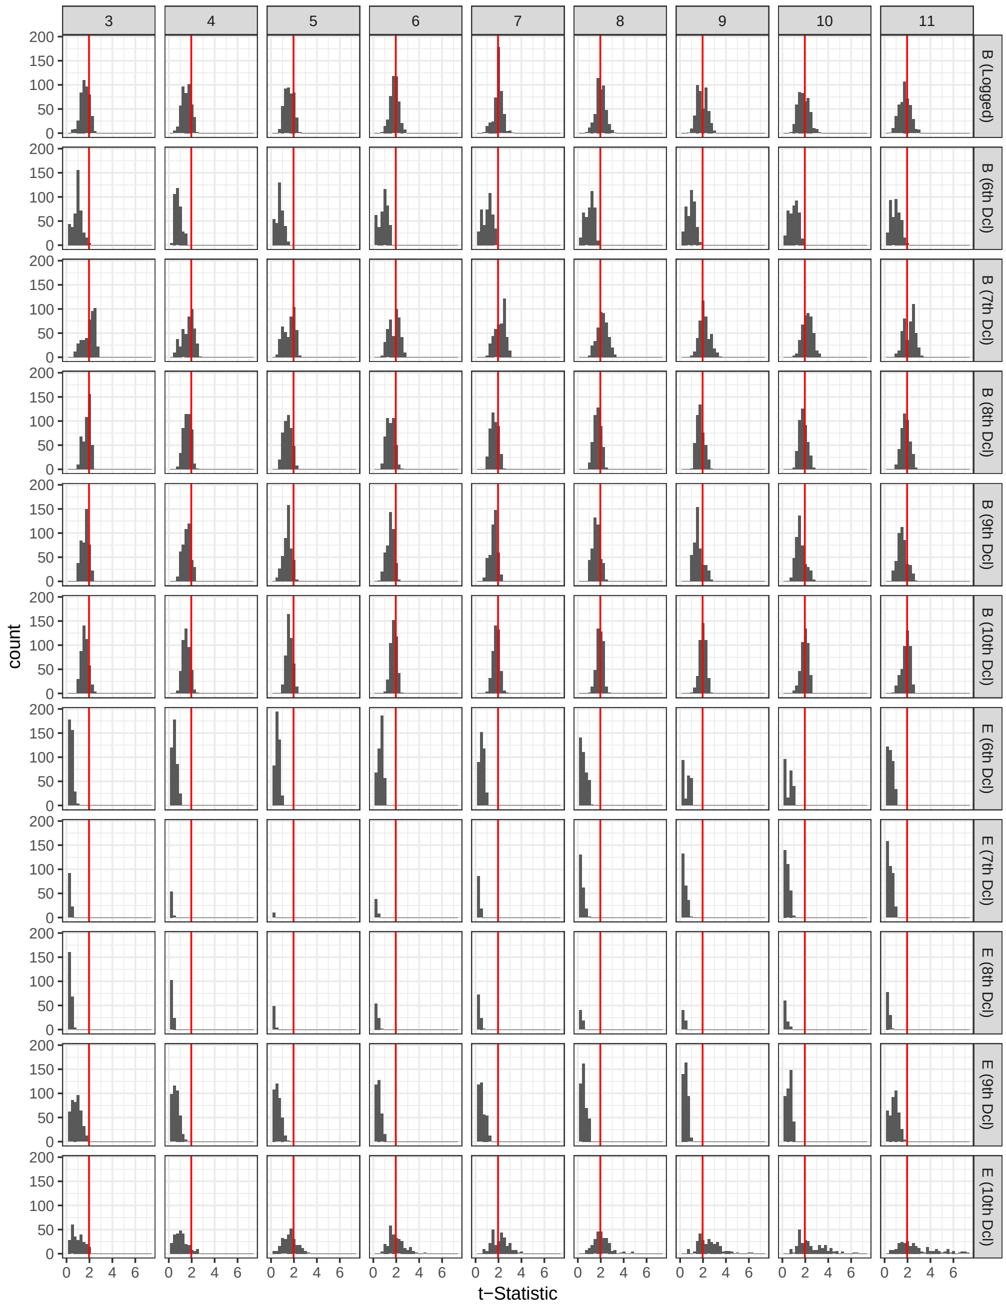

This figures summarizes the t statistics for the effects of betweenness (logged and in deciles) and eigenvector (in deciles) centrality within rail networks on fraction of coverage, across 450 different samples, 3 different sets of covariates, and 3 different measures of coverage. These are reported across different definitions of the coverage window.

Figure B15: Lynching coverage (issues) as a function of rail network centrality — Lynching Events, Random Effects

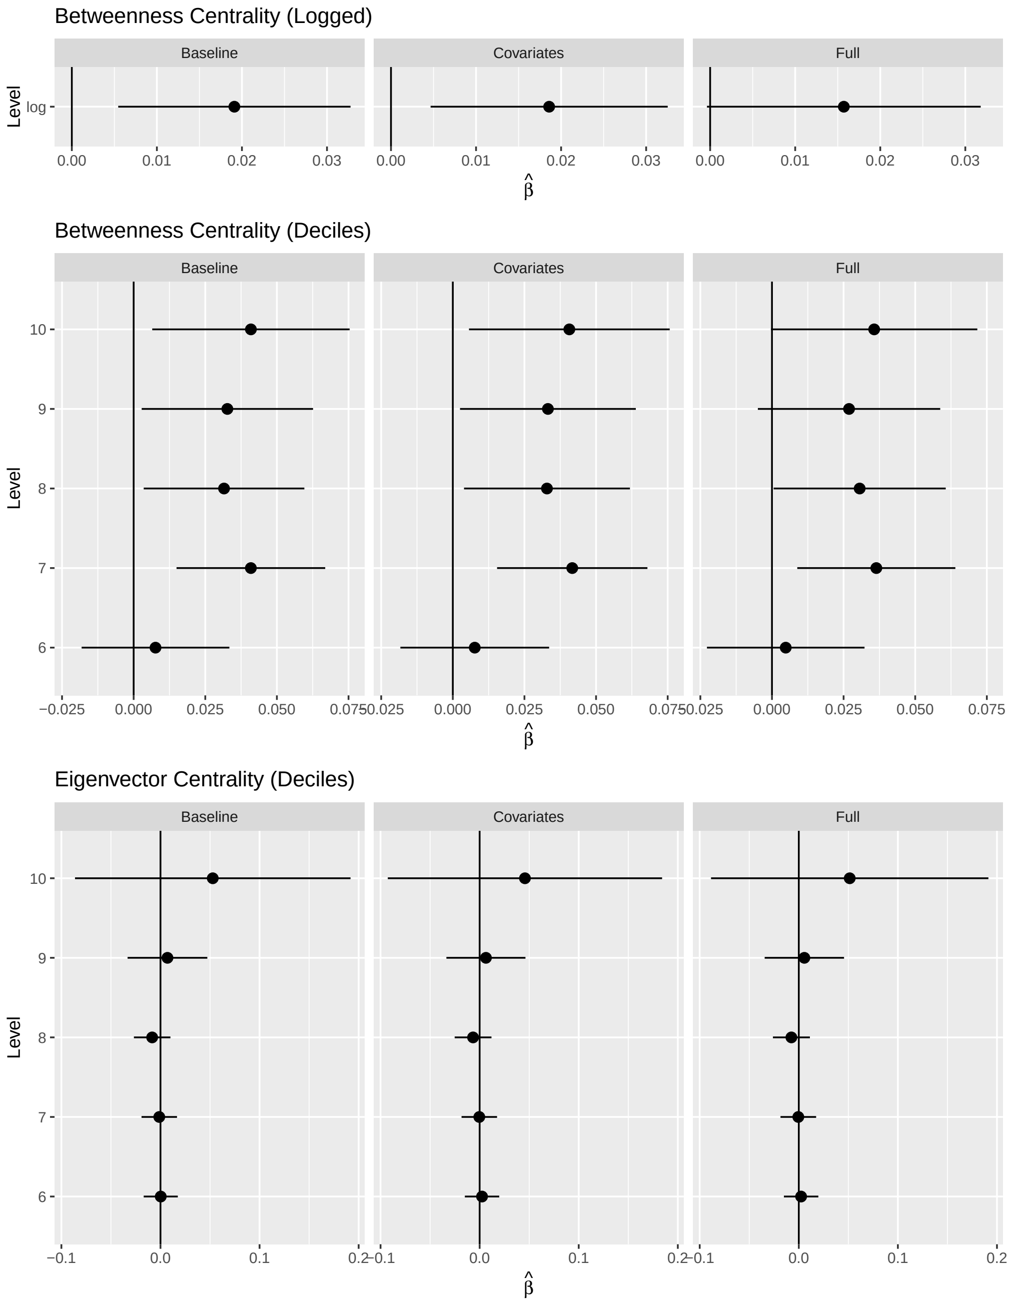

This figure shows the effects of betweenness centrality (logged and in deciles) and eigenvector centrality (in deciles) on the fraction of issues covering lynching using data from 1880 to 1900 using the random effects within estimator. *Baseline* models include lynching county and year random effects. N is 3231 lynchings in 1239 counties. *Covariate* models add logged population, logged urban population, logged agricultural and manufacturing output, percent black, percent urban for lynching counties. *Full* models further add random effects for degree centrality (rail lines connected to a county and its direct neighbors). N is 3044 in 1125 counties.

Figure B16: Lynching coverage (papers) as a function of rail network centrality — Lynching Events, Random Effects

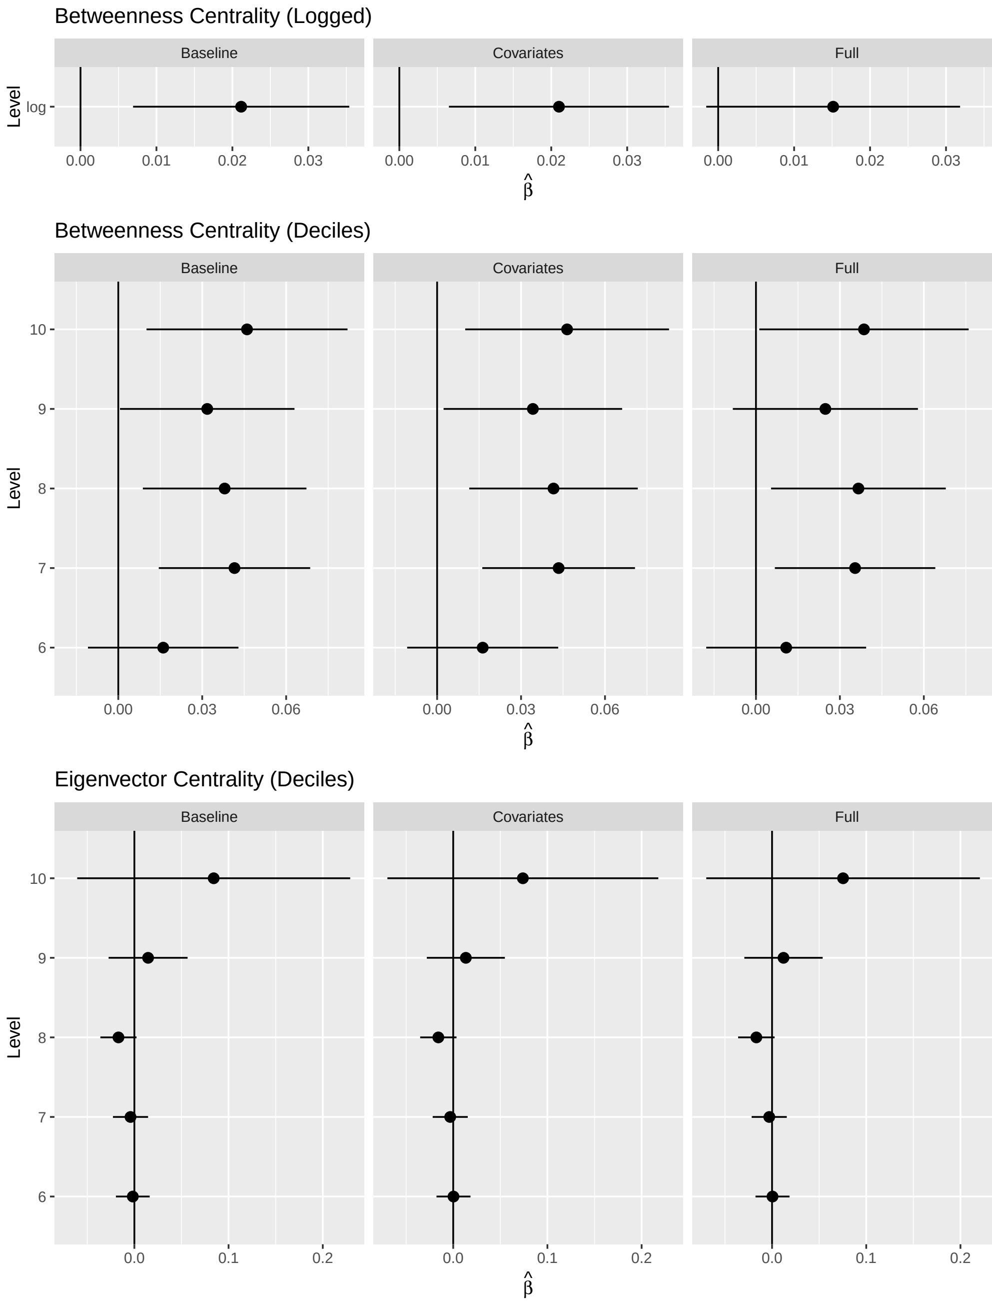

This figure shows the effects of betweenness centrality (logged and in deciles) and eigenvector centrality (in deciles) on the fraction of newspapers covering lynching using data from 1880 to 1900 using the random effects within estimator. *Baseline* models include lynching county and year random effects. N is 3231 lynchings in 1239 counties. *Covariate* models add logged population, logged urban population, logged agricultural and manufacturing output, percent black, percent urban for lynching counties. *Full* models further add random effects for degree centrality (rail lines connected to a county and its direct neighbors). N is 3044 in 1125 counties.

Figure B17: Lynching coverage (counties) as a function of rail network centrality — Lynching Events, Random Effects

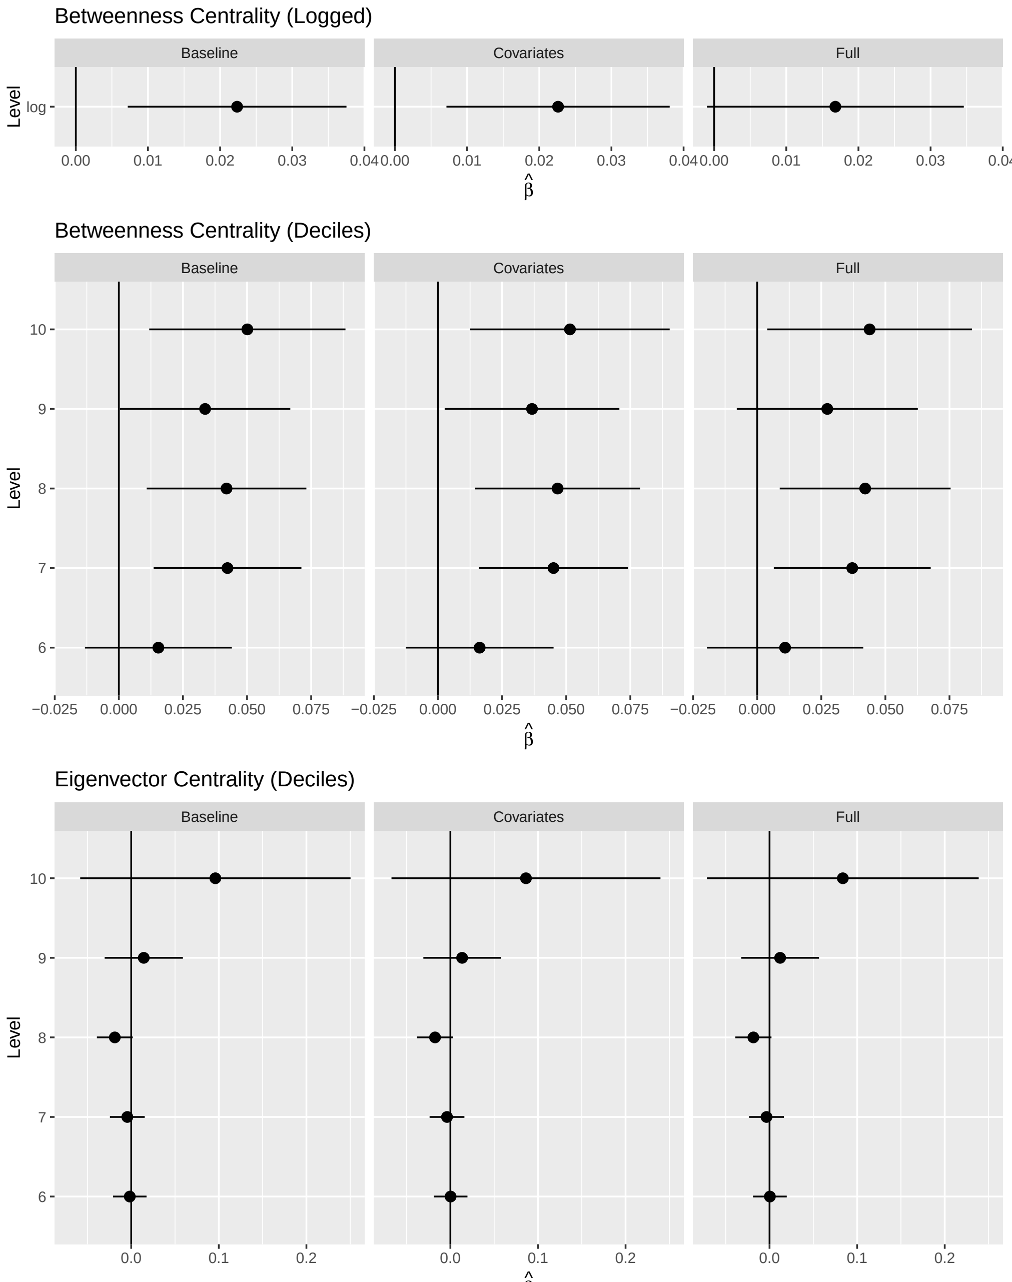

This figure shows the effects of betweenness centrality (logged and in deciles) and eigenvector centrality (in deciles) on the fraction of publication counties covering lynching using data from 1880 to 1900 using the random effects within estimator. *Baseline* models include lynching county and year random effects. N is 3231 lynchings in 1239 counties. *Covariate* models add logged population, logged urban population, logged agricultural and manufacturing output, percent black, percent urban for lynching counties. *Full* models further add random effects for degree centrality (rail lines connected to a county and its direct neighbors). N is 3044 in 1125 counties.

B.1.4 Effects Across Distance

The results in the paper and this appendix show that increasing centrality in rail networks increased the coverage of lynching, they do not reveal how this increase in coverage was distributed geographically. It could be that all increases in coverage were concentrated locally, and thus increasing centrality in rail networks did not create more geographically distant audiences. I address this by replicating the analyses reported in the main body of the paper[4] but interacting centrality and distance. Figure B18 shows the heterogeneous effects of betweenness and eigenvector centrality on coverage across 300-mile distance bins for the dyad data. In both cases, the effects of increased centrality is not concentrated locally but across a range of distances.

The same is true when examining the event-level data. Figures B19, B20, and B21 show the effect of betweenness and eigenvector centrality on coverage across distances. Because the data are aggregated by lynching event, it is impossible to include an interaction with distance. Instead, I generate dependent variables that correspond to the coverage rates in one hundred overlapping distance windows. The numerator and denominator for the coverage rate in each window were created by smoothing the coverage rates between 0 and 3000 miles using a kernel density estimator with a bandwidth of 300 miles. Here, higher betweenness centrality increases coverage across all distances, whereas the effects of eigenvector centrality appear to be concentrated within 1000 miles of the event.

[4] All reported lynchings between 1880 and 1910, all archives, 7 day window, and no covariates.

Figure B18: Lynching coverage across distance as a function of rail network centrality — Dyads

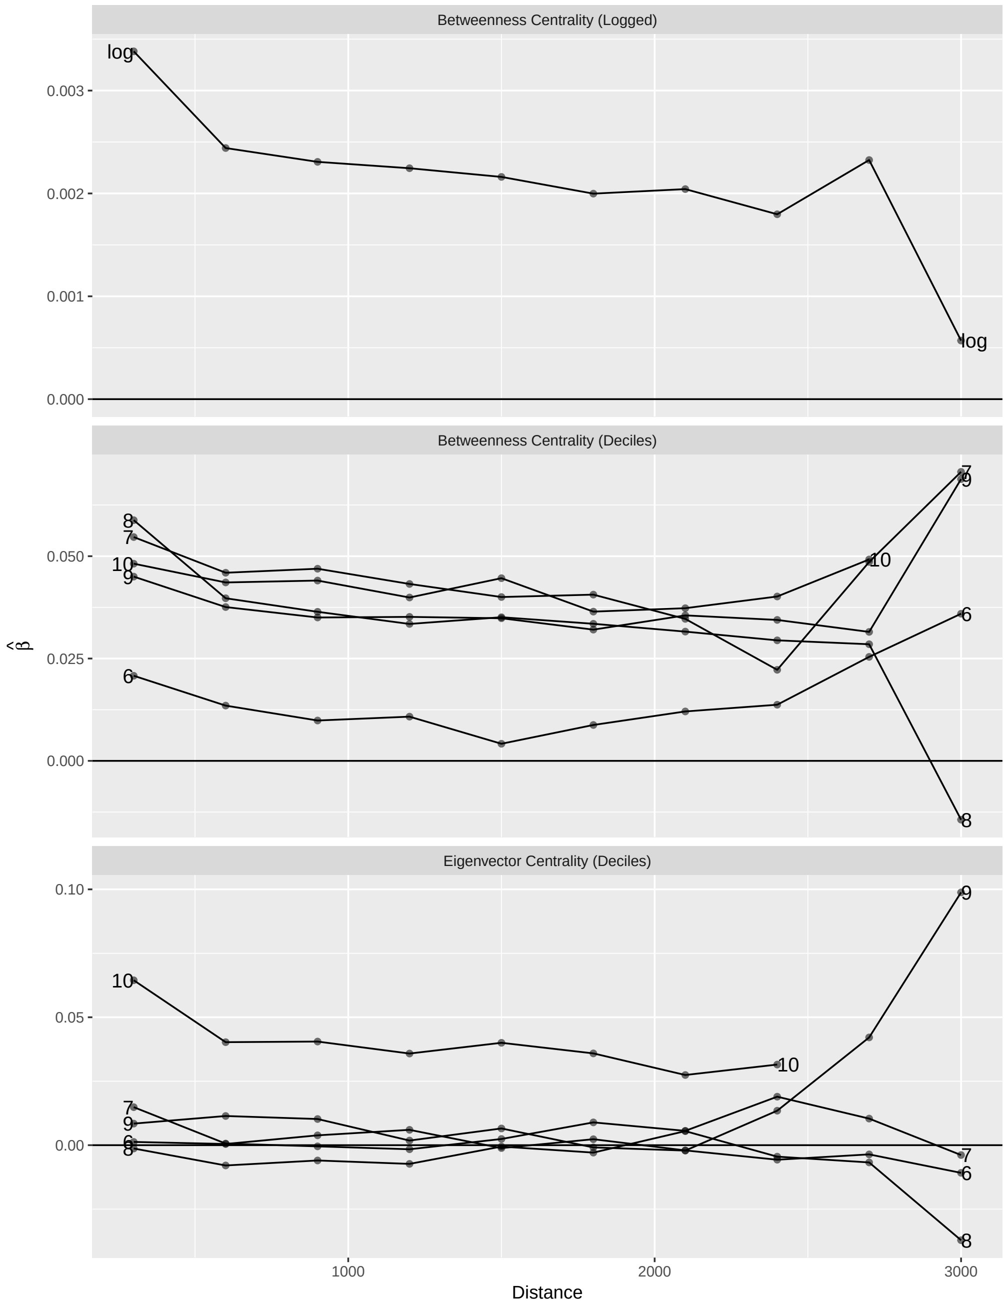

This figure shows the effects of betweenness centrality (logged and in deciles) and eigenvector centrality (in deciles) on probability of coverage using data from 1880 to 1910 across 300-mile distance bins. These are *Baseline* models which include lynching county, publication county, and year fixed effects. N is 9,934,593, across 3231 lynchings in 1239 counties and 5873 newspapers in 1138 counties. Figures only show point-estimates without standard errors. Lines are labeled to indicate the decile number (where appropriate).

Figure B19: Lynching coverage across distance as a function of log betweenness centrality — Lynching Events

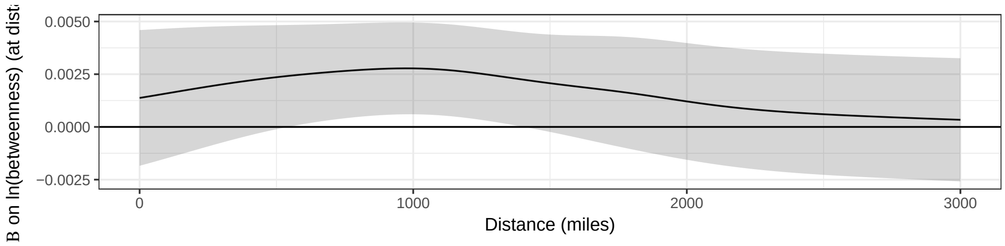

This figure shows the effects of betweenness centrality (logged) on the fraction of newspaper issues covering lynching using data from 1880 to 1910 using a *baseline* model with lynching county and year fixed effects. N is 3231 lynchings in 1239 counties. This shows the results of 100 regressions of coverage rates across moving distance bins. All models cluster standard errors by lynching county.

Figure B20: Lynching coverage across distance as a function of betweenness centrality deciles — Lynching Events

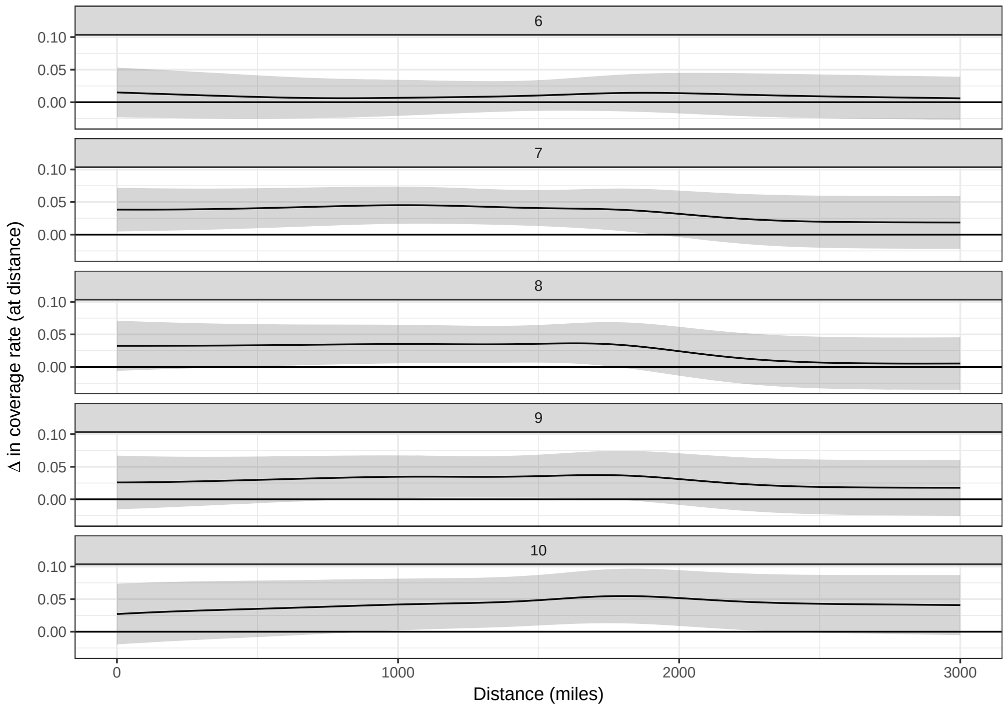

This figure shows the effects of betweenness centrality (deciles) on the fraction of newspaper issues covering lynching using data from 1880 to 1910 using a *baseline* model with lynching county and year fixed effects. N is 3231 lynchings in 1239 counties. This shows the results of 100 regressions of coverage rates across moving distance bins. All models cluster standard errors by lynching county.

Figure B21: Lynching coverage across distance as a function of eigenvector centrality deciles — Lynching Events

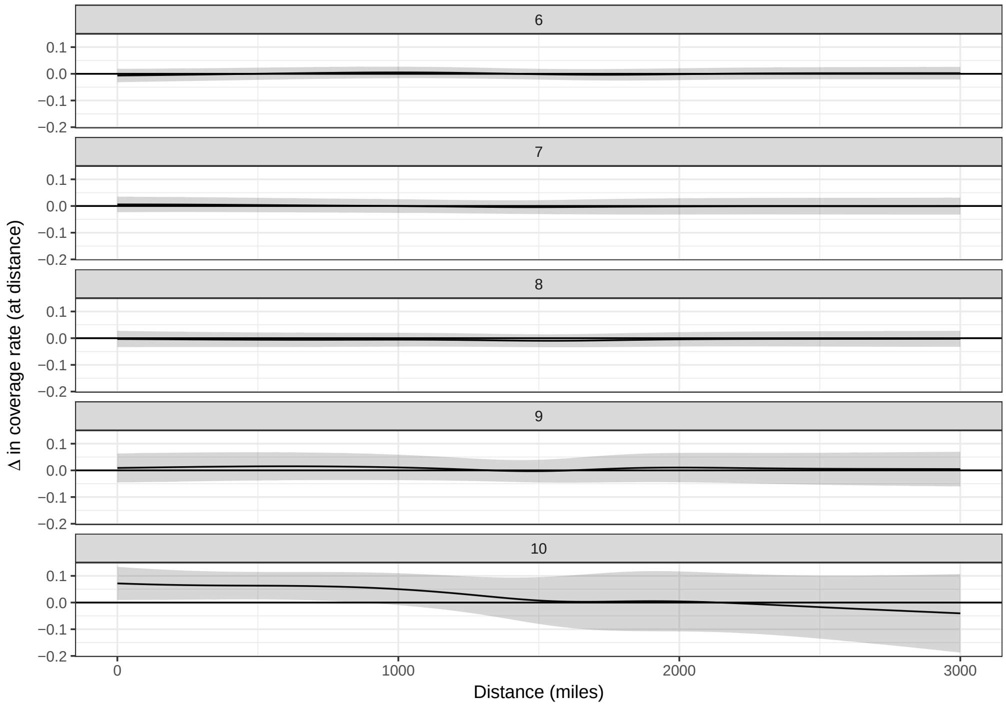

This figure shows the effects of eigenvector centrality (deciles) on the fraction of newspaper issues covering lynching using data from 1880 to 1910 using a *baseline* model with lynching county and year fixed effects. N is 3231 lynchings in 1239 counties. This shows the results of 100 regressions of coverage rates across moving distance bins. All models cluster standard errors by lynching county.

B.2 Coverage

In this section, I consider the robustness of the relationship between distance from a lynching and the degree to which press coverage is critical. In the main body of the paper, I include data on all lynchings in my data between 1880 and the 1930s, newspapers from all archives, and using coverage within a window of 7 days. I also use four specifications that attempt to capture different sets of confounding variables.

1. The measure of discourse that I use is zero when no keywords are found. More distant coverage may be briefer and simply use fewer words, and lead the discourse score to tend toward zero. To address this possibility, I include in all specifications dummies for the total number of keywords found in an article.

2. Lynching events that were spectacles and carried out in a horrific manner received more coverage—and thus were likely to be reported at greater distance—and more easily condemned (Seguin 2016). This would produce a pattern where distant coverage is more critical, but only as a result of unobserved differences between lynchings. But the effects of distance persist when adding fixed effects for each unique lynching event.

3. Newspapers may have had fairly consistent positions on lynching (for instance, the Chicago Tribune committed itself to condemning lynching). While this is not a particular problem for my argument (as long as newspapers with consistent anti-lynching positions were farther away from most lynchings, and newspapers with pro-lynching positions were closer to most lynchings), it would be strong evidence for the importance of distance if, *within the same newspaper* coverage was more critical of more distant events. Including these fixed effects rendered the effects less significant, though the direction of the effect was the same.

4. Finally, the relationship between anti-lynching discourse and distance could be a function of time. In later years, trends in both technology and public attitudes might produce a correlation between distant coverage and criticism of lynching. However, I find that distance continues to predict more critical coverage when including year fixed effects instead of lynching-event and newspaper fixed effects.[5]

I report the results of these tests in the main body of the paper. I also drop newspaper fixed effects in favor of an indicator for whether a paper was African American. While African American papers were more critical (effect size, p value), this does not explain the effect of distance. To further probe the robustness of my results, I considered several different subsamples:

- **Lynching Sample** I consider the five different samples of lynching events discussed below in Section D.1. These vary across the quality of the source (academic vs. not) and the set of states considered (all vs. former slave states).

- **Archive** My data on coverage comes from four different digitized newspaper archives. One concern might be that the results are due to particular attributes of one of these archives. I consider the robustness of my results to keeping all four archives and dropping each archive in turn (five different configurations in total).

- **Window** My definition of coverage uses the presence of keywords within a window of time. In the main reported specifications, I examine coverage within a 7 day window. But as discussed below in Section D.3.2, this is a somewhat arbitrary choice. I repeat the analysis across windows of 3 to 11 days.

For each of these 225 different samples, I also considered different model specifications (all models include fixed effects for the number of matching keywords on a page):

- Different sets of additional fixed effects: None; Lynching-event fixed effects; Lynching-event and Newspaper fixed effects; Year fixed effects

- Linear vs. Logged Distance

- Two different measures of lynching discourse. As per section D.4.1, the keywords meant to capture the "innocence" anti-lynching discourse did a poor job. Thus, I use one version of my dictionary index that includes the 'innocence' keywords and one that excludes them.

[5] Lynching-event fixed effects also eliminate any differences in coverage due to changes over time.

In total, this gives a total of 16 different model specifications.

The only lack of robustness comes from the use of the newspaper fixed effects (Figure B22). This is a very hard test. What this shows is that the newspapers that were further from lynchings were systematically different in their coverage of lynching. This does not refute my argument at a general level. Remarkably, despite losing significance, even under this specification, distance was still associated with more negative coverage (albeit not significant).

Excluding specifications that include newspaper fixed effects, the effects of distance was significant at $p < 0.05$ for nearly every sample and specification. The effects were consistent across the use of linear or log distance, the coverage window, the use of archives, and different lynching samples (for the last of these, see Figure B23). Moreover, the inclusion or exclusion of "innocence" keywords made little difference (see top versus bottom panel).

Figure B22: Association between Distance and Lynching Discourse: t Statistics by model specification

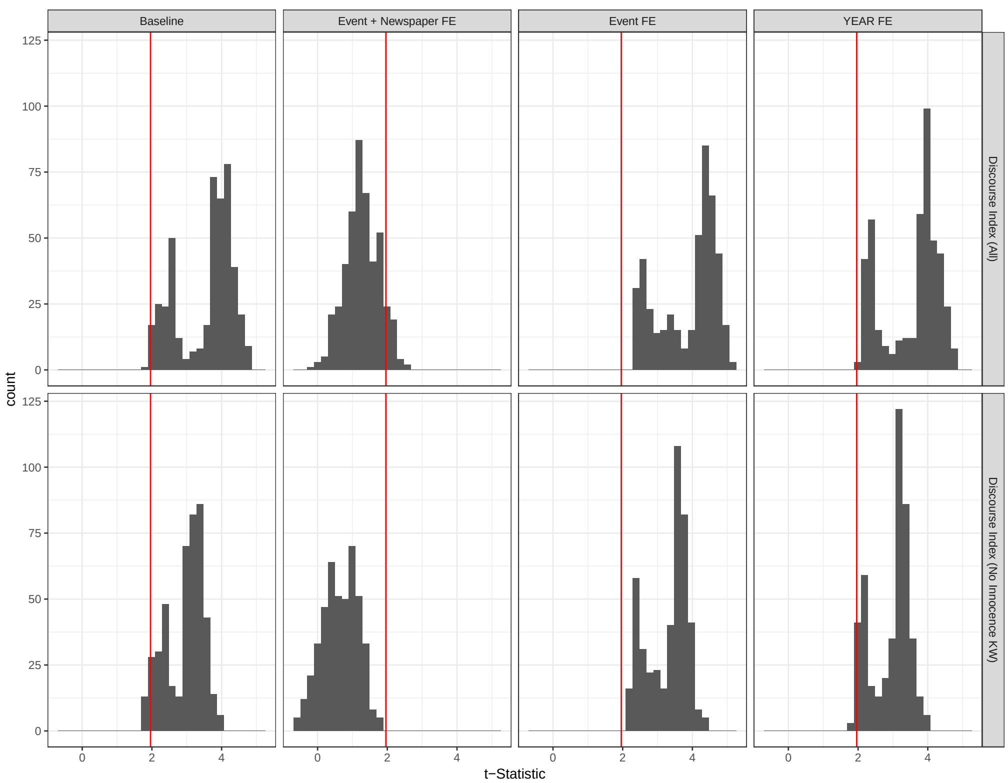

This figure summarizes the t statistics of regressions of the two different lynching discourse indices on distance from the lynching, across 450 different samples and 16 different model specifications.

Figure B23: Association between Distance and Lynching Discourse: t Statistics by lynching sample

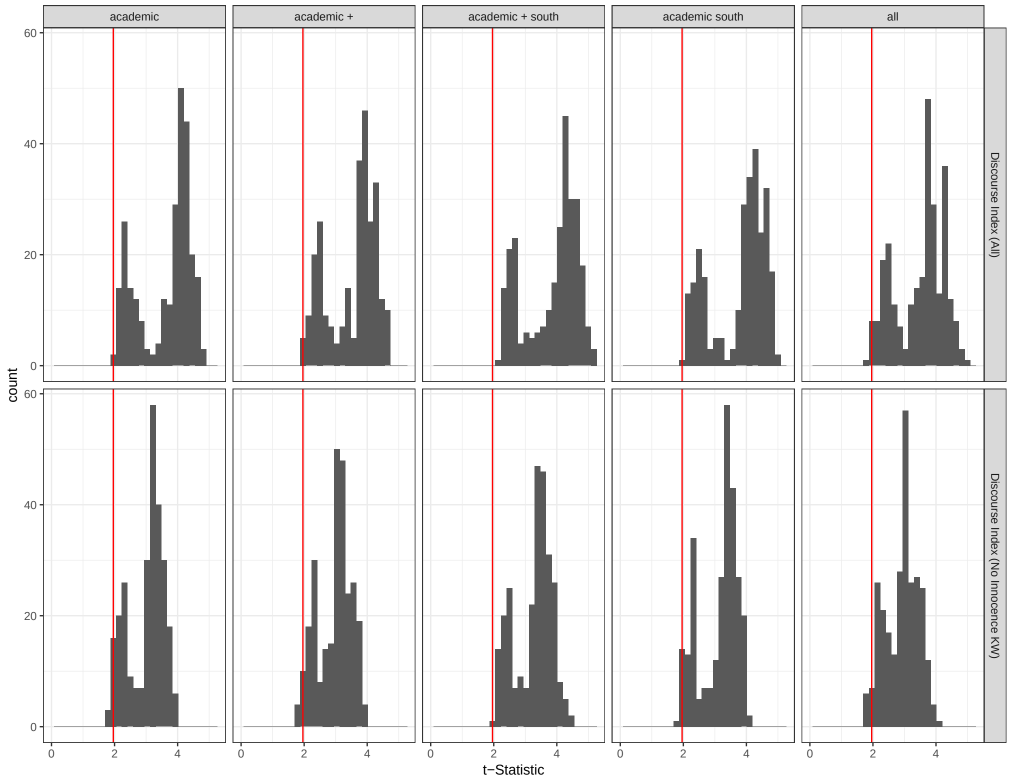

This figure summarizes the t statistics of regressions of the two different lynching discourse indices on distance from the lynching, across 450 different samples and 12 different model specifications. This excludes specifications that include newspaper fixed effects.

B.3 Lynchings

In this section, I show the robustness of my results on the relationship between exposure to publicity (media access) and the incidence of lynching. This works in two parts. First, I show that the main findings from the paper actually pertain to greater publicity. Second, I show that these results are robust to several possible specifications. Finally, I show that these results do not depend on the use of a linear probability model.

B.3.1 Exposure to Publicity

In the main body of the paper, I report the effect of "access" to (daily) newspaper circulation on the incidence of *any lynchings* in a county-year. However, one might be concerned that this reflects, not exposure to publicity, but exposure to places with greater population. Additionally, one might also be concerned that this 'media access' might be capturing exposure to local audiences, rather than distant audiences. To address both of these concerns, I estimated additional models using the same specifications as those given in the paper. First, Tables B2 and B5 show that access to *population* does not consistently affect the incidence of lynching (and the same is true when using urban population). Second, Table B3 shows that the effects of circulation persist (and in fact are stronger) when excluding newspapers within the state. Thus, it is access to audiences *outside* the state that is driving the result. Finally, Tables B1 and B4 show that this effect is weaker, but also present when using weighted access to daily newspapers (weighting all dailies as the same, regardless of circulation). Using circulation as the metric seems more useful, as it captures the size of the audience more directly, whereas the number of daily papers ignores vast differences in the audiences these papers reach.

Table B1: Effects of log Access to Daily Newspapers on Probability of Lynching

	(1)	(2)	(3)	(4)	(5)	(6)
Log Dailies Access	-0.007	-0.007	-0.021^{*}	-0.009^{*}	-0.009	-0.026^{**}
	(0.004)	(0.005)	(0.008)	(0.004)	(0.006)	(0.009)
County FE	X	X	X			
Year FE	X	X	X	X	X	X
Lagged DV				X	X	X
Covariates		X	X		X	X
Local Rail Network			X			X
N	47,649	41,316	41,316	47,649	41,316	41,316
Adjusted R^2	0.068	0.067	0.067	0.071	0.070	0.071

$^{*}p < .05$; $^{**}p < .01$; $^{***}p < .001$
Estimates obtained using OLS, with standard errors clustered by county.

B.3.2 Different Samples

I also consider the sensitivity of these results to both different samples of the data and different specifications. First, it is important to see if the results are sensitive to the sample of lynchings

Table B2: Effects of log Access to Population on Probability of Lynching

	(1)	(2)	(3)	(4)	(5)	(6)
Log Population Access	−0.006	−0.005	−0.025**	−0.006	−0.006	−0.027**
	(0.004)	(0.006)	(0.009)	(0.004)	(0.006)	(0.010)
County FE	X	X	X			
Year FE	X	X	X	X	X	X
Lagged DV				X	X	X
Covariates		X	X		X	X
Local Rail Network			X			X
N	47,649	41,316	41,316	47,649	41,316	41,316
Adjusted R^2	0.068	0.067	0.067	0.071	0.070	0.071

$^{*}p < .05$; $^{**}p < .01$; $^{***}p < .001$
Estimates obtained using OLS, with standard errors clustered by county.

Table B3: Effects of log Access to (out of state) Newspaper Circulation on Probability of Lynching

	(1)	(2)	(3)	(4)	(5)	(6)
Log Circulation Access	−0.028***	−0.041***	−0.067***	−0.033***	−0.048***	−0.079***
	(0.005)	(0.007)	(0.010)	(0.006)	(0.008)	(0.011)
County FE	X	X	X			
Year FE	X	X	X	X	X	X
Lagged DV				X	X	X
Covariates		X	X		X	X
Local Rail Network			X			X
N	47,649	41,316	41,316	47,649	41,316	41,316
Adjusted R^2	0.068	0.068	0.068	0.071	0.071	0.072

$^{*}p < .05$; $^{**}p < .01$; $^{***}p < .001$
Estimates obtained using OLS, with standard errors clustered by county.

Table B4: Effects of log Access to (Out of State) Daily Newspapers on Probability of Lynching

	(1)	(2)	(3)	(4)	(5)	(6)
Log Dailies Access	-0.011^{**}	-0.014^{*}	-0.036^{***}	-0.013^{**}	-0.016^{**}	-0.042^{***}
	(0.004)	(0.006)	(0.008)	(0.004)	(0.006)	(0.009)
County FE	X	X	X			
Year FE	X	X	X	X	X	X
Lagged DV				X	X	X
Covariates		X	X		X	X
Local Rail Network			X			X
N	47,649	41,316	41,316	47,649	41,316	41,316
Adjusted R^2	0.068	0.067	0.067	0.071	0.070	0.071

$^{*}p < .05$; $^{**}p < .01$; $^{***}p < .001$
Estimates obtained using OLS, with standard errors clustered by county.

Table B5: Effects of log Access to (out of state) Population on Probability of Lynching

	(1)	(2)	(3)	(4)	(5)	(6)
Log Population Access	-0.006	-0.005	-0.023^{**}	-0.007	-0.006	-0.026^{**}
	(0.004)	(0.006)	(0.009)	(0.005)	(0.006)	(0.010)
County FE	X	X	X			
Year FE	X	X	X	X	X	X
Lagged DV				X	X	X
Covariates		X	X		X	X
Local Rail Network			X			X
N	47,649	41,316	41,316	47,649	41,316	41,316
Adjusted R^2	0.068	0.067	0.067	0.071	0.070	0.071

$^{*}p < .05$; $^{**}p < .01$; $^{***}p < .001$
Estimates obtained using OLS, with standard errors clustered by county.

I employ. This is important for two reasons: first, as discussed below, some of the data sources for lynchings have been verified by academic historians and sociologists, whereas others were not. It is important to ensure that the results I find are not an artefact of reporting errors by the NAACP and Chicago Tribune. Second, it is important to ensure that the results hold up when only looking at lynchings in the South. While I claim that my argument applies more widely, much of the historical case I make is about Southern lynchings (which are the bulk of events). To this end, I consider five different samples described in Section D.1. In each sample, I restrict my focus to lynching events and counties within states which have different sources and different histories of antebellum slavery. I also consider two different time periods: to maintain comparability with the analyses of coverage of lynching events, I examine both the period from 1880 to 1900 and from 1880 to 1910, even though the railroad travel-times used to compute media access do not change between 1880 and 1910. This makes ten samples in all.

Additionally, I consider 72 different model specifications:

- **Type of Access** I measure media access as a travel-time weighted sum of *circulation* of daily newspapers and the *number* of daily newspapers nationwide. Circulation is my preferred measure, as it captures the size of an audience. But the number of newspapers is also a reasonable measure of audience size. (2 options)

- **Access in/out of state** I also vary whether the measure of media-access includes newspapers within the same state as a county or excludes newspapers within the same state. My argument is that exposure to distant audiences can be costly, so my preferred measure is out-of-state access. (2 options)

- **Access Covariates** Access to circulation or daily papers could simply reflect access to larger populations. If this were true, then "media access" might also be capturing something proximity to economically or politically important places, rather than the possible spread of information to audiences. To address this possibility, I consider the inclusion/exclusion of controls for log access to population and log access to urban population. (2 options)

- **Covariates** I consider three different sets of other covariates: none; logged manufacturing output, logged agricultural output, logged total economic output per capita, logged total population, logged urban population, percent black, and percent black squared; the previous covariates and dummies for the railroad degree centrality (number of rail links) of a county and the degree centrality of its immediate neighbors. (3 options)

- **Model specifications** I consider three different model specifications: year and county fixed effects; lagged dummies for any lynching for each of the five preceding years (LDV) and year fixed effects; all interactions between lagged dummies for any lynching for each of the five preceding years (flexible LDV) and year fixed effects. (3 options)

I estimate all 720 models using a linear probability model, with errors clustered by county and year. First, the effects of access to counts of daily newspapers has inconsistent effects. These vary strongly across model specification. The effects tend to vary depending on sample and model specification (not shown). By contrast, the effects of access to circulation tell a more consistent story. While the effects of access to circulation overall is closer to 0 for the lagged DV specifications in some samples, access to out-of-state circulation consistently predicts fewer lynchings. Reassuringly, the effect of out-of-state circulation is strongest for the academic and academic south samples of lynching: the states for which lynching lists are most correct show consistent, negative, and significant effects of out-of-state circulation access on lynching (Figure B24).

The effects of access to out-of-state circulation do vary by model specification as well (Figure B25). Lagged dependent variable specifications show more inconsistent effects in the absence of covariates, whereas models with year and county fixed effects show consistently significant, negative effects. This may be attributable to lynchings being rare and lagged dependent variables insufficiently capturing county-level attributes that might affect the propensity to experience lynching. Lagged dependent variable models that include county-level covariates consistently show significant, negative effects of access to out-of-state circulation (Figure B26).

Figure B24: Effect of Circulation Access on Lynching: t Statistics by lynching sample

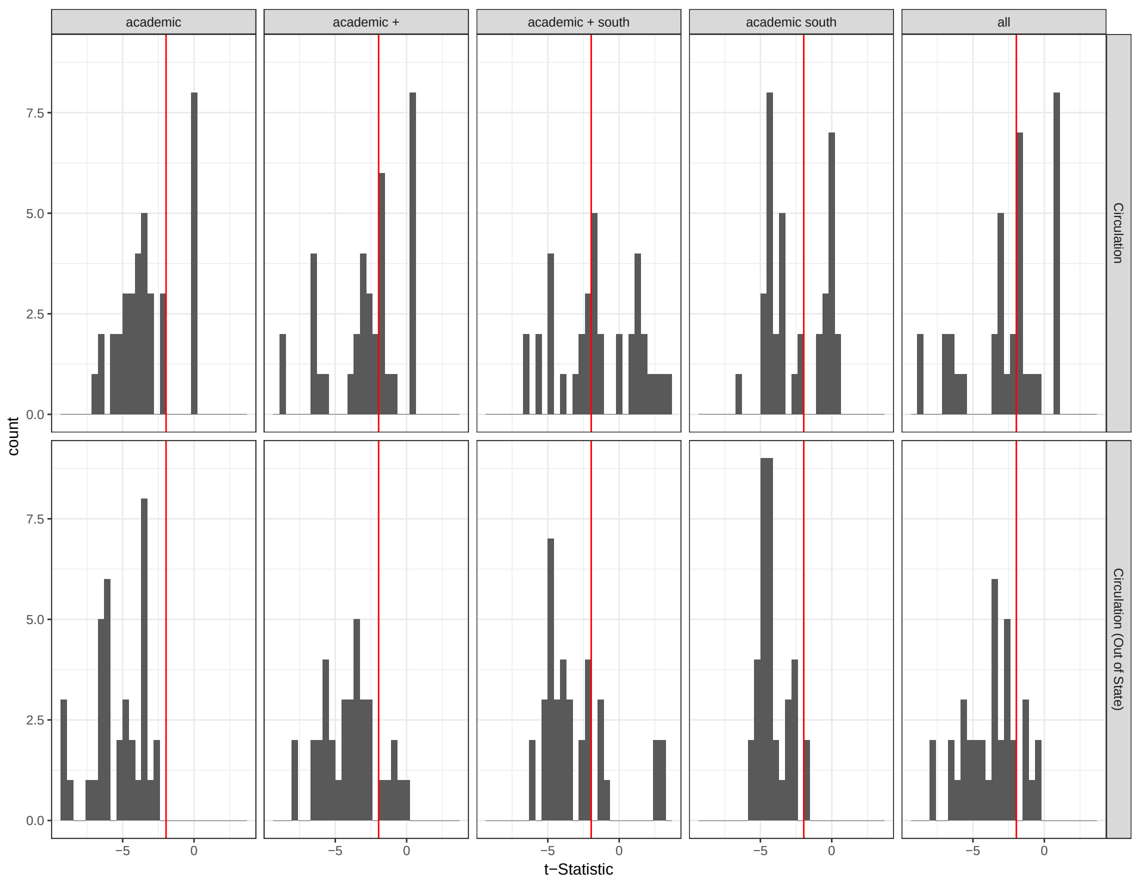

This figure summarizes the t statistics for the effect of out-of-state circulation access on the incidence of any lynching in a county-year for different lynching samples and time periods by different samples of lynchings (and states).

Figure B25: Effect of Circulation Access on Lynching: t Statistics by model specification

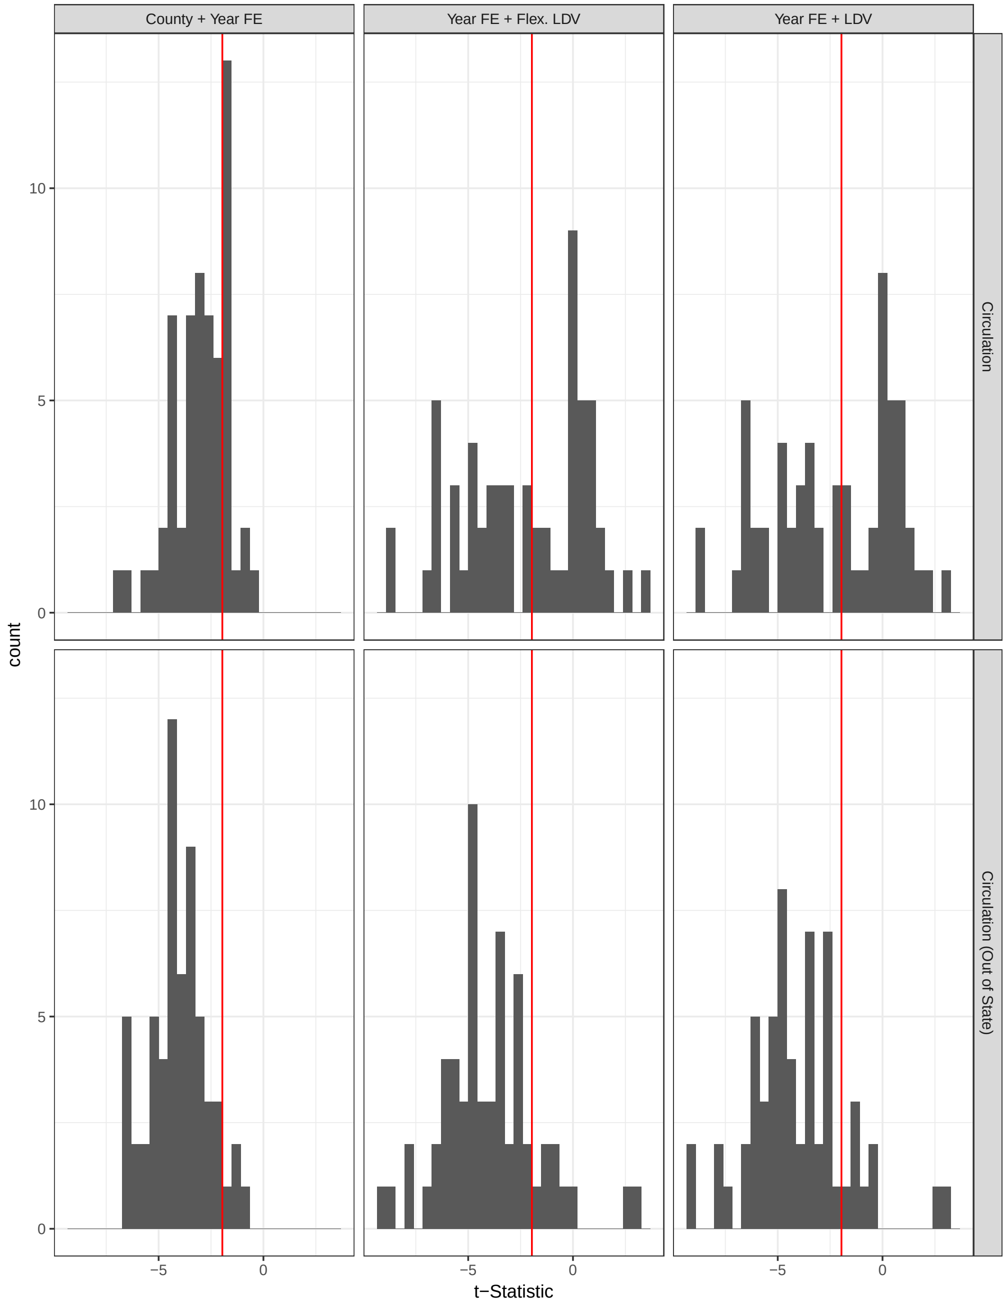

This figure summarizes the t statistics for the effect of out-of-state circulation access on the incidence of any lynching in a county-year for different lynching samples and time periods by model specification.

Figure B26: Effect of Out of State Circulation Access on Lynching: t Statistics by covariates and model specification

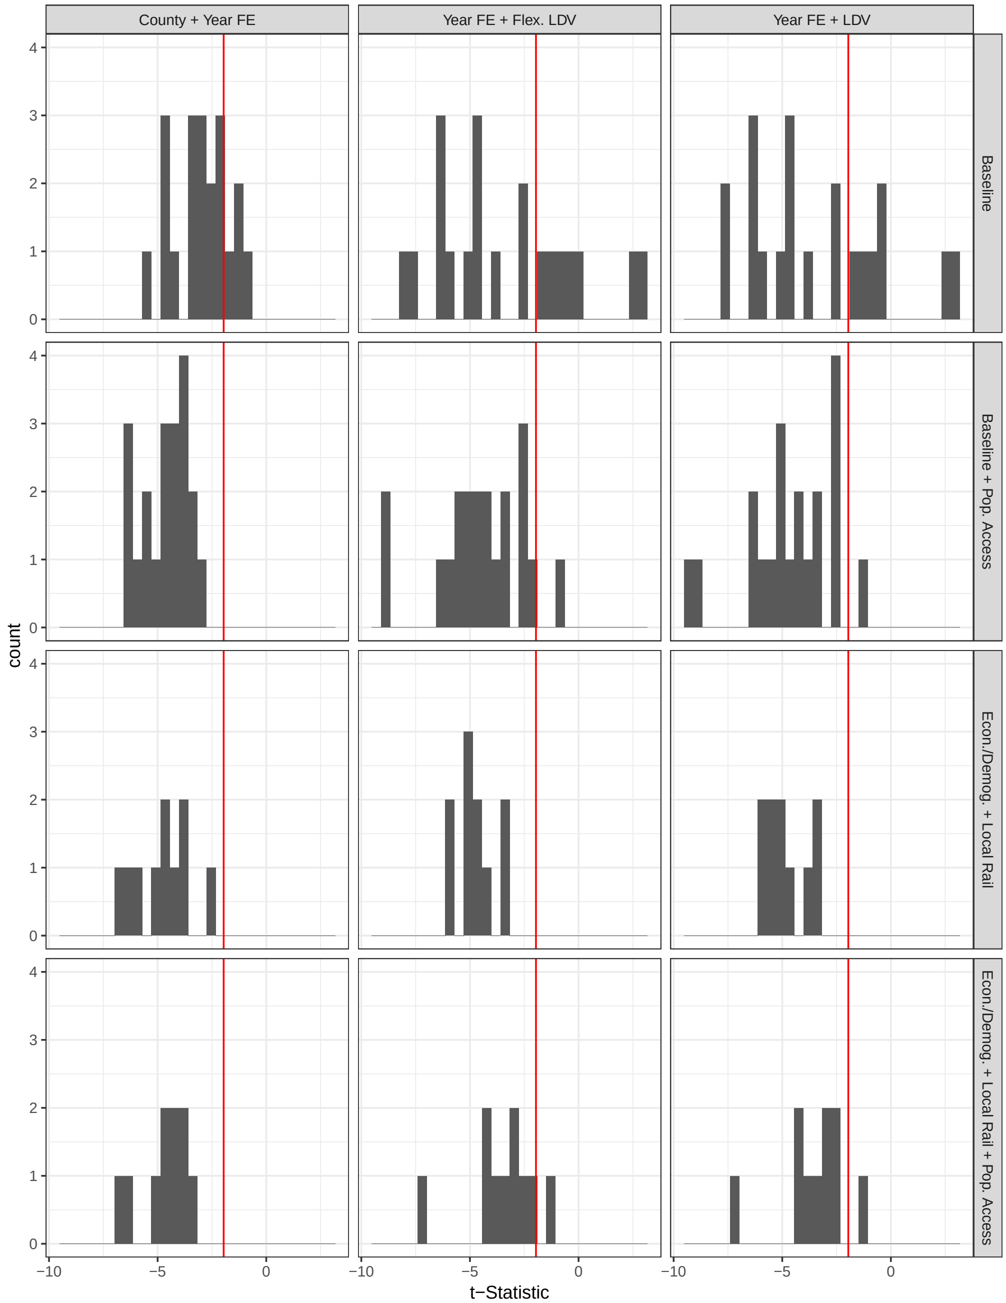

This figure summarizes the t statistics for the effect of out-of-state circulation access on the incidence of any lynching in a county-year for different lynching samples and time periods by covariates and model specification.

B.3.3 Random Effects and Logit

Finally, I consider whether these effects hold up to using logit to estimate the incidence of *any lynchings* in a year. My preferred specification is a linear probability model, because the inclusion of fixed effects does not complicate convergence of maximum likelihood estimates and least squares is more robust to misspecification errors (Angrist and Pischke 2008). But because lynchings are relatively rare, the linear probability model may not perform as well when the probability is near zero. Thus, I also estimate the county and year fixed effect models reported in the main body of the paper using logit with random effects (within estimators). For the "academic south" sample of states and lynchings, there are 1121 county-years out of 20874 with *any* lynchings between 1880 and 1900. This is an incidence rate of about 5 percent which is beyond the thresholds usually required for rare events logit.[6] Figure B27 shows that, across three different sets of covariates, access to circulation is associated with significantly lower probability of observing a lynching ($p < 0.0001$).

[6]The linear probability model does not suffer from bias due to rare events, which is another reason it is my preferred specification.

Figure B27: Effect of Out of State Circulation Access on Lynching: Random Effects Logit

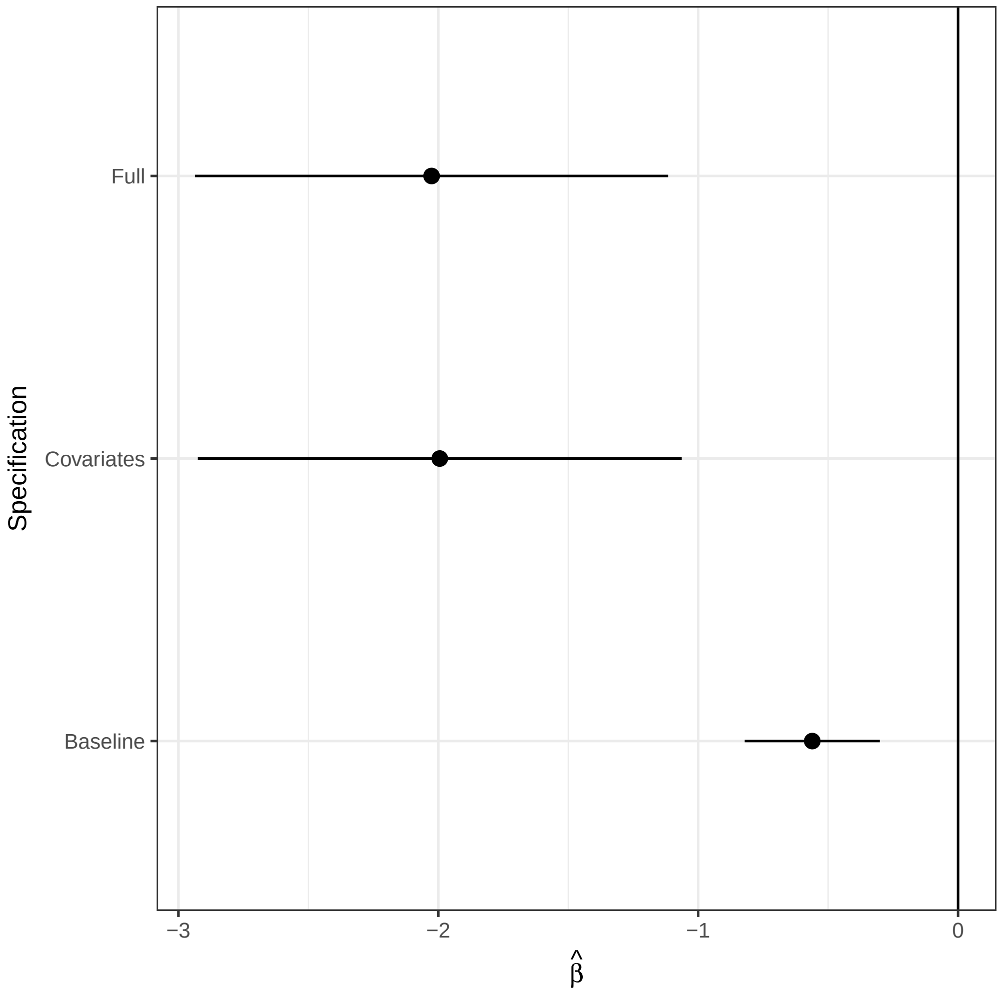

Effects of out-of-state circulation access on incidence of any lynchings for southern academic sample between 1880 and 1900. Estimates come from a logit random effects within estimator. All variables are normalized to help convergence, so scale of coefficients are not comparable with results from fixed effects.

C Background

C.1 Changes in Publicity

From early in the history of the Republic, towns and cities across the United States were stitched together through postal routes and the newspaper exchanges they carried (John 2009). Decades before the telegraph, these networks enabled news to travel across the nation and fostered national political debates (King and Haveman 2008; Kielbowicz 1983). While the United States had certain institutional features of a national public even prior to the nineteenth century, the nature of this public—and the reach and inclusivity of the publicity it entailed—changed drastically in both quality and quantity during the second half of the nineteenth century.

First, the kind of the news that circulated in the early 19th century did little to produce the conditions of reach discussed above. The news that traveled through newspaper exchanges was almost exclusively important news on national or international politics and economics. As a result, exchanges primarily involved moving political news from the national and state capitals and economic news from major centers of commerce to outlying areas (John 2009; Kielbowicz 1983; Russo 1980; Pred 1973). This ensured that news from the centers of economic and political power could spread to hinterland and periphery, but this did not necessarily increase the reach of news from most localities.[7] This was the case for several reasons. Many newspapers initially were funded not through circulation but through patronage of political parties, which both encouraged trafficking in national political news and discouraged collecting original news for local readers (Kielbowicz 1989). Moreover, local news simply not understood to be "news" since it would have been known to most people without the paper (Russo 1980; Huntzicker 1999). In other words, despite the postal system in this early period, local events were not reaching wider audiences, a necessary condition for the effects of publicity discussed above.

Second, the timeliness of news changed as well. In the first half of the nineteenth century, the speed of news was both much slower and highly uneven compared to the end of the century. News from urban centers of politics and commerce moved faster than from outlying localities, because outlying papers published less frequently and postal routes *returning* to major urban areas were slower and more infrequent, making this news slower and outdated (Pred 1973). This speed differential may have further hampered the dissemination of local news to wider audiences. The slow speed of news also limited other conditions required for the effects of expanding publicity. News might take weeks or months to reach a wider audience, leaving audiences little to do in response, as events had long since passed, and any criticism would take even longer to reach locals. Moreover, the slow speed of news also meant that news reached audiences in a staggered sequence, which prevented audiences nationwide from simultaneously responding to events. This, then, inhibited the growth of a story and made it less likely that scandal surrounding violence could erupt. This state of affairs was particularly true in the South which, as compared to the North and West, had fewer newspapers per capita, papers that published less frequently, and poorer and fewer postal routes (Pred 1973; Kielbowicz

[7]While newspaper exchanges brought outlying papers back to major urban areas, their main purpose was to disseminate news from centers of power. Postal officials and urban editors bemoaned the deluge of rural exchanges that slowed the post and usually ended up as waste basket lining (Kielbowicz 1982).

1985).

Finally, the sheer quantity of news was far less. Newspapers printed fewer pages of material and fewer issues, due to both lack of content and the limits of printing technology (Clark 1948; Russo 1980). And the volume of news that traveled was far less: the post office carried exchanges between newspaper editors for free, but sending private letters or newspapers and periodicals to subscribers was expensive(Kielbowicz 1986).[8] This was broadly because the use of wagons and coaches over often difficult roads placed hard limits the amount of mail that could be carried (States 1885). Ultimately, these conditions limited the reach of publicity, and concomitantly, limited the possibility for greater inclusivity. Several changes over the course of the century drastically increased both reach and inclusivity of publicity.

First, local events entered the press for several reasons. This happened first in urban areas, where larger populations simultaneously made it difficult to learn of local events by word of mouth alone and made it feasible to generate revenue based on circulation and advertising instead of political patronage. This made it profitable to cover local events for a local audience (Huntzicker 1999). When the expansion of the railroad and changes in Post Office policy both expedited and reduced the cost of mail, urban papers began directly competing with outlying newspapers by delivering directly to subscribers. Unable to compete with urban dailies in coverage of national and international news, rural and suburban papers began to cover local news to keep circulation up (Russo 1980; Kielbowicz 1986).

Second, changes in technology and postal policies created new audiences, increased both the volume of news and the speed at which it moved, and helped create a truly national print culture. While the first commercial railroads and telegraph were introduced in the United States in the decade between 1828 and 1838, in 1870 these technologies had only *just* connected the Atlantic and Pacific coasts, and much of the country lacked direct access to these networks. The cost of using the telegraph for news was beyond the reach of most individual newspapers, and even membership in the Associated Press was restricted primarily to urban dailies (Kielbowicz 1987; Blondheim 1994). But in the 1870s, this began to change: multiplexing made it possible for telegraph wires to send multiple messages simultaneously, bringing down costs (Kielbowicz 2016), and the Railway Post Offices system that streamlined and revolutionized the United States Postal Service was in its infancy (Carpenter 2000). Between 1870 and 1900, the miles of railroad in the United States nearly tripled from 38,000 to 131,000 miles (Perez-Cervantes 2014). The vast majority of this increase occurred during the 1880s and 1890s. By 1911, the rail network in the US had reached its peak (Atack 2013). And between 1880 and 1910 the mileage of Western Union's telegraph wires alone nearly sextupled, and the Associated Press and other wire services created nation-wide news markets (Blondheim 1994).

Railroad and telegraph helped to spread news further and faster. Railroads permitted news to travel by post at a much faster rate and in much greater volume than before, while telegraph permitted instantaneous transmission of news across great distances. Even though other postal routes continued, railroads fundamentally reshaped these networks, turning them into spurs off of railroad lines (States 1885; Carpenter 2000). In contrast, express services predating these changes did reduce the travel time of news, but the expense of these endeavors was considerable and as a result they could only be used by major urban dailies to collect vital

[8] King and Haveman (2008) discussion of the post office and the spread of abolitionism is an exception here: religious periodicals uniquely had favorable postal rates (Kielbowicz 1986)

market and political news, reinforcing the primacy of news from the center over news from the periphery (Kielbowicz 1985; Blondheim 1994). Telegraph and railroad also enabled greater news collection: railroads helped bring news from outlying areas to regional cities, and from there it could be shared on the telegraph network (Kielbowicz 1987).

The reduced cost and increased speed of news not only ensured that it reached wider audiences while it was still current, the it also created news cycles and amplified coverage and increased demand for news from distant places (Blondheim 1994; Kielbowicz 2016). With the news moving fast, news cycles emerged in which not only would an event be reported, but responses to the event and reception of the coverage became part of the story itself. This served to amplify the amount and duration of coverage given to events, and also made it possible for, in this case, criticism of a lynching to become part of the story. At the same time, the speed of news also helped make it possible for people to imagine themselves as part of a unified audience that stretched across the country (Kielbowicz 2016). During this time, newspaper and telegraph offices often attracted large crowds to hear the latest news, whether it was sports or politics. This permitted people across the country to "experience" distant events simultaneously, which could lead to outpouring of sympathy and aid in the wake of natural disasters, or waves of racial violence as when African American boxer Jack Johnson defeated the "Great White Hope" in a highly-anticipated and nationally-telegraphed fight.[9]

The combination of growing railroad networks and the introduction of second class mail brought a national print culture across the country. With low second class mail rates and the capacity and speed of railroads, people across the nation were able to subscribed to newspapers and periodicals printed in distant urban centers. Between 1880 and 1920 the volume of second class mail increased at a rate twenty times the size of the population (Kielbowicz 1990). This brought localities across the country into contact with a veritable avalanche of mass culture and advertising aimed a truly national audience. This both ensured locals knew of any criticism of their community, but also may have helped imagine distant audiences as peers or equals within a national community.[10]

C.2 Railroads and Publicity

Why use railroads as an indicator for access to the national public? In this section, I discuss the motivation for using railroad network centrality and railroad travel time.

- **Railroad Centrality** Rail network centrality indicates how well connected a county is to others via the railroad. This can capture differential access to publicity in two ways. First, betweenness centrality indicates the frequency with which a county appears on the shortest path between other counties. As a county is on the shortest path between more pairs of counties, it is likely that it also saw more rail traffic. This is because train travel was still costly. So it was likely that freight, passenger, and mail traffic between any pair of counties would move along the shortest possible path. This would mean

[9] And this type of coverage also occurred for lynchings, such as the case when a reporter for a Denver newspaper worked with telegraphers to transmit live reports on a lynching as it happened (Kielbowicz 2016).

[10] That this process was both real and powerful is perhaps best evidenced by the attempts of local publishers, businesses, and religious leaders to end subsidies for urban publications as a means to bolster local interests and defend against the changes urban mass culture was making to life in the hinterland (Kielbowicz 1990; Fuller 2003).

that more trains might pass through a county, picking up news as either passengers or railway employees interacted with locals. Given that Kielbowicz (1987) shows that news from rural areas continued to be collected via mail, railway postal routes may have been particularly important. While annual data on the location of railway postal routes exist in reports to Congress, digitizing them for the purpose of this study would be prohibitively expensive in both time and money.

Betweenness centrality also serves a proxy for telegraph networks. Because telegraph networks were frequently coterminous with the railroad (even if offices were spaced out), being located on more shortest paths might mean that a county is also on direct telegraph trunk lines between many other places. This might make it more likely that news would spread. Moreover, while anywhere located along a telegraph line theoretically could "access" it, because it ran along rail lines, in practice, telegraph companies only operated offices at locations spaced out along the lines in order to avoid cannibalizing their sales (Garcia-Jimeno et al. 2018). In counties that were on more shortest paths, it seems more likely that rail (and telegraph lines) would intersect and telegraph offices would be present and possibly more geographically dispersed through the county.

Eigenvector centrality might capture a different important aspect of access to publicity. Places with higher eigenvector centrality are connected to places that themselves have more connections. In this case, that would translate to: counties with more rail links to counties that themselves have more rail links have higher rail network eigenvector centrality. This could create greater publicity as any news that travels by rail out of a county with high eigenvector centrality would then potentially spread more quickly, because its neighbors are themselves highly connected. By contrast, news that comes out of counties with low eigenvector centrality would spread more slowly, as its neighbors are also poorly connected to the network.

- **Railroad Travel Time** Railroad travel time between two places (conditional on travel time by other modes) maps onto increases in publicity in two ways. First, it directly captures access to news. If information travels by rail (through passengers and railway employees but also the post), then decreasing travel time means that news travels faster possibly and more frequently. This should, reasonably, lead to greater transmission of information between two places. It might also, indirectly, capture increasing interest among audiences. The easier it is to travel to a place, the more likely it is that people might develop an interest in it for personal or economic reasons. This increase in interest could also lead to greater coverage.

I exclude closeness centrality—the inverse of the average distance from a county to all other counties—since I examine this more explicitly using travel time data. I also exclude degree centrality—the number of rail links directly connected to a county—as local rail construction is likely endogenous to local economic and political conditions.

C.3 Partisanship of Press

One possible reason why pro-lynching discourse in the South was challenged less frequently could be the distribution of the partisan press. Compared to the North and West as well as

Border States, Southern states saw nearly complete dominance of the Democratic press. While in the 1870s, the Republican press had a foothold in the region, taking on white supremacy with the backing of state patronage (Abbott 2004), following Redemption only a tiny fraction of both daily and weekly papers remained (Figures C1 and C2). By contrast, other regions saw closer parity between the number of Democratic and Republican papers during the same period.

Given that Southern Democrats embraced white supremacy and that Republicans in the South had backed civil rights for African Americans, the dearth of Republican newspapers in the South may have been one way by which opposition to lynching was silenced (Abbott 2004). And by the same token, the presence of more Republican papers in the North and West meant that, as news spread further, lynchings gained the attention of more critical editors.

Figure C1: Partisanship of Daily Papers by Region over time

Derived from data by (Gentzkow and Shapiro 2014). Includes daily papers nationwide.

Figure C2: Partisanship of Southern Daily and Weekly Papers over time

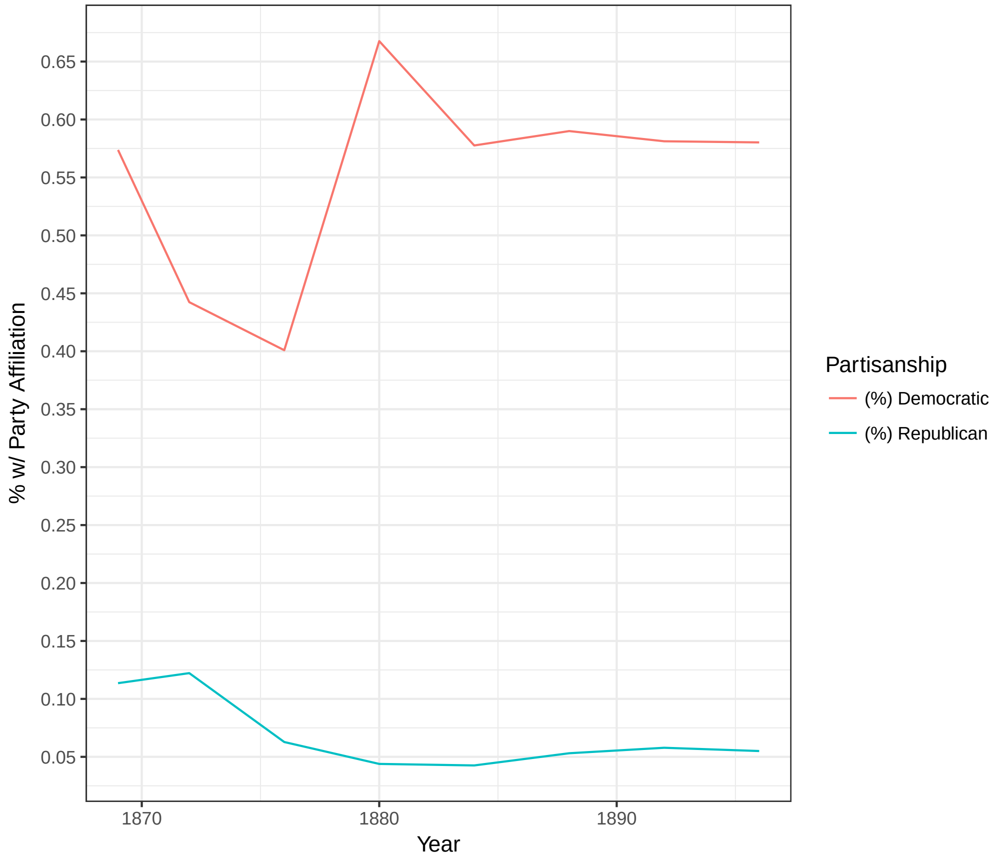

Derived from data by (Gentzkow and Shapiro 2014). Includes both daily and weekly papers from the US South.

D Data

D.1 Lynchings

I compile data on lynchings nation-wide. This is for two reasons: one methodological and one theoretical. First, from a methodological perspective, nationwide data permits me to model coverage as a function of events in and out of the South, which could occur simultaneously. Second, theoretically, lynchings occurred in many parts of the country. While the vast majority were in the former Confederacy and border states (i.e. former slave states), the mechanisms by which I propose lynching gained more publicity should be operative outside the South. Thus, it makes sense to include these states as well. The procedure I used to compile lynching events nationwide is detailed below.

1. I started with a lists of lynchings from two groups of sources. The first group of sources are academic, and come from the work of academics and historians who have verified the occurrence of and facts about lynching events in specific states or groups of states. Data on most of the former Confederacy come form Tolnay and Beck (1995), data on Georgia and Virginia come from Brundage (1993), data on Kentucky come from Wright (1990), and data on numerous states in the North and West come from Pfeifer (2011). To this, I also add lists of lynchings in Texas and the West collected by Carrigan (2005), though these are not an exhaustive record for the relevant states. The second group of sources are non-academic and contemporaneous with lynching. Both the NAACP and Chicago Tribune collected their own lists of lynching events, which cover the entire country. However, these lists are flawed and include events that were not lynchings (Tolnay and Beck 1995). These sources and the states they cover are listed in Table D1.

2. For each of these several lynching lists, I standardized the key identifying information: location (state, county, and town), date, and victim names and races. This provided me a single list of all reported lynchings and their source.

3. I then matched identical lynchings across sources, to both pool information (not all agreed on location or names) and to eliminate duplicate events. This procedure worked as follows: using a computer script, I looped through each lynching for a given source and found all potential matches from other sources. The criteria for potential matches were: matches on state, year, and month; exact date match, or near date match; exact county name match; exact victim name match; fuzzy county name match quality; fuzzy victim name match quality. Using these match criteria, I then created potential matches. First, I counted as potential matches those where the state and date matched exactly and uniquely across sources. Then, I added non-unique state and date matches. For events without any match, I classified as possible matches as those with the best name and best date similarity, best name similarity, and best date similarity.

   These potential matches were then evaluated manually. For each potential match, I and research assistants evaluated whether the names, locations, and dates were the same, allowing for tolerance on misspellings, name variants, and minor date errors (usually a few days or transposition of numbers). Then, using these manually classified matches, I created clusters of event records that were connected on a graph. (That is, if lynching 1 was matched to 2, and 2 was matched to 3, but 3 and 1 were not matched, (1,2,3) were counted as an event cluster, because they created a connected graph.)

I then manually inspected these clustered events, starting with the largest clusters. For each cluster, I manually investigated whether they were the same event or not, and split them accordingly. The criteria here was again using names, dates, locations, and the existence of multiple events from the same source (deferring to academic sources). After manually classifying these event clusters, I generated a unique ID for each cluster of events. These represent the lynching events used in my analysis, and correspond to reports in (potentially) multiple academic and non-academic sources.

4. For each event cluster, I then geocoded its recorded location according to each source using the Google Maps API. When this geocoding failed or returned a "route" (a road of some kind), I manually geocoded this event. This entailed using location, name, and date criteria from the source to identify the most local newspaper coverage of this event. Using this information, I then consulted historical gazetteers and maps to provide coordinate (where possible). I then used these coordinates (from different sources), to match events to 2000 County Boundaries (same as the railroad and census data). If the county matches disagreed across sources, then I again used local newspaper accounts and academic sources to identify the most specific location for the event and then obtained the latitude, longitude, and 2000 county.[11]

   I took as the date for each event the date given by an academic source (where available) and the earliest date when sources disagreed. This is because the NAACP and Chicago Tribune lists often report the date of the newspaper coverage of the lynching, rather than the date of the event. This is evident when comparing NAACP and Tribune event dates to those from academic sources for the same lynchings.

5. Finally, if in the process of manually geocoding an event, I had found newspaper evidence that no lynching occurred, then I dropped this from my final analysis.

Rather than provide a complete set of programming files to reproduce this process exactly, in the replication a file I provide a list of the lynching events I identify matched to each of the corresponding original source records. This file also includes the date used and the 2000 County FIPS code. This would permit others evaluate the sensitivity of my results to changes in these classifications and matching decisions that I made.

When using this data in analysis, I am sensitive to two sets of concerns. First, the academic sources may be more valid than the non-academic sources, because they exclude some events that did not occur. However, because my argument about publicity of lynchings should apply also to attempted, threatened, or rumored lynchings, it may be appropriate to include falsely reported lynchings in the analysis of coverage. On the other hand, there is no available systematic list of threatened or attempted but incompleted lynchings.[12] As a result, the data on non-lynchings is unrepresentative. Thus, I consider some samples in which only lynchings

[11] During this process, I was unable to find any digitized newspaper account of a few lynchings. This does not mean the event lacked coverage at all. The relevant coverage may not have been digitized, or the quality of digitization made keyword searches ineffective.

[12] Beck et al. (2016) have compiled this data for much of the South, but, to my knowledge, it has not been released for public use.

reported in academic sources are included and others that include all reported lynchings. Second, my argument is primarily about Southern lynching (but see Campney (2013) for a similar logic in Kansas) and the South differed from the North substantially in its access to rail and telegraph in the 1880s. Thus, it is important to ensure that the effects I find are not driven by events outside the South. Using the event data I compiled, I created five different sets of lynching:

1. Full Sample: This sample includes all reported lynchings, including those later determined by academic researchers to be either false, unconfirmed, or perhaps not lynchings. This is the default sample for the coverage analyses.

2. Academic Sample: This sample includes lynchings only from those states for which academics have compiled *complete* lists of confirmed lynchings. This excludes some states, like Texas, where some lynchings have been verified, but others have not.

3. Academic Sample (South): This sample includes lynchings only from states with legalized slavery at the start of the Civil War where academics have compiled *complete* lists of confirmed lynchings. This includes, then, states in the former Confederacy and Border States.

4. Academic Sample +: This sample combines both academic and non-academic sources. For states where academics have compiled complete lists of confirmed lynchings, only those academic sources are used. For states without or with incomplete academic sources, I supplement using events reported by the NAACP or Chicago Tribune.

5. Academic Sample + (South): This sample combines both academic and non-academic sources for states that had legal slavery at the start of the Civil War. For states where academics have compiled *complete* lists of confirmed lynchings, only those academic sources are used. For states without or with incomplete academic sources, I supplement using events reported by the NAACP or Chicago Tribune.

The inclusion of each state these samples is given in Table D2.

D.2 Railroads

All measures of railroad centrality and travel times are created using data from (Perez-Cervantes 2014) and (Atack 2013). The Perez-Cervantes data used a database of railroad construction projects to identify segments of track constructed in each year between 1840 and 1900. In total, it covers more than 90 percent of the known mileage of railroads. Some of this gap is attributable to segments that were double track.

From these lines, edges of a network of counties were created by identifying which lines crossed shared county borders (using 2000 county boundaries). These edges were given the length of distance between each county population centroids. In addition to the railroad network, I construct similar edges between counties whose borders share navigable waterways or canals using the data from Atack (2013). Finally, I also created edges between counties that were directly adjacent and treated these as travel on overland roads. These simplified networks were used to calculate centrality measures and travel times.

Table D1: Lynching Sources and States Covered

Source	Type	Coverage	States
(Tolnay and Beck 1995)	Academic	Complete	AL, AR, FL, GA, KY, LA, MS, NC, SC, TN
(Brundage 1993)	Academic	Complete	GA, VA
(Wright 1990)	Academic	Complete	KY
(Pfeifer 2011)	Academic	Complete	AK, AZ, DE, IA, ID, IL, IN, ME, MI, MN, MT, ND, NE, NJ, NV, NY, OH, OR, PA, SD, UT, WA, WI, WY
(Carrigan 2005)	Academic	Partial	TX
(Carrigan and Webb 2013)	Academic	Partial	AZ, CA, CO, LA, NE, NM, NV, OK, TX, WY
Chicago Tribune	Non-Academic	Complete	AK, AL, AR, AZ, CA, CO, CT, DE, FL, GA, IA, ID, IL, IN, KS, KY, LA, MD, ME, MI, MN, MO, MS, MT, NC, ND, NE, NJ, NM, NV, NY, OH, OK, OR, PA, SC, SD, TN, TX, UT, VA, WA, WI, WV, WY
NAACP	Non-Academic	Complete	AK, AL, AR, AZ, CA, CO, DE, FL, GA, IA, ID, IL, IN, KS, KY, LA, MD, ME, MI, MN, MO, MS, MT, NC, ND, NE, NJ, NM, NV, NY, OH, OK, OR, PA, SC, SD, TN, TX, UT, VA, WA, WI, WV, WY

Table D2: Lynching Samples: States Covered and Source Type

Sample	Academic Source States	Non-Academic Source States
Full	AL, AR, AZ, DE, FL, GA, IA, ID, IL, IN, KY, LA, ME, MI, MN, MS, MT, NC, ND, NE, NJ, NV, NY, OH, OR, PA, SC, SD, TN, UT, VA, VT, WA, WI, WY	AL, AR, AZ, CA, CO, CT, DC, DE, FL, GA, IA, ID, IL, IN, KS, KY, LA, MA, MD, ME, MI, MN, MO, MS, MT, NC, ND, NE, NH, NJ, NM, NV, NY, OH, OK, OR, PA, RI, SC, SD, TN, TX, UT, VA, VT, WA, WI, WV, WY
Academic	AL, AR, AZ, DE, FL, GA, IA, ID, IL, IN, KY, LA, ME, MI, MN, MS, MT, NC, ND, NE, NJ, NV, NY, OH, OR, PA, SC, SD, TN, UT, VA, VT, WA, WI, WY	
Academic +	AL, AR, AZ, DE, FL, GA, IA, ID, IL, IN, KY, LA, ME, MI, MN, MS, MT, NC, ND, NE, NJ, NV, NY, OH, OR, PA, SC, SD, TN, UT, VA, VT, WA, WI, WY	CA, CO, CT, KS, MA, MD, MO, NH, NM, OK, RI, TX, WV
Academic (South)	AL, AR, FL, GA, KY, LA, MS, NC, SC, TN, VA	
Academic (South) +	AL, AR, FL, GA, KY, LA, MS, NC, SC, TN, TX, VA	MD, MO, TX

Travel Times I obtained travel times between counties with and without railroads by converting the length of edges between counties into travel times associated with the specific mode of transportation. Road speeds were taken to be 30 miles per day, canal speeds were taken to be 4 miles per hour, river speeds were taken to be 12.5 miles per hour, and railroad travel times were taken to be 25 mph (1880–1885), 30 mph (1886–1895), and 35 mph (1896–1910). These times are all average speeds identified by economic historians of the period (Perez-Cervantes 2014).

To obtain the shortest travel time between places, I use Dijkstra's algorithm to find the shortest path, where each edge was weighted by its length in time. I estimated the travel time using only overland routes ("distance time"), using overland and water routes ("no-rail time"), and using all routes including railroads ("rail time").

Centrality I calculate centrality measures as centrality within the *rail* network. This excludes all non-rail edges. For both betweenness and eigenvector centrality, counties with no rail links take a value of zero. I use the travel times for each edge in order to compute betweenness centrality, since this is pertinent in correctly identifying the number of times a county appears on a shortest path between other counties. There is no clear way of implementing weights in eigenvector centrality. When converting these centrality scores to deciles, I compute deciles for all counties within each year. This is especially important for eigenvector centrality, since it's scores may not be directly comparable for different networks.

D.3 Newspaper Coverage

Data on newspaper coverage of lynchings come from four digital historical newspaper archives:

- *America's Historical Newspapers:* This archive is maintained by Readex and permits a variety of sophisticated keyword searches over pages, as well as articles within pages. This is the only source that separates out articles on the page. For the purposes of the analyses in this paper, only simple keyword searches of pages are used. This archive is largely static and has not seen much expansion. The quality of the underlying text used for keyword searches appears to be high. In total, this archive provided 835 newspaper titles.

- *Chronicling America* This archive is publicly hosted by the Library of Congress and is composed of newspaper digitized by state libraries under a federal grant. The extent to which states pursued this digitization is highly variable, and this archive originally included newspapers only through 1923 (though, it has since been updated). The quality of the underlying text used for keyword searches appears to be high. In total, this archive provided 1798 newspaper titles.

- *NewspaperArchive* This archive can be accessed via personal or institutional subscription. It is constantly updated with new material and includes many small-town newspapers. However, the quality of the transcription is often poor. This archive provided 2351 newspaper titles.

- *Newspapers.com* This archive can be accessed via personal subscription. It is constantly updated with new material and includes many small-town newspapers. However, the quality of the transcription is highly variable. This archive provided 3784 newspaper titles.

Data used in this paper was collected from these archives in the summer of 2015. The following data was accessed from these archives:

1. Each archive was indexed in its entirety for the years 1880 through 1940. This indexing obtained a list of each newspaper held by the archive (including its place of publication), as well as a list of each issue (its paper and date) held by the archive. This indexing makes it possible to have a "denominator" when looking at coverage, because it includes in the data newspaper issues that could have matched a keyword search but did not.

2. Keyword searches for words and phrases corresponding to lynching were executed in each of these archives. The keywords used to find lynching-related coverage are: 'lynching', 'lynched', 'lynchings', 'lyncher', 'lynchers', 'lynches', 'lynch law', 'judge lynch', 'lynch mob', 'lynching mob', 'mob law', 'mob rule', 'mob violence'. These keywords were chosen to minimize false positives, though this may result in more false negatives (discussion of lynching using other words).

3. Keyword searches for words and phrases corresponding to pro- and anti-lynching discourses were executed in each archive. This is covered in more detail below.

Using newspaper places of publication, each paper was matched to a state and county.

D.3.1 Representativeness of Digitized Papers

In total, I collected data on 8768 newspapers. While this includes papers from all over the country, from cities and country towns, weeklies and dailies, it still is not representative. In what follows, I briefly describe what is included and excluded in this coverage data. Most obviously, while there is a large number of titles, not all titles have a complete run of issues digitized. Some titles have only one issue, while others have thousands. To get a sense of how what is digitized might be biased across space and time, I matched the digitized newspapers to the Library of Congress's list of American newspapers.

This list gives all newspaper titles known to the Library of Congress through its and other library holdings. This list includes the title, place of publication, and known publication dates for all newspapers. While this is likely incomplete, it most likely includes most papers. Using this database, I created a list of newspapers and newspaper issues that existed by state-year, and year. This was done using the following procedure:

- For each paper in the LOC database, I identified its publishing schedule as daily, weekly, monthly, or quarterly. Then, using its reported start and end date, I created a list of each year in which the publication should have existed. Because there is uncertainty about the start and end date of some publications, I expand the start dates backward in time to the earliest possible date and expand end dates forward to the latest possible date.[13]

[13] For example, if the start year of a newspaper was given as '187?', I would code it as '1870'. If it had an end year as '19??', I would code it as '1999'.

Then, for each year, the paper was given the expected number of issues depending on the publishing cycle (365 for dailies, 52 for weeklies, 12 for monthlies, etc). Thi resulted in a list of newspaper-years and the expected number of issues in that year.

- I then used location, title, and year to match digitized newspapers in my sample to the LOC list. For each digitized newspaper in my sample, I calculated the number of digitized issues for each year it was available. I then merged the digitized records to the complete listing of newspapers.
- Using this merged data, I was able to compute the percent of newspaper-years for which at least 1 issue of the newspaper was digitized, and the average fraction of issues digitized in each newspaper-year.

Because of the choices I make in coding the years in which a paper was published (assuming the widest possible interval) and its frequency of publication (assuming a complete publishing run each year), my estimates of how much material has been digitized is likely biased downward.[14]

Figures D1 and D2 are maps that show, by state, the average fraction of newspaper-years and newspaper-issues that are digitized for the years 1880 to 1940. It is very evident that the coverage is uneven. Some states, like Indiana, North Carolina, and Utah, have much higher digitization rates than others. One notable pattern is that Southern states have lower digitization rates, generally, than the rest of the country. Figure D3 shows, by year, the fraction of newspapers with any digitization and fraction of issues digitized. Over most of the period, only around 4 to 5 percent of all newspapers and newspaper issues are available in the archives I use. The stark dropoff in the 1920s is due to the Chronicling America archive stopping in 1923. The fraction of *issues* digitized is likely higher than the fraction of distinct newspapers digitized, because daily papers are digitized more than less frequent publications.

D.3.2 Matching Coverage to Events

Both analyses of the effects of railroad access on coverage rates and the relationship between distance and criticism of lynching require matching lynching events to coverage in newspaper issues. Between 1880 and 1910, I identify a maximum of 3233 possible lynching events. While Seguin (2016) was able to match coverage to lynching events based on searches of three major newspapers, it is highly impractical to attempt to do this for the thousands of papers in my sample.[15] Rather than manually investigate whether each newspaper issue covers a specific lynching event, I use a simple set of rules for coding "coverage." For each lynching event, I consider all newspaper issues published within a range of days (from 0 to w, where w is a window between 3 and 11) as *capable* of covering that event. If a newspaper issue within 0 to w days since the lynching uses a lynching keyword (see above), then I classify that as "coverage" of the lynching event. In the main specifications reported in the body of the paper, I set w as 7 days.

This is an obvious simplification and could suffer from a few issues. First, the use of lynching keywords could undercount coverage that doesn't use those words to described lynchings

[14]Though, because the LOC archive is incomplete, there is also an upward bias.

[15]For more on this, see D.3.3

Figure D1: Map of Estimated Digitization Rate (newspaper-years) by State

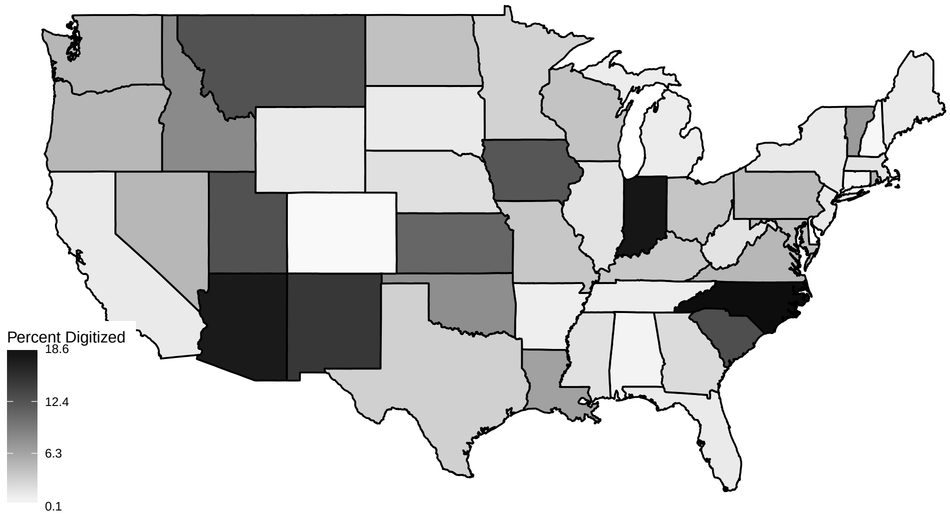

This figure shows, by state, the average fraction of newspapers published in a given year with *any* digitized issues (in the newspaper archives used in this paper) between 1880 and 1940.

Figure D2: Map of Estimated Digitization Rate (of newspaper issues) by State

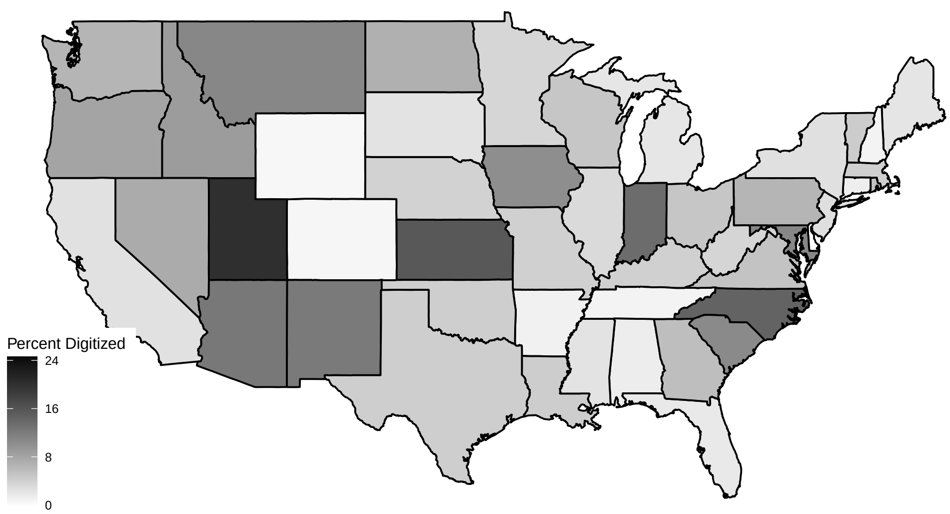

This figure shows the fraction of newspaper issues digitized (and available in my data) by state between 1880 and 1940. This is estimated based on newspapers, newspaper publication years, and newspaper publication frequency given by the Library of Congress.

(not common, but this happened). This could be a problem, particularly for analyses of dis-

Figure D3: Digitization Rates by Year

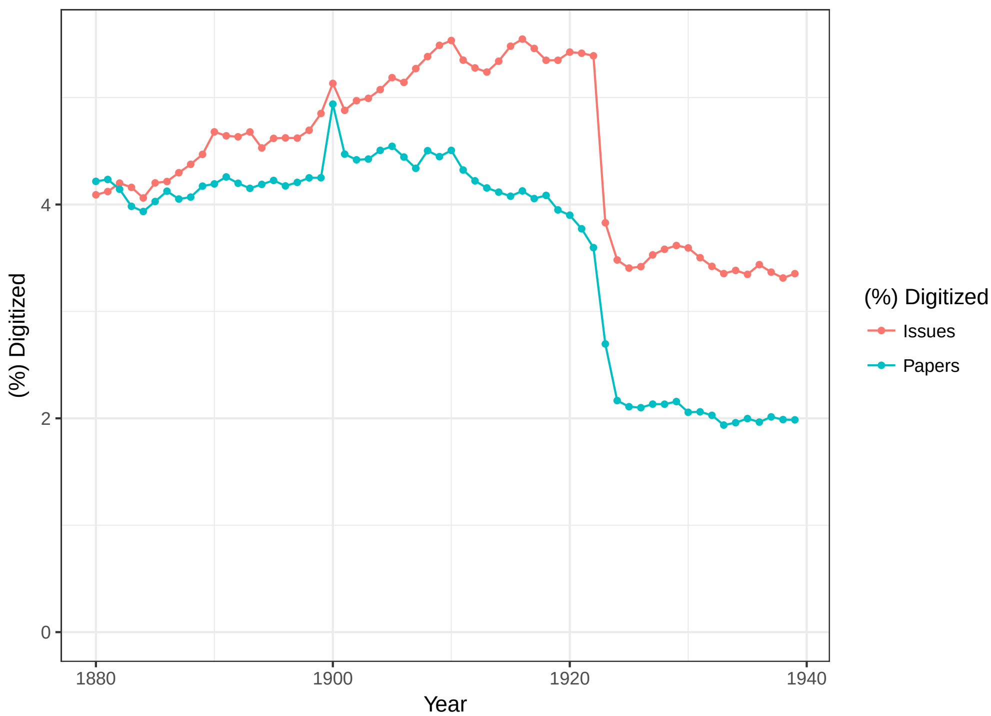

This figure shows the fraction of newspaper issues digitized and the average fraction of newspapers published in a given year with *any* digitized issues by year between 1880 and 1940. This is estimated based on newspapers, newspaper publication years, and newspaper publication frequency given by the Library of Congress.

couse, because newspapers that refused to classify a lynching as a "lynching" often endorsed a set of criteria for "warranted" lynchings and contrasted them from "unwarranted" lynchings (Jean 2005). Thus, coverage of lynchings that do not use lynching-related words may actually be more pro-lynching.

Second, and more obviously, using keyword searches within a short window of time could induce both false positives and false negatives. False positives (classifying coverage of an event when it does not occur) would occur because any discussion of lynching would be classified as coverage. This could count of coverage of one lynching event: coverage of other lynchings, reports of threatened or attempted lynchings, or discussions and editorials on lynching as a general phenomenon. Conversely, false negatives (classifying as not coverage when coverage does occur) could occur because coverage does not use a lynching keyword, the lynching keyword was incorrectly transcribed, or the coverage occurred *outside* the coverage window w.

D.3.3 Validating Coverage Matching

I validate the use of this noisy measure of coverage using two different sources of information.

- **Classification of Event-Specific Coverage** At an early stage in this project, I attempted to manually collect data on newspaper coverage of specific lynching events.

Because this was both time-consuming and very expensive, I dropped this approach in favor of the simpler procedure outlined above. Before stopping, I compiled data on coverage of 34 lynching events chosen at random from the Tolnay and Beck list of lynchings. Data on coverage for a specific lynching followed the following procedure. For each lynching, I would conduct a series of keyword searches in four archives (Chronicling America, America's Historical Newspapers, NewspaperArchive, and ProQuest) using lynching related words as well as specific places, names, and dates. Based on a few articles in these archives, I generated a list of keywords that fit into the following categories: place names, full names (victims and other people named in an article), last names, unusual first or last names, races explicitly mentioned, and violence keywords regarding lynching and the specific form of violence used in the event. I then generated keyword searches for the period of up to 1 month following the lynching. The searches were for: any of the rare places, any rare names, any of the names and any of the violence keywords, any of the names and any of the race keywords, any of the names and any of the places, any of the places and any of the violence keywords, or any of the race keywords and any of the violence keywords. Then, I or RAs would go through all of the search results and record the results that match the lynching event.

By matching this to the full data obtained from the newspaper archives, it is possible to evaluate the extent of false positives and false negatives induced by using keywords to identify coverage, and how the rates change as a function of the window choice.

- **Hand Classification of Pages with Lynching Keywords**

  As part of the revisions to this paper, I also hired a team of RAs to hand classify the content of 2000 (1796 were completed) newspaper pages that (1) matched a keyword search for lynching-related words and (2) appeared within 7 days of a lynching event (This is discussed in more detail in D.4.1. But the RAs, as part of their classifications, coded whether the article discussed lynching in any way and *how* the article discussed lynching. This makes it possible to evaluate the composition of the "coverage" generated by the matching procedure I adopt in the paper.

First, I evaluate the choice of window when matching. Using the manually coded coverage of specific lynching events, I calculated the fraction of total coverage for each lynching that occurs within a given number of days since the event. Figure D4 shows that, on average, about 80% of coverage occurs within 7 days of a lynching, and nearly 100% within 14 days. This emphasizes that most coverage of an event occurs within a week. While there are certainly exceptions, extending the window further would create new problems when using keywords to classify coverage.

Second, I evaluate the performance of using keyword searches to classify whether newspapers cover a lynching or not. I match the hand-coded coverage of lynching events to the keyword search results. For each lynching event, I calculate the True Positive Rate (TPR, or sensitivity, is the fraction of newspaper issues manually coded as covering the event labeled as coverage using keywords and the given window) and the Positive Predictive Value (PPV, or precision, is the fraction of newspaper issues labeled as covering the event using keywords and the given window that are manually coded as coverage). In other words, TPR tells us how well using the keyword-window approach performs with respect to false negatives (higher

TPR implies fewer false negatives); PPV tells us how well the keyword-window approach performs with respect to false positives matches of coverage (higher PPV indicates fewer false positives).

Figure D5 shows the average TPR and PPV across different windows. Unsurprisingly, expanding the window increases the TPR, because it includes as coverage reports that take place longer after the event. However, this comes at a cost; as the window expands, PPV drops precipitously. Thus, there is a clear trade-off between avoiding false negatives and avoiding false positives (Figure D5). How do we decide a window that minimizes this tradeoff? Common ways of evaluating overall performance on TPR and PPV are to take the arithmetic, geometric, or harmonic average of the two.[16] Figure D6 shows that the harmonic average peaks 2 to 3 days out, the geometric average 6 days out, and the arithmetic average 13 days out. The main window of seven days is in the middle. But, based on these results, I test for the robustness of my results using windows between 3 and 11 days after a lynching.

These charts also permit evaluating the overall performance of the keyword-window approach (not just the relative performance of different choices of window). Figure D5 shows that both the TPR and PPV are fairly low: on average, the keyword-window method achieves a maximum TPR of just under 60 percent and a maximum PPV of just over 20 percent. The problem of false negatives largely appears to be an issue of digitization quality. Figure D7 shows the same validation exercise, only using newspapers from America's Historical Newspaper, which has the cleanest corpus of text. Here the TPR is nearly 80 percent. Thus, false negatives are probably mostly at random due to the quality of archived newspapers, the scans made, and the OCR technology used to transcribe them.

The problem of false positives is likely attributable to four issues. (1) The manual matching of coverage to specific lynching events took place between 2013 and 2014. But the keyword searches were executed in the summer and fall of 2015. NewspaperARCHIVE and ChroniclingAmerica continued to add new content during that time. As a result, some of the "false positives" may in fact be true coverage of the event that was not available when manual coding took place. This can be seen when looking at Figure D7: America's Historical Newspapers is rarely updated by comparison, and the maximum PPV was 0.3, rather than 0.2 for the overall sample. (2) Some false positives would result from keywords picking up news that is not in fact about lynching. This appears to be fairly rare. Based on the classification of newspaper pages printed within 7 days of a lynching that contain lynching keywords, 88 percent of these pages covered lynching in some way (and a few percent were simply unreadable). The main culprits of this were the phrase "Judge Lynch" referring to an actual person and "mob violence" or "mob law" referring to something else. (3) Some false positives could be coverage of other events that did not become lynching. Coders found 9% of these pages contained coverage of lynching attempts that were thwarted, 13% contained coverage of lynchings that were threatened but not attempted, and 10% contained coverage about "fears" that a lynching might occur. By contrast, nearly 45% of these pages contained coverage of actual lynching events that had happened. While some of this may have been of different events, it does suggest that a large plurality of the content picked up using the keyword approach actually refers to specific lynching events. (4) Finally, 20% of the pages classified discussed lynching editorially or covered lectures that did so. While this coverage would be classified as "false positives"

[16]In the machine learning literature, where concern for both TPR and PPV are common, the geometric and harmonic average of the two are called the F and G scores respectively.

because it didn't refer to specific lynching events, undoubtedly much of editorializing about lynching was in response to news of actual lynching events. Moreover, editorial comment on lynching is clearly relevant to my analysis of discourse; discarding it because it cannot be matched to a specific event runs contrary to the spirit of the argument I make.

Figure D4: Fraction of Lynching-Event Coverage Issued by Days Since Event

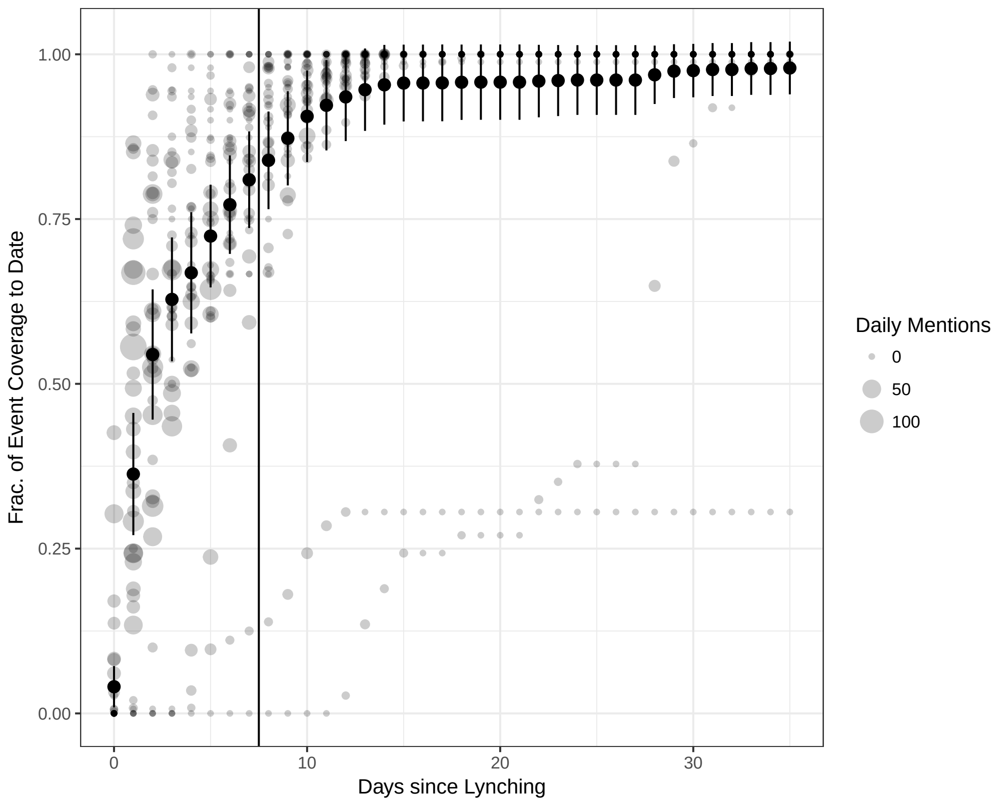

This figure shows the average fraction of coverage about specific lynching events completed within so many days since the event. Average reflects manual coding of coverage of 34 lynching events.

D.4 Discourse

Building on the data collection of newspaper *coverage* of lynchings, I also create data on the *content* or discourse about lynching in this coverage. These data, like the coverage data, make use of keywords for classification. In this section, I describe both the creation and validation of measures of pro- and anti-lynching discourse used in the paper.

Due to both the size of the data and the poor quality of the underlying text, I employed a dictionary-based classification scheme to indicate support for or opposition to lynching. To develop this dictionary, I first created a typology of major pro- and anti-lynching discourses from the historiography of lynching. This typology is in Table D3. Some of these discourses involve explicit endorsement/condemnation of lynching, while others are rhetorical devices or diagnostic frames that provide support for pro-/anti-lynching arguments. For each of these discourses, I created a list of keywords and phrases that corresponded to the discourse.

Figure D5: TPR and PPV for using keyword-windows to match coverage to events

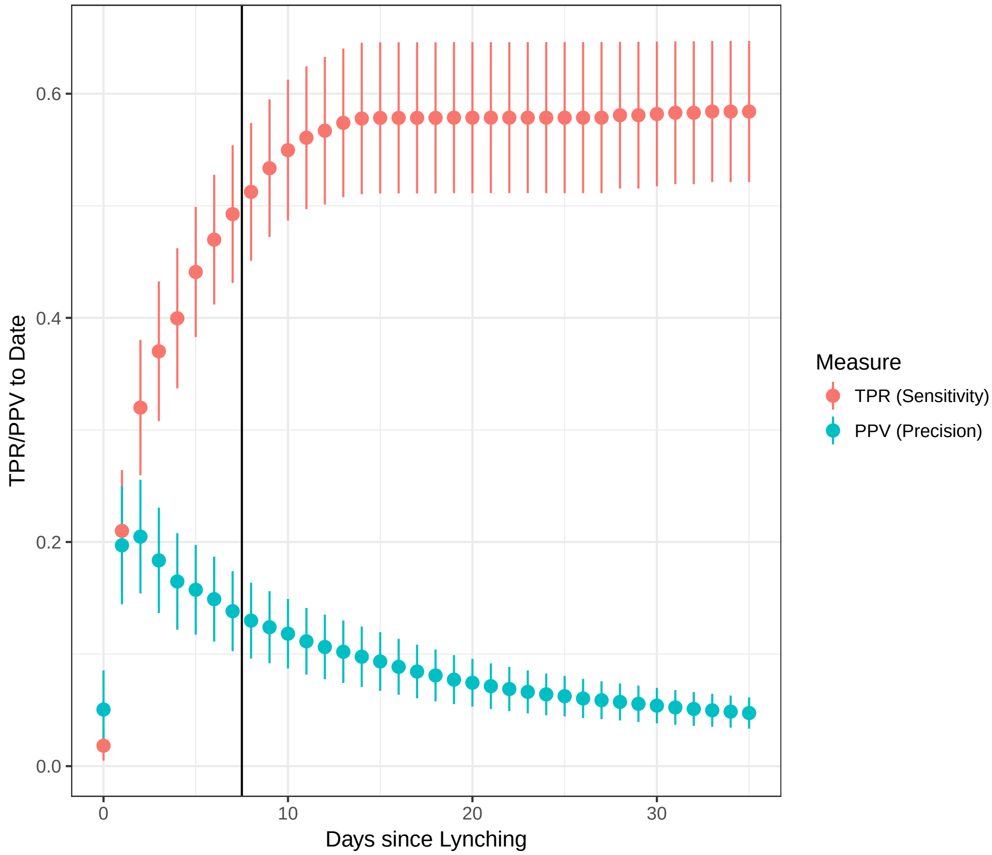

This figure shows the average TPR and PPV for classifying coverage of lynching events using lynching keywords and different windows (days since the event). Average reflects manual coding of coverage of 34 lynching events.

Like the discourses themselves, these derive from the historiography of lynching (particularly (Perloff 2000; Jean 2005; Wood 2009)) and my own reading of several thousand newspaper articles about lynching. The keywords used here also closely resemble those employed by Seguin (2016) to classify news coverage as pro- and anti-lynching. Table D3 also shows these corresponding keywords.

I then aggregate these keywords into a general index showing support and opposition to lynching. Following Grimmer and Stewart (2013), I create a simple dictionary index. In this index, I take the difference of the sum of anti-lynching keywords/phrases and the sum of pro-lynching keywords/phrases. For each article classified as "lynching coverage" as described above, I also calculate this overall lynching discourse measure. For a given article k, this computes the discourse score as the difference between sum of anti-lynching dictionary words and the sum of pro-lynching dictionary words Grimmer and Stewart (2013). Thus, n_a and n_p are the total number of keywords in the anti- and pro-lynching dictionaries. $Word_ik$ and $Word_jk$ indicate whether word i or j is present (1) or absent (0) in the article.

Figure D6: Performance of keyword-windows to match coverage to events: averaging of PPV and TPR

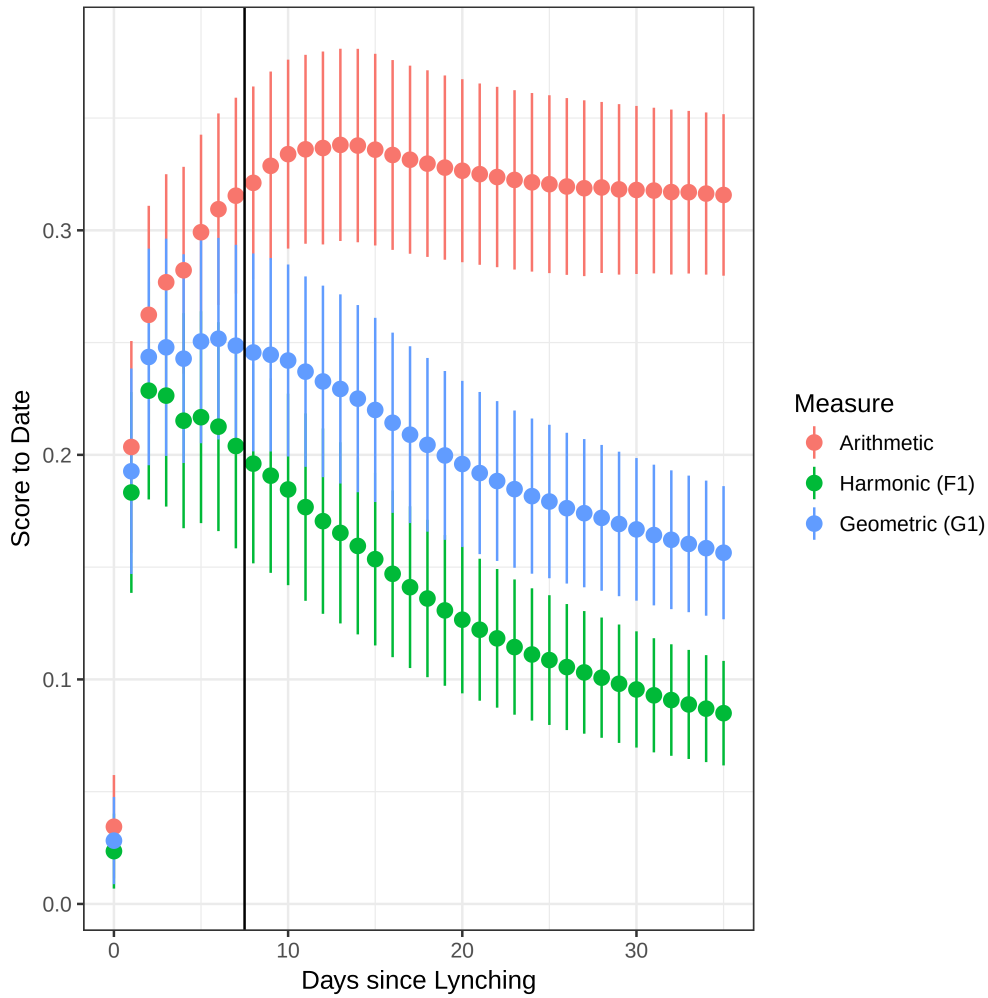

This figure shows the average of TPR and PPV (often called F- or G- scores) for classifying coverage of lynching events using lynching keywords and different windows (days since the event). Average reflects manual coding of coverage of 34 lynching events.

$$scaledDiscourse_k = \left(\sum_{i=1}^{n_a} AntiLynchingWord_{ik} \right) - \left(\sum_{j=1}^{n_p} ProLynchingWord_{jk} \right) \quad (2)$$

However, using keywords raises many question. Keywords might be a very noisy indicator. There is no guarantee that the keywords are used in the article(s) about lynching. Even if the keywords are used in an article about lynching, the meaning of words depends on context. Given that it is infeasible to read hundreds of thousands of articles, how can we be sure that the discourse keywords reflect meaningful differences? The solution I employ is to validate the use of keywords against manual classifications made by research assistants.

Table D3: Lynching Discourses and Associated Keywords

Discourse	**Description**	**Keywords**
Victim Dehumanized	Lynching apologists dehumanized the victims of lynching, particularly black men, as a way to justify the violence perpetrated against them.	brute, fiend, beast, monster, demon, savage
Victim "Guilty"	News accounts often assumed the guilt of those lynched, implying that the victim earned their punishment	murdered, slayed, stole, outrage, outraged, ravish, ravished, rape, raped, assault, assaulted, culprit, firebug, ravisher, assaulter, murderer, slayer, scoundrel, ruffian, desperado, rapist, villain, thief, outrager, confessed
Sexual Threat	Lynching was most widely justified as a response to sexual violence, particular by black men against white women. This was repeated so frequently that many assumed that almost all lynchings were for sexual crimes.	womanhood, white women, white womanhood, white woman, white maidens, white maiden, white girls, white girl, stain, honor, defend the honor, purity, protect the honor, rape, raped, rapist, ravish, ravished, ravisher, outrage, outraged, outrager, assault, assaulted, assaulter
Victim "Innocent"	Opponents of lynching questioned whether victims were guilty, given that the mob may have ulterior motives or simply lack the capacity to ascertain the truth.	wrong negro, wrong man, was innocent, maintained innocence, innocent man, innocent, forced confession, acquitted, alleged, allegedly, denied guilt
Shameful	Lynching was also rejected as uncivilized and barbaric, and shameful to communities and the nation as a whole.	backward, barbaric, barbarous, condemn, condemned, decried, decry, denounce, denounced, despicable, drunk, embarrassing, embarrassment, gruesome, heinous, stain, shameful, shame, uneducated, unruly
Lawlessness	Critics of lynching emphasized how lynching was a crime itself and threatened to undermine law and order.	unruly, anarchy, anarchic, lawless, lawlessness

Figure D7: TPR and PPV for using keyword-windows to match coverage to events: *America's Historical Newspapers* only

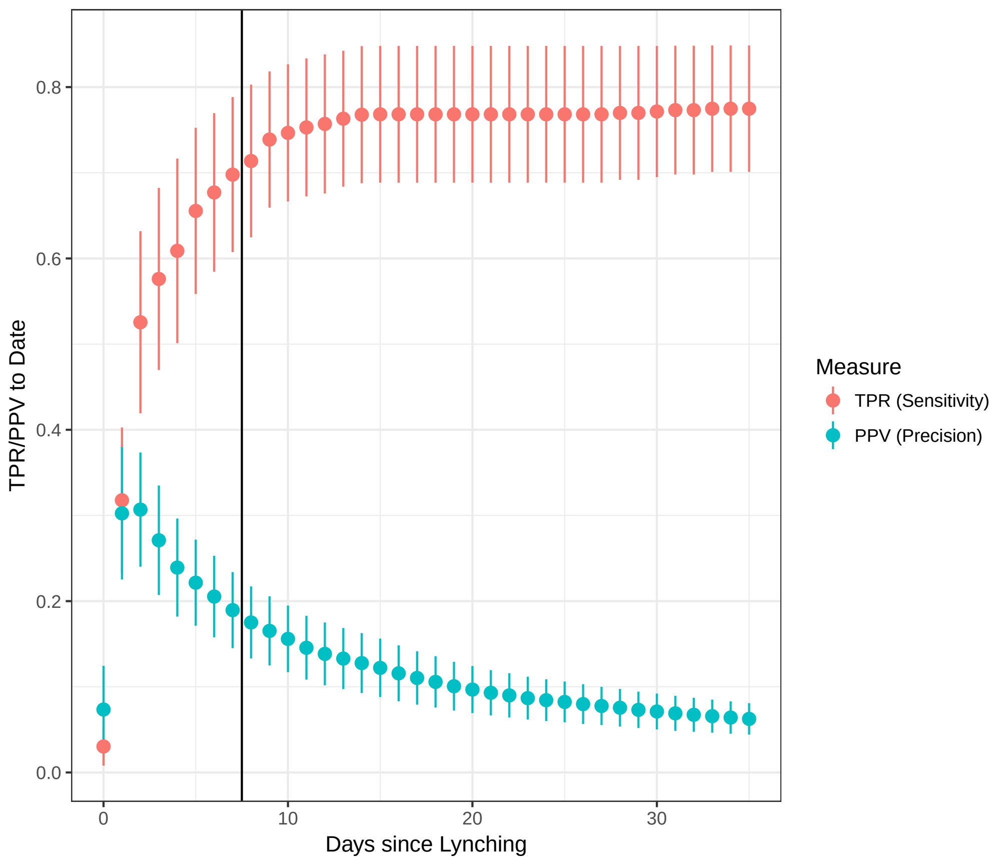

This figure shows the average TPR and PPV for classifying coverage of lynching events using lynching keywords and different windows (days since the event) when using newspapers from *America's Historical Newspapers* only. Average reflects manual coding of coverage of 34 lynching events.

D.4.1 Discourse Validation

To validate the discourse measures, I developed a coding scheme for the presence and absence of various pro- and anti-lynching discourses. I took a random sample of 2000 (1796 of which were completed) newspaper pages classified as "lynching coverage" based on the presence of lynching keywords and their occurrence within seven days of a lynching event. And I employed research assistants to read and classify these news pages using the schema I developed.

Coding Scheme I first created a coding scheme to help classify the presence and absence of various lynching discourses. Again, drawing on the historiography of lynching, I created a typology of pro- and anti-lynching discourses. To the the discourses I used when generating keywords, I also added the following:

Pro: Invocations of popular sovereignty; claims that the legal system was corrupt or inefficient; lynching is a "natural" response

Anti: Humanizing victims; Refuting claims of sexual threat; Inclusion of African American voices

To simplify the classification task and demand less interpretation by coders, I generated a few dozen questions that asked whether certain things did or did not appear on the page. Each RA conducted a test run of 25 articles to generate questions pertaining to the coding scheme. We then met and discussed answers and clarifications to these questions, and altered the coding instrument to improve clarity. Because the same newspaper might include both pro- and anti-lynching discourses, the coding scheme did not make any of answers to these questions mutually exclusive. The questions, their possible answers, whether they indicate pro/anti-lynching discourse, and which specific pro-/anti-lynching discourse the questions correspond to are listed in Table D4.

Table D4: Lynching Article Classification: Questions, Answers and Discourses

Main Question	Sub Question	Answers	Valence	Discourse
Look or search for lynching-related coverage on this page. Is lynching discussed anywhere on this page?		Can't tell (impossible to read) | No | Yes	N/A	
Can you read any of the news/editorial about lynching?		No, it is illegible. | Yes	N/A	
In what context does (do) the article(s) address lynching?		A specific lynching event that was attempted by stopped | A specific event in which no explicit lynching threat has occurred, but there is a fear of a possible lynching. | A specific lynching event that was threatened but has not (yet) happened | Editorial/Opinion/Lecture | A specific lynching event that has happened	N/A	
Which of the following are true of how the lynching victim was described?	Victim's name given (any victim's name given, if multiple victims)	No | Yes	Anti Lynching	Humanized
	Described as an animal, subhuman, or lacking reason or empathy (negative connotation: e.g. "savage", "brute", "fiend", "monster")	Yes | No	Pro Lynching	Dehumanize
	Victim is humanized (e.g. aspects of life prior to lynching discussed, victim's family mentioned, victim pleads or claims innocence, other people testify to victim's good character)	Yes | No	Anti Lynching	Humanized
	Described as criminal or of bad character (e.g. "notorious") prior to events leading to lynching	Yes | No	Pro Lynching	Guilt

	Victim's guilt asserted (e.g. "confessed", "guilty", "no doubt", described unambiguously as perpetrator of the crime – described as "murderer" without qualification)	Yes | No	Pro Lynching	Guilt
	Alleged crime was trivial: (e.g. talking back, self-defense. This must be obviously something that is not a criminal act.)	Yes | No	Anti Lynching	Innocence
	Victim explicitly "innocent" or guilt in question (crime is "alleged", "uncertain" if person is correct)	Yes | No	Anti Lynching	Innocence
	Victim's alleged "crime" is sexual in nature (relating in some way to interactions across sexes, even if not a crime by legal standards. Does not need to be explicitly "assault", "rape", "outrage")	Yes | No	Pro Lynching	Sexual Threat
Which of the following are true of how the mob (lynchers) was described?	Mob is portrayed as composed of the entire community (e.g., "the men of the community" or "citizens" instead of mob)	Yes | No	Pro Lynching	Popular Sovereignty
	Mob described as marginal group (small group, outsiders, wore masks, acted in secret)	Yes | No	Anti Lynching	Shameful/Barbaric
	Mob portrayed as orderly (e.g. mob or their actions described as quiet, calm, deliberate, sober, rational)	Yes | No	Pro Lynching	Popular Sovereignty
	Mob portrayed as having unruly behavior (disorderly, damaging property, out of control, drunk)	Yes | No	Anti Lynching	Lawless
	Mob is portrayed as backward (poor, uneducated)	Yes | No	Anti Lynching	Shameful/Barbaric
	Mob is portrayed as local elite (e.g. "best citizens", "town leaders", respected)	Yes | No	Pro Lynching	Popular Sovereignty
	Mob described as enforcing the law (e.g. mob held trial, actions described as justice	Yes | No	Pro Lynching	Popular Sovereignty
	Mob characterized as, e.g., "lawless", "anarchic", or "criminal" (as threat to law and order, public safety)	Yes | No	Anti Lynching	Lawless

	Mob described as "barbaric" or "savage" (described in terms suggesting they are uncivilized, less than fully human)	Yes | No	Anti Lynching	Shameful/Barbaric
	Violent actions of mob use passive voice (person was lynched, person was burned)	Yes | No	Pro Lynching	Natural Response
	Violent actions of mob use active voice (e.g., mob broke open the cell, mob dragged the body, mob shot him ..., e.g.)	Yes | No	Anti Lynching	Shameful/Barbaric
	Violence of mob described in detail (explicit detail about the violence and its effects on the body, not just "was shot" or "was burned" or "riddled with bullets")	Yes | No	Anti Lynching	Shameful/Barbaric
	Mob use of violence is portrayed as "natural" or "inevitable" response to alleged crime (must be clear about inevitability of lynching)	Yes | No	Pro Lynching	Natural Response
	Mob portrayed as having support of local community (e.g., crowd was large, no intervention to stop lynching, police do not attempt to stop the event)	Yes | No	Pro Lynching	Popular Sovereignty
	Mob reported to have received criticism in their community (community members denounce after the fact, attempts to stop the lynching, police work to prevent lynching)	Yes | No	Anti Lynching	Lawless
Which of the following is true about the article's discussion of the lynching event:	Lynching event attributed to failure to convict victim at trial. (It must explicitly state that the lynched person was tried for a crime and that they were not convicted and this was invoked as a reason to lynch)	Yes | No	Pro Lynching	Courts Inefficient
	Article (or person in it) explicitly states that the specific lynching was justified or the right thing	Yes | No	Pro Lynching	Explicit Justification
	Article (or person in it) explicitly condemns the specific lynching	Yes | No	Anti Lynching	Explicit Condemnation

	Police/law enforcement described as attempting to stop lynching (e.g. protect lynching victim, refuse to open jail cell, attempt to disperse mob)	Yes | No	N/A	
	Police/law enforcement described as using major effort to stop lynching (police/law enforcement persist despite threat of violence against them, move lynching victim to a different county, arrest/detain members of the mob)	Yes | No	Anti Lynching	Lawless
Which of the following is true about the article's discussion sexual violence and lynching?	article (or person in it) describes lynching as "natural" or "understandable" response to crime (sexual transgression, or any alleged crime involving interaction between sexes)	Yes | No	Pro Lynching	Sexual Threat
	article (or person in it) mentions threat of rape/invoke need to protect women/chivalry	Yes | No	Pro Lynching	Sexual Threat
	articles (or person in it) states lynching only way to stop rape	Yes | No	Pro Lynching	Sexual Threat
	articles (or person in it) suggests lynching will stop when rape stops/ stopping rape will stop lynching	Yes | No	Pro Lynching	Sexual Threat
	article (or persons in it) discusses women rejecting protection from lynching	Yes | No	Anti Lynching	No Sexual Threat
	article (or person in it) states alleged interracial rapes are actually consensual	No	Anti Lynching	No Sexual Threat
	article (or person in it) rejects connection between rape and lynching (specific to the event, or in general)	Yes | No	Anti Lynching	No Sexual Threat
Which of the following is true about the article's discussion of law and order and lynching?	article (or person in it) describes legal system as too slow/ineffective/corrupt	Yes | No	Pro Lynching	Courts Inefficient
	article (or person in it) attributes lynching (in general) to failure to convict	Yes | No	Pro Lynching	Courts Inefficient
	article (or person in it) describes lynching as lawless or anarchic (threat to public order, law and order)	Yes | No	Anti Lynching	Lawless

	article (or person in it) describes lynching as illegal and/or calls for legal action against lynchers	Yes | No	Anti Lynching	Lawless
	article (or person in it) claims that lynching encourages lawlessness/crime	Yes | No	Anti Lynching	Lawless
	article (or person in it) argues that lynching denies defendants rights to a trial	Yes | No	Anti Lynching	Lawless
Which of the following is true about the article's discussion of race and lynching?	article includes comment/perspective on lynching from African American (quote or paraphrase of statement given by person identified as "black", "colored", or "negro". Famous people would include: Frederick Douglass, Booker T. Washington, Ida Wells, W. E. B. Du Bois)	Yes | No	Anti Lynching	Inclusion
	article includes comment/perspective on lynching from an organization (or its representative) identified as being African American	Yes | No	Anti Lynching	Inclusion
	article (or person in it) advocates violent resistance to lynching	Yes | No	Anti Lynching	Inclusion
	article (or person in it) explicitly portrays lynching (in general, not a specific case) as a racial issue or problem (e.g., lynching is described as a "Negro problem" or lynching elicits fear of "race war")	Yes | No	N/A	
	article (or person in it) explicitly portrays lynching as violence based on or motivated by race (e.g., used to suppress African Americans, targeted because of race)	Yes | No	Anti Lynching	Inclusion
Which of the following is true about the article?	article (or person in it) explicitly states that lynching (even if only in some circumstances) is justified	Yes | No	Pro Lynching	Explicit Justification
	article (or person in it) expicitly condemns lynching (for any reason)	Yes | No	Anti Lynching	Explicit Condemnation
	article (or person in it) explicitly states that lynching is barbaric/uncivilized	Yes | No	Anti Lynching	Shameful/Barbaric

	article (or person in it) explicitly states that lynching is a source of shame or embarassment for the town, state, country	Yes | No	Anti Lynching	Shameful/Barbaric
	article (or person in it) explicitly states that lynching is bad for business	Yes | No	Anti Lynching	Shameful/Barbaric
	article (or person in it) accuses lynchers/supporters of lynching of hypocrisy	Yes | No	Anti Lynching	Explicit Condemnation

Reliability Of the 1796 pages read by RAs, 736 were triple-coded. This was designed to permit evaluation of coder reliability. Unfortunately, many of the items of the coding scheme were rare. This created two problems: first, coders were not exposed to enough examples in the trial period to generate improvements in the coding scheme; second, with highly imbalanced classifications, inter-coder reliability scores like Krippendorff's alpha tend to have lower scores as agreement is likely to occur by chance. Overall, I find that there was widespread agreement about the absence of various discourse attributes, but less agreement on their presence. Thus, the inter-coder reliability scores were low (not shown).

Nevertheless, these manual codes are nevertheless helpful for three reasons:

1. Despite the lack of reliability, it is nevertheless possible to measure the correlation between keyword and manual classifications for specific discourses *and* to evaluate the overall lynching discourse measure.

2. The manual classifications and keywords measure the *dependent variable*. Insofar as this suffers from random measurement error, then my analyses should not be biased, only have increased standard errors. Thus, even if there is noise in both the keywords and the manual classifications, we should be reassured if their relationship is strongly statistically significant.

3. The triple coding of the classifications means that the uncertainty of the manual classification will be directly reflected in measuring the correlation between the manual classification and the keyword classification.

If there is still a correlation between these manually coded measures of pro- and anti-lynching discourse and keywords, that is evidence that they keywords capture meaningful variation in discourse.

D.4.2 Evaluation

I evaluate the performance of the keyword measures by correlating the keyword counts against the manual classifications. I do this in two steps; first, I examine whether the count of keywords for each specific discourse correlated to manual classifications of articles into that discourse. I simplify this task by summing the 'yes' answers across questions corresponding to each specific pro-/anti- lynching discourse. Second, I examine whether the sum of all anti-lynching keywords, sum of all pro-lynching keywords, and the difference of these two sums correlate with differences in manual classification of pro- and anti-lynching discourse.

Pro-Lynching Discourse Table D5 shows the Spearman rank correlation between counts of keywords corresponding to the dehumanization, guilty, and sexual threat pro-lynching discourses and the count of manually coded 'yeses' for questions corresponding to these three discourses and to explicit justification. While each of the pro-lynching discourse keywords are positively and significantly correlated with all of the manual codes, these correlations are strongest for the corresponding discourses: dehumanizing keywords have their strongest correlation with dehumanizing manual classifications; guilt keywords with the guilt classifications; sexual threat keywords with the sexual threat classifications. This shows that pro-lynching keywords are picking up the expected dimensions of discourse. Table D6 shows that the pro-lynching keywords primarily correlate with the pro-lynching keywords, and *not* anti-lynching discourse. While there is a positive correlation between articles classified as portraying lynching as shameful/barbaric and guilt and sexual threat keywords, this is a much weaker relationship. And guilt and sexual threat keywords are *negatively* related to explicit condemnations of lynching.

Table D5: Rank correlation of keyword and manual coding indices (count) of pro-lynching discourses

	Dehumanize MC	Guilt MC	Sexual Threat MC	Explicit Just. MC
Dehumanize KW	0.17***	0.05**	0.04*	0.08***
Guilt KW	0.11***	0.25***	0.18***	0.01
Sexual Threat KW	0.06***	0.21***	0.27***	0.01

Spearman rank correlation of pro-lynching discourse indices:
count of keywords (KW) and manually coded (MC) attributes.
N is 1606 pages of 'lynching' coverage.
*$p < 0.05$; **$p < 0.01$; ***$p < 0.001$

Table D6: Rank correlation of pro-lynching keywords and manually coded anti-lynching attributes (count)

	Lawless MC	Innocent MC	Shame/Barbaric MC	Explicit Condemn. MC
Dehumanize KW	0.00	-0.07***	0.03	0.03
Guilt KW	-0.03	0.01	0.05**	-0.12***
Sexual Threat KW	-0.03	0.03	0.06***	-0.05**

Spearman rank correlation of pro-lynching and anti-lynching discourse indices:
count of keywords (KW) and manually coded (MC) attributes.
N is 1606 pages of 'lynching' coverage.
*$p < 0.05$; **$p < 0.01$; ***$p < 0.001$

Anti-Lynching Discourse Table D7 shows the Spearman rank correlation between counts of keywords corresponding to the lawlessness, innocence, and shame/barbaric anti-lynching discourses and the count of manually coded 'yeses' for questions corresponding to these three discourses and to explicit condemnation. While the overall correlation of anti-lynching keywords to their corresponding discourses are weaker than for pro-lynching discourses, the substantive pattern is the same: lawless keywords, innocence keywords, and shame/barbaric keywords are all positively and significantly correlated with the manual classification of articles into their corresponding discourse. And lawless and shame keywords are also positively and significantly correlated with explicit condemnation of lynching. Conversely, Table D8 shows the correlation of anti-lynching keywords with manual classification of pro-lynching discourses. Lawless keywords are significantly and negatively correlated with pro-lynching classifications. Shame keywords are weakly and not significantly correlated with pro-lynching classifications. But, innocent keywords are positively and significantly correlated with some pro-lynching discourses. In fact, this is as strong as innocent keywords are correlated with the "innocence" manual classifications.

Table D7: Rank correlation of keyword and manual coding indices (count) of anti-lynching discourses

	Lawless MC	Innocent MC	Shame/Barbaric MC	Explicit Condemn. MC
Lawless KW	0.14***	-0.02	0.09***	0.14***
Innocent KW	0.01	0.06***	0.00	-0.01
Shame/Barbaric KW	0.05*	-0.03	0.08***	0.10***

Spearman rank correlation of anti-lynching discourse indices:
count of keywords (KW) and manually coded (MC) attributes.
N is 1606 pages of 'lynching' coverage.
* $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

Table D8: Rank correlation of anti-lynching keywords and manually coded pro-lynching attributes (count)

	Dehumanize MC	Guilt MC	Sexual Threat MC	Explicit Just. MC
Lawless KW	-0.06**	-0.07***	-0.07***	-0.03
Innocent KW	-0.01	0.04*	0.06**	-0.02
Shame/Barbaric KW	0.03	0.01	0.02	0.04*

Spearman rank correlation of pro-lynching and anti-lynching discourse indices:
count of keywords (KW) and manually coded (MC) attributes.
N is 1606 pages of 'lynching' coverage.
* $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

Overall Discourse Index Table D9 shows the correlation of all anti-lynching and pro-lynching manual classifications with each set of pro- and anti-lynching keywords, all pro- and all anti- lynching keywords, and the index used in the paper (Anti- minus Pro-lynching keywords). Among anti-lynching keywords, all but the innocence keywords are positively and significantly correlated with anti-lynching classifications and either negatively or not significantly correlated with pro-lynching keywords. Among pro-lynching keywords, all are significantly and positively correlated with all pro-lynching classifications and negatively correlated with anti-lynching classifications. This is reassuring: pro-lynching keywords pick up on actual discourses that justify or endorse lynching; anti-lynching keywords pick up on discourses that refute or condemn lynching. The exception are keywords associated with innocence. As a result, in robustness checks I consider measures of lynching discourse that also exclude these keywords.

The overall keyword measures are also reassuring. The sum of all pro-lynching keywords is strongly and positively correlated with pro-lynching classifications (negative with anti-lynching classifications), and the sum of all negative keywords is positively and significantly correlated with anti-lynching classifications (and not correlated with pro-lynching classifications). The keyword index used in the paper works as intended as well: it is positively and significantly correlated with anti-lynching classifications and negatively and significantly correlated with pro-lynching classifications. It is notable, however, that there is a stronger correlation with the pro-lynching classifications.

To assess whether this is due to noise, or whether there is a difference in the slope, I regressed (standardized) pro- and anti-lynching manual classifications on the keyword index used in the paper, including dummies for the total number of keywords matched. Table D10 shows that after accounting for the total number of keywords, increases in the lynching keyword index are associated with increases in anti-lynching classifications and decreases in pro-lynching classifications of similar magnitude. That is, a increase of one anti-lynching keyword vs pro-lynching keywords result in an approximately 0.09 standard deviation increase anti-lynching classifications and 0.11 standard deviation decrease in pro-lynching classifications. Not only, then, does the keyword index capture the relevant dimensions of discourse, but it appears to be related to similarly sized shifts in pro- and anti-lynching discourse.

D.5 Census

Census data was taken from the Historical, Demographic, Economic, and Social Data: The United States, 1790-1970 (ICPSR 3). Census data was matched to county boundaries from the year 2000 (to align with the data on the railroad network). This was done using the following approach: for each county in the census file for a given decade, I matched it to a county sharing the same name in the same state in 2000. If this match existed, then the 2000 county received the entire census record of the name match from the census. If such a name match did not exist between a county in 2000 and a county in the census year, then I used the Newberry Library's Atlas of Historical County Boundaries to create weights for counties in each of the census years between 1870 and 1940 where the weight reflects the area of the county that intersects a particular county from the year 2000. Then, I resolved each census year into 2000 county boundaries by taking the sum of count variables. When taking the sum, I weighted by the fraction of the census county contained in the year-2000 county. If there

Table D9: Rank correlation of overall lynching discourse keywords and manually coded attributes

	All Anti-Lynching MC	All Pro-Lynching MC
Lawless KW	0.13∗∗∗	-0.09∗∗∗
Innocent KW	0.02	0.04∗
Shame/Barbaric KW	0.07∗∗∗	0.02
All Anti-Lynching KW	0.09∗∗∗	0.02
All Anti-Lynching KW (alt)	0.10∗∗∗	-0.02
Dehumanize KW	-0.01	0.06∗∗∗
Guilt KW	-0.03	0.24∗∗∗
Sexual Threat KW	0.00	0.23∗∗∗
All Pro-Lynching KW	-0.02	0.27∗∗∗
(Anti - Pro) KW	0.07∗∗∗	-0.28∗∗∗

Spearman rank correlation of pro-lynching and anti-lynching discourse indices: count of keywords (KW) and manually coded (MC) attributes.
N is 1606 pages of 'lynching' coverage.
*$p < 0.05$; **$p < 0.01$; ***$p < 0.001$

Table D10: Relationship between keyword discourse measure and manual coding

	Anti-Lynching MC		Pro-Lynching MC	
	(1)	(2)	(3)	(4)
$(Anti - Pro)$ Keywords	0.028***	0.085***	−0.117***	−0.105***
	(0.008)	(0.010)	(0.008)	(0.010)
Total Keywords FE		X		X
N	2,873	2,873	2,873	2,873
Adjusted R^2	0.004	0.036	0.073	0.076

*p < .05; **p < .01; ***p < .001
Estimates obtained using OLS. Dependent variables have been standardized.

was no data in a county that overlapped the year-2000 boundary for a given census year, then the data was taken as missing.

This resulted in a panel of year-2000 county boundaries in several census years. To obtain yearly data, I linearly interpolated between censal years to obtain an annual panel.

D.6 African American Newspapers

In robustness checks on the relationship between distance and coverage of lynching, I show that the results cannot be explained by whether a paper is white or published by African Americans. I obtain this data from the Library of Congress U.S. Newspaper Directory. This includes bibliographic and historical details on newspapers from the United State. One piece of information that is included is whether the paper was published by African Americans. I match newspapers from my digitized archives to this list by paper name, place of publication, and years of publication.

References

Abbott, Richard (2004). *For Free Press and Equal Rights: Republican Newspapers in the Reconstruction South.* Athens, GA: University of Georgia Press.

Angrist, Joshua D, and Jörn-Steffen Pischke (2008). *Mostly Harmless Econometrics: An Empiricist's Companion.* Princeton University Press.

Atack, Jeremy (2013). On the Use of Geographic Information Systems in Economic History: The American Transportation Revolution Revisited. *The Journal of Economic History 73*(2), 313–338.

Beck, E. M., James Massey, and Stewart E. Tolnay (1989). The Gallows, the Mob, and the Vote: Lethal Sanctioning of Blacks in North Carolina and Georgia, 1882 to 1930. *Law \& Society Review 23*(2), 317–31.

Beck, E. M., Stewart E. Tolnay, and Amy Bailey (2016). Contested Terrain: The State versus Threatened Lynch Mob Violence. *American Journal of Sociology 121*(6), 1856–1884.

Bell, Andrew, and Kelvyn Jones (2015). Explaining Fixed Effects: Random Effects Modeling of Time-Series Cross-Sectional and Panel Data. *Political Science Research and Methods 3*(1), 133–153.

Blondheim, Menahem (1994). *News Over the Wires: The Telegraph and the Flow of Public Information in America, 1844–1987.* Cambridge, MA: Harvard University Press.

Brundage, W. Fitzhugh (1993). *Lynching in the New South: Georgia and Virginia, 1880-1930.* Chicago: University of Illinois Press.

Campney, Brent (2013). Ever Since the Hanging of Oliphant": Lynching and the Suppression of Mob Violence in Topeka, Kansas. *Great Plains Quarterly 33*(2), 71–86.

Carpenter, Daniel (2000). State Building through Reputation Building: Coalitions of Esteem and Program Innovation in the National Postal System, 1883–1913. *Studies in American Political Development 14*(2), 121–155.

Carrigan, William (2005). *The Making of a Lynching Culture: Violence and Vigilantism in Central Texas, 1836–1916.* Champaign, IL: University of Illinois Press.

Carrigan, William, and Clive Webb (2013). *Forgotten Dead: Mob Violence against Mexicans in the United States, 1848–1928.* New York, NY: Oxford University Press.

Clark, Thomas (1948). *The Southern Country Editor.* Indianapolis, IN: Bobbs-Merrill.

Fuller, Wayne (2003). *Morality and the Mail in Nineteenth Century America.* Chicago, IL: University of Illinois Press.

Garcia-Jimeno, Camilo, Angel Iglesias, and Pinar Yildirim (2018). Women, Rails, and Telegraphs: An Empirical Study of Information Diffusion, Social Interactions, and Collective Action.

Gentzkow, Matthew, and Jesse M Shapiro (2014). *United States Newspaper Panel, 1869-2004 (ICPSR 30261).* Inter-University Consortium for Political and Social Research.

Grimmer, Justin, and Brandon Stewart (2013). Text as Data: The Promise and Pitfalls of Automatic Content Analysis Methods for Political Texts. *Political Analysis 21*(3), 267–297.

Hagen, Ryan, Kinga Makovi, and Peter Bearman (2013). The Influence of Political Dynamics on Southern Lynch Mob Formation and Lethality. *Social Forces 92*(2), 757–787.

Huntzicker, William (1999). *The Popular Press: 1833–1865.* Westport, CT: Greenwood Press.

Jean, Susan (2005). "Warranted" Lynchings: Narratives of Mob Violence in White Southern Newspapers, 1880-1940. *American Nineteenth Century History 6*(3), 351–372.

John, Richard (2009). *Spreading the News: The American Postal System from Franklin to Morse* (3rd ed.). Cambridge, MA: Harvard University Press.

Keil, Thomas, and Gennaro Vito (2009). Lynching and the Death Penalty in Kentucky, 1866–1934: Substitution or Supplement? *Journal of Ethnicity in Criminal Justice 7*(1), 53–68.

Kielbowicz, Richard (1982). Newsgathering by Printers' Exchanges Before the Telegraph. *Journalism History 9*(2), 42–48.

Kielbowicz, Richard (1983). The Press, Post Office, and Flow of News in the Early Republic. *Journal of the Early Republic 3*(3), 255–280.

Kielbowicz, Richard (1985). Speeding the News by Postal Express, 1825–1861: The Public Policy of Privileges for the Press. *The Social Science Journal 22*(1), 49–63.

Kielbowicz, Richard (1986). Origins of the Second-Class Mail Category and the Business of Policymaking, 1863–1879. *Journalism Monographs 96*, 1–26.

Kielbowicz, Richard (1987). News Gathering by Mail in the Age of the Telegraph: Adapting to a New Technology. *Technology and Culture 28*(1), 26–41.

Kielbowicz, Richard (1989). *News in the Mail: The Press, Post Office, and Public Information, 1700–1860s.* New York, NY: Greenwood Press.

Kielbowicz, Richard (1990). Postal Subsidies for the Press and the Business of Mass Culture, 1880–1920. *The Business History Review 64*(3), 451–488.

Kielbowicz, Richard (2016). Electrifying news!: Journalists, audiences, and the culture of timeliness in the United States, 1840–1920. *Time & Society*.

King, Marissa, and Heather Haveman (2008). Antislavery in America: The Press, the Pulpit, and the Rise of Antislavery Societies. *Administrative Science Quarterly 53*(3), 492–528.

Perez-Cervantes, Fernando (2014). Railroads and Economic Growth: A Trade Policy Railroads and Economic Growth: A Trade Policy Approach. In *Bank of Mexico Working Papers.* Bank of Mexico.

Perloff, Richard A. (2000). The Press and Lynchings of African Americans. *Journal of Black Studies 30*(3), 315–330.

Pfeifer, Michael J (2011). *The Roots of Rough Justice: Origins of American Lynching.* Champaign, IL: University of Illinois Press.

Pred, Allan (1973). *Urban Growth and the Circulation of Information: The United States System of Cities, 1790–1840.* Cambridge, MA: Harvard University Press.

Russo, David (1980). The Origins of Local News in the U.S. Country Press, 1840–1870s. *Journalism Monographs 1*(65), 1–43.

Seguin, Charles (2016). *Making a National Crime: The Transformation of US Lynching Politics 1883–1930.* Dissertation, University of North Carolina at Chapel Hill, Chapel Hill, NC.

States, United (1885). *History of the Railway Mail Service: A Chapter in the History of Postal Affairs in the United States.* Washington, DC: United States Government Printing Office.

Tolnay, Stewart E., and E. M. Beck (1995). *A Festival of Violence: An Analysis of Southern Lynchings, 1882-1930.* Chicago, IL: University of Illinois Press.

Wood, Amy Louise (2009). *Lynching and Spectacle: Witnessing Racial Violence in American, 1890-1940.* Chapel Hill, NC: University of North Carolina Press.

Wright, George C. (1990). *Racial Violence in Kentucky, 1865–1940: Lynchings, Mob Rule, and "Legal Lynchings".* Baton Rouge, LA: LSU Press.