The Specter of Slavery

Online Appendix

Patrick Crawford

pcrawfo6@gmu.edu

November 5, 2026

Click here for the most recent version

Appendix

Contents

A Data Construction and Topic Measurement **3**

A.1 Corpus Assembly and Filtering . . . 3

A.2 OCR Cleaning and Preprocessing . . . 4

A.3 Syndicated Article Reuse . . . 5

A.4 Keyword Dictionaries . . . 6

A.5 Page- and Article-Level Co-occurrence . . . 7

A.6 Panel Construction . . . 8

B Text Corpus and OCR Quality Filtering **10**

B.1 The Digitized Corpus . . . 10

B.2 Union-State Restriction . . . 10

B.3 OCR Quality Filtering . . . 11

B.4 Filter Bias Diagnostics . . . 12

C Additional Results **19**

C.1 Full Event-Study Coefficients . . . 19

C.2 Grouped Pre-Election Windows . . . 19

C.3 Full Party-by-Cycle Table . . . 20

C.4 Papers with a Rival in Town . . . 20

C.5 What the Campaign Displaced . . . 21

C.6 County Vote Shares . . . 23

D Robustness Checks **26**

D.1 The Retired Embedding Measure . . . 26

D.2 The Emancipation Proclamation . . . 26

D.3 Alternative Coverage Measures . . . 26

D.4 The Earlier Slavery Dictionary . . . 27

D.5 Time-Index Sign Convention . . . 27

References **28**

A Data Construction and Topic Measurement

A.1 Corpus Assembly and Filtering

LCCN whitelist. Newspaper inclusion starts from a hand-compiled metadata table keyed by Library of Congress Control Number (LCCN) that records each newspaper's state, county, city, first and last issue dates, primary language, and partisan alignment for each of the three election years. The alignments are the author's hand codes from each paper's own front-page ticket notices (the data section of the paper). Two restrictions are applied at the metadata stage, before any article text is touched. First, only newspapers whose recorded language is English are retained; non-English titles—principally German-language papers common in the Midwest—are excluded because the text pipeline downstream (Section A.4) operates on English. Second, the analysis panel keeps a title-month only if it can be attributed to a congressional election in the title's state within twelve months (Section A.6). The panel applies no Union-state list. In practice the attribution rule keeps out most papers from seceded states, but not all: the cleaned metadata contains 20 titles from states that seceded, and the newspaper panel includes papers from Tennessee, attributed to its August 1861 congressional vote, and, in 1864, from occupied Virginia and South Carolina. A hardcoded list of Union-state names (enumerated in Section B.2) is used only in the OCR bias diagnostic of Section B.

Date range. The underlying AmericanStories corpus (Section B.1) covers article-level text from 1860–1865. The newspaper-month panel construction additionally restricts to calendar years 1859–1865 when assembling LCCN×month rows; the 1859 lower bound is a defensive bound in the panel-construction step rather than a binding restriction, since no article text exists before 1860 in the underlying corpus.

Newspapers with multiple metadata rows. Among the retained titles, 4 LCCNs appear twice in the raw metadata, 3 of them because the newspaper moved to a new city within the same county partway through the sample period. These are collapsed to a single row per LCCN, taking state, county, and city from the earliest-dated row (the paper's original location), the minimum `first_issue` and maximum `last_issue` across rows as the paper's full operating window, and the first non-missing (non-`"?"`) partisan alignment value across rows for each election year.

Partisan affiliation coding. The metadata table codes each newspaper's alignment separately for the 1860, 1862, and 1864 elections, because a paper's partisan identity could shift across the sample period. The analysis panel maps the hand-coded ticket labels to parties with an explicit table: `"Republican"`, `"Union/Republican"`, and `"National Union"` map to Republican; `"Democratic"`, `"National Democratic"`, and `"Southern Democratic"` map to Democratic; `"Constitutional Union"` and `"None"` are assigned to neither party. Every label in the data appears in the table. Each newspaper-month takes the affiliation of the cycle of its attributed congressional election. An earlier version of the coding matched labels with a

pattern that also coded 1860 Constitutional Union (Bell) papers as Republican; that error is corrected in every estimate in the paper.

Issue density. There is no minimum-issue-count or minimum-article-count threshold applied to determine which LCCNs enter the panel: any newspaper passing the language and attribution rules above is included regardless of how many articles it contributes. This matters because the within-newspaper baseline used to construct R_{jt} (the newspaper's own mean and standard deviation in months far from any election) is unreliable when computed from very few observations. Newspapers whose baseline rests on fewer than three such months are left undefined for R_{jt} (Section A.6) but retained for the absolute measures. No minimum-density cut is applied to the panel as a whole.

A.2 OCR Cleaning and Preprocessing

Source and granularity. The raw text comes from the AmericanStories dataset (Section B.1), which is already segmented at the article level rather than the page level: each record's `article_id` encodes the source page and an article index within that page, and LCCN is recovered by splitting `article_id` on underscores and taking the fourth field. No further page-to-article segmentation is performed downstream; the unit of observation throughout the pipeline is the AmericanStories article record as delivered.

Language filtering is newspaper-level, not article-level. There is no per-article language detection model in the pipeline. English-language restriction happens once, at the LCCN whitelist stage (Section A.1), on the assumption that a newspaper's language of publication is stable over its run—reasonable for this period, since bilingual mastheads were rare and a paper's language was a fixed editorial choice tied to its readership, not something that varied issue to issue.

OCR correction layers. Raw AmericanStories text is corrected in four layers during corpus preparation, from cheapest/most-mechanical to most aggressive. (1) Ligature normalization maps the single-character typographic ligatures for fi, fl, ff, ffi, and ffl, produced by period typesetting and OCR, to their plain-letter equivalents; this always runs. (2) A regex table of `tl`/`li`/`wl` substitutions corrects a specific, high-frequency OCR confusion in which the long-s/rounded-letter shapes common in 19th-century typefaces are misrecognized (`tlie` → `the`, `wliat` → `what`, `otlier` → `other`, and similar patterns); this also always runs and is independent of any dictionary. (3) A line-join heuristic reattaches words hyphenated across a line break in the original typesetting. (4) Word-level correction, which replaces words with dictionary or model suggestions, is not applied to the corpus used here. Its correction dictionary is built from modern English, and it risks "correcting away" period vocabulary and spelling (`freedmen`, `slaveholder`) that carries issue signal; a comparison against hand-typed transcriptions found that it made recognition errors worse, not better. The corpus used

in the paper therefore applies layers (1)–(3) only.

Quality metrics and the minimum-length filter. OCR quality is scored with the `ocroscope` package, which reports a character-recognition score and the share of non-alphabetic characters for each article; together with article length, these define the quality filter. In the earlier version of the corpus, the one the diagnostics in Section B describe, an article passed if it had at least 20 words, an OCR quality score of at least 95, and a non-alphabetic share of at most 10 percent; failing articles were routed to a separate low-quality split rather than deleted, which is the split Section B.3 analyzes. The corpus used in the paper scores the deterministically cleaned text once with `ocroscope` and applies the same kind of length and quality filter, but its thresholds were set by inspection and are not documented.

A.3 Syndicated Article Reuse

Concern. Wire-service copy and popular editorials were routinely reprinted verbatim, or near-verbatim, across many newspapers in this period. Because the coverage measures are *shares*—the fraction of a newspaper-month's articles that contain two different dictionary entries for slavery or for campaign coverage (Section A.4)—uniform reprinting across a paper's issues does not by itself bias the share: a syndicated item counts once in the numerator and once in the denominator. Reprinting matters for the measure only if syndicated content is disproportionately on or off slavery relative to a paper's original content, or if the reprint volume moves systematically with the electoral calendar. Neither is assumed away, so this is a limitation rather than a resolved problem.

Detection. A first text-reuse pass identifies syndicated and duplicate articles, the problem computational work on reprinted texts in Chronicling America's antebellum newspapers addresses (Smith et al., 2015), using shingling with MinHash and locality-sensitive hashing (LSH) banding. Full pairwise comparison is infeasible at this scale: the earlier version's roughly 8.5 million article records imply on the order of 3.6×10^{13} candidate pairs; MinHash approximates the Jaccard similarity of two documents' word-shingle sets from a compact fixed-size signature, so similarity can be estimated in time proportional to signature length rather than document length (Broder, 1997). Each article (filtered to articles that pass the quality filter and have at least 50 words) is reduced to 5-word shingles and a 128-permutation MinHash signature; `MinHashLSH` generates candidate near-duplicate pairs at a loose threshold (0.4), each candidate is re-checked against the exact Jaccard similarity of its signature, and pairs at or above 0.6 are kept as graph edges. Edges are grouped into clusters via union-find, and a cluster counts as genuinely syndicated only if it spans at least 2 distinct newspapers, so that a paper re-running its own earlier item is not treated as syndication.

Status. On a development subsample of the earlier version (1,169,035 articles) this MinHash pass found

23,346 clusters, of which 27 contained duplicates spanning newspapers, covering 62 articles—about 5 in every hundred thousand. The full-corpus detection pass and a canonical-text merge collapsing verified clusters have been designed but not applied; the panels retain every reprint as a separate article. The detector found few cross-paper clusters at that threshold, but its recall on uncorrected OCR text is unknown: two copies of one item set by different compositors and read by the OCR engine with different errors can fall below a Jaccard similarity of 0.6. The low count is therefore a lower bound on reprinting, not an estimate of it. A second pass to address that weakness is specified in the replication package. It keeps the clustering, verification, and cross-paper rule above but replaces the word shingles with embeddings from a neural bi-encoder trained to detect duplicate news text, linking each article to those among its 50 nearest neighbours with a cosine similarity of at least 0.9. It has not been run, so the paper reports no estimate from it, and syndication remains the limitation stated above.

A.4 Keyword Dictionaries

Construction. Every measure in the paper comes from one keyword dictionary with 23 topics and 4,581 entries. Of the topics, 11 are political issues: slavery and abolition, the campaign (elections), race and equality, the tariff and land policy, finance, immigration and nativism, the draft, war, temperance, labor, and Copperheads and states' rights. The other 12 are the ordinary content of a newspaper: shipping and transport, agriculture, religion, obituaries, local crime, fiction, commercial news, weather, government business, legal notices, literary matter, and advertisements. Each entry sits in a tier. Strong entries name the subject (1,808 across topics), formula entries are the stock phrases of a kind of item (628), and medium entries belong to the topic but also appear elsewhere (1,712). A further 433 common words are listed to document the topic's vocabulary but are not counted. The full dictionary, with every entry, its tier, its spelling variants, and its accepted misreadings, is in the replication package.

Matching. Text is lowercased, apostrophes are dropped, digit runs are read as one number symbol, and each entry is matched as a whole word or an exact sequence of words. Spelling variants of the period (*Negroe*, *despatch*) and inflected forms count as the entry they belong to. Misreadings by the optical-character-recognition engine are listed explicitly: for each entry word of eight letters or more (six inside a phrase), common confusions such as *rn* for *m*, *cl* for *d*, *l* for *i*, and *c* for *e* were generated, and a misreading was kept only if it appears at least 5 times in a 96,000-article sample of the corpus, is at most a quarter as common as the correct spelling, and is neither an English word nor another entry. A match inside a longer match of the same topic is not counted twice. An article is a hit for a topic if it contains at least two different counted entries for that topic; a misreading counts as the entry it garbles.

An earlier version. An earlier version of the measure used a smaller set of regular-expression dictionaries, 19 for issues and 8 for content, in which a single matching term made an article a hit. Re-estimated on those dictionaries, the 1860 break in slavery coverage is about 27 percent larger as a share of its baseline. On the articles that both versions flag, the break is the same size as on the earlier version.

Aggregation. For each newspaper j and calendar month t, absolute salience $s_{slav,jt}$ is the share of that month's articles that are slavery hits, and $s_{pol,jt}$ is the share that are campaign hits. Two variants are carried through: a word-weighted share, which counts the words in hit articles over all words published so that the measure reflects column inches rather than article count, and the ratio $s_{slav,jt}/(s_{slav,jt} + s_{pol,jt})$, which requires no baseline normalization. The within-paper relative measure R_{jt} is built from the two article shares (Section A.6).

Validation. The data section of the paper reports three checks: accuracy against 5,000 articles drawn at random and labelled blind; whether each topic peaks at dated events (the war topic at Antietam and Gettysburg, the draft topic near the New York draft riots, the slavery topic not at the Emancipation Proclamations); and the campaign topic's movement on the electoral calendar. An earlier attempt to build the salience measure from sentence-transformer embeddings and BERTopic clusters carried essentially no article-level topical signal on this OCR'd corpus and was abandoned; that diagnostic is reported in Section D.

A.5 Page- and Article-Level Co-occurrence

The composed constructs (`race × crime`, `slave × abolition`, and the rest), built from the nine term families of the earlier regular-expression dictionary, count an item when terms from both families appear in it. The families are race (terms for people of color), slavery (the institution and its defenders), abolition (antislavery politics), crime (violence and wrongdoing), insurrection (servile revolt), equality (social and political equality, amalgamation, miscegenation), labor (white labor and economic competition), influx (migration and "Africanization"), and citizenship (rights and the franchise). Each family's terms are softened against two common optical-character-recognition confusions, among `i`, `l`, `1`, and `!`, and `w` read as `vv`. They are not used in the paper's regressions; this comparison is a standalone measurement exercise. Table A.1 applies the rule at two levels over the full corpus: within a single article, and within a newspaper page, as a page-level corpus would have to. Page-level counts exceed article-level counts by a factor of 22.4 to 95.9 across the 14 constructs, because a page holds many unrelated items and two vocabularies land on the same page far more often than in the same article.

Table A.1: Co-occurrence of Dictionary Families at the Article and Page Level

Construct	Articles containing both	Articles on a matching page	Ratio
race × crime	37,802	2,336,712	61.8
race × labor	12,636	746,625	59.1
race × equality	11,466	256,971	22.4
race × influx	5,433	520,774	95.9
race × citizen	6,874	381,980	55.6
race × insurrection	10,149	813,264	80.1
slave × crime	44,637	2,272,563	50.9
slave × abolition	63,196	1,596,702	25.3
slave × labor	16,839	771,267	45.8
slave × insurrection	13,919	832,940	59.8
slave × citizen	7,690	398,417	51.8
abolition × crime	20,284	1,418,897	70.0
abolition × insurrection	7,345	560,985	76.4
crime × insurrection	12,094	857,663	70.9

Note: Each row counts articles in the full corpus. Articles containing both: articles in which at least one term from each of the two dictionary families appears. Articles on a matching page: articles on a newspaper page (identified by title, date, edition, and page number) where terms from both families appear anywhere on the page, in the same article or not. The ratio is the second count over the first. The labor family is the least precise of the nine vocabularies, so constructs that include it carry more incidental matches. Source: author's dictionary scan of the AmericanStories corpus, 1860–1865.

A.6 Panel Construction

Election-relative indexing. The panel has one row per newspaper and calendar month, so no month is counted twice. Each month is attributed, separately, to the nearest congressional and the nearest presidential election in the paper's state, from a hand-compiled table of election dates for the 1860, 1862, and 1864 contests; ties go to the upcoming election. Event time is the calendar month minus the month of the attributed election, so negative values fall before the election and zero is the month containing election day. A month more than twelve months from any congressional election (twenty-four for the quadrennial presidential elections) has no attributed election of that type. Presidential elections are held on a uniform November date; congressional election dates differ across states, and in the 1860 cycle they run from December 1859 to October 1861, which is what allows calendar-month fixed effects in the event study to be identified separately from election proximity.

Within-paper normalization. The relative measure asks whether slavery grew *relative to campaign coverage* rather than in absolute terms, comparing the two series within each newspaper against its own non-election behavior. Let $\mathcal{B}_j$ denote newspaper j's *baseline* months, those at least seven months from both its nearest congressional and its nearest presidential election, and let $q \in \{slav, pol\}$ index the slavery and campaign dictionary shares. The within-paper z-score of topic q in newspaper j in month t is

$$z_{q,jt} = \frac{s_{q,jt} - \bar{s}_{q,j}}{\tilde{\sigma}_{q,j}}, \qquad \bar{s}_{q,j} = \frac{1}{|\mathcal{B}_j|} \sum_{t \in \mathcal{B}_j} s_{q,jt}, \qquad \tilde{\sigma}_{q,j} = \max\{\sigma_{q,j}, \underline{\sigma}_q\}, \tag{A.1}$$

where $\sigma_{q,j}$ is the standard deviation of $s_{q,jt}$ over $t \in \mathcal{B}_j$. Because a small paper can have zero slavery or campaign articles across its entire baseline window, the floor $\underline{\sigma}_q$ is the tenth percentile of the across-paper distribution of strictly positive baseline standard deviations of s_q. Each series is thus centered and scaled on the same paper's own non-election behavior, so the normalization absorbs permanent differences in level and volatility across newspapers. The normalized relative salience is the difference of the two,

$$R_{jt} = z_{slav,jt} - z_{pol,jt}. \tag{A.2}$$

A rise in R_{jt} means slavery coverage climbed faster than the paper's campaign coverage, relative to that paper's own normal: editors reprioritized rather than adding slavery on top of a larger campaign section. R_{jt} is left undefined for papers with fewer than three baseline months. The newspaper panel is 7,929 newspaper-months from 242 newspapers in 26 states, as described in the data section of the paper.

County panel. The county-level panel joins the newspaper panel to the county returns of Clubb et al. (1987) and 1860 census aggregates from NHGIS (Schroeder et al., 2025) via a Chronicling America LCCN-to-county crosswalk, matching on state as well as county name. Each newspaper is assigned to its city's county; a county's pre-election coverage is the mean of s_{slav} over the six months before the election across the newspapers indexed to it. Republican vote share is the Republican or National Union share of the county's presidential vote in 1860 and 1864, and the Republican share of the two-party congressional vote in 1860, 1862, and 1864, with uncontested races removed. Controls are 1860 county totals of population, enslaved and free Black population, foreign-born white population, farm cash value, and real estate value, each interacted with election dummies. The estimation sample with census controls is 313 county-elections in 105 counties.

Table B.1: Corpus Composition by Sample Restriction

Sample	Total articles	Retained	Filtered	Filter rate	Newspapers (LCCNs)
Full corpus	8,486,860	6,194,997	2,291,863	27.0%	408
Union states	6,084,017	4,497,209	1,586,808	26.1%	241

Note: The full corpus includes all article records in the AmericanStories dataset for 1860–1865. The Union-states sample retains only records whose newspaper LCCN matches a title in one of the Union states listed in the next subsection. Clean articles are those passing the OCR quality thresholds and minimum-length filter described in the text. The filter rate is the share of Union-state articles classified as low quality and excluded.

B Text Corpus and OCR Quality Filtering

This appendix describes the upstream corpus construction that the main data section treats as given. It covers the raw digitized text, the OCR quality filters applied to select clean articles, the Union-state restriction, and diagnostic evidence that the filter does not systematically distort the sample in ways that would threaten identification. The counts and filter diagnostics below were computed on an earlier version of the corpus. The estimates in the paper use a later version of the corpus, 6,309,545 articles from 406 titles, which keeps the deterministic cleaning layers, drops the model-based spelling correction, and applies its own quality filter; the diagnostics have not been recomputed on that version.

B.1 The Digitized Corpus

The newspaper text comes from Chronicling America, the Library of Congress's digitized historical newspaper archive. I access it through the AmericanStories dataset (Dell et al., 2023), which restructures the raw Chronicling America page scans into article-level records with associated OCR quality metrics. Each record contains the extracted article text, the newspaper LCCN, publication date, page number, and two post-correction quality scores: `ocr_quality_proc`, a character-recognition score on a 0–100 scale, and `nonchar_proc`, the share of non-alphabetic characters in the processed text (expressed as a percentage).

The full corpus spans 1860–1865 and contains 8,486,860 articles before any filtering. Table B.1 reports the composition of the corpus at each stage of the sample restriction.

B.2 Union-State Restriction

The corpus diagnostics in this appendix are restricted to newspapers published in states that remained in the Union throughout the study period. The sample includes 19 free states: California, Connecticut, Illinois, Indiana, Iowa, Kansas, Maine, Massachusetts, Michigan, Minnesota, New Hampshire, New Jersey, New York, Ohio, Oregon, Pennsylvania, Rhode Island, Vermont, and Wisconsin. It also includes four border states that

did not secede: Delaware, Kentucky, Maryland, and Missouri. West Virginia is included from its admission to the Union in June 1863 onward. The analysis panel does not apply this list: it keeps any title it can attribute to a congressional election within twelve months (Section A.1), so it also contains Tennessee and West Virginia papers attributed to their 1861 congressional votes and, in 1864, a few papers from occupied Virginia and South Carolina.

Newspapers published in Confederate states, U.S. territories, and the District of Columbia, as well as articles whose newspaper LCCN could not be matched to a state, are dropped. This removes 2,402,843 articles (28.3 percent of the full corpus), yielding a Union-states corpus of 6,084,017 articles.

The restriction applies to the diagnostics in this appendix only. The analysis panel uses a different rule, attribution to a congressional election within twelve months (Section A.1), and so includes some papers from seceded states: Tennessee's in the 1860 cycle and 1864, and a few from occupied Virginia and South Carolina in 1864. Border-state papers are retained because those states held contested congressional elections throughout the war.

B.3 OCR Quality Filtering

Not all digitized articles are suitable for topic measurement. Chronicling America's digitization program spans newspapers with widely varying print quality, paper condition, and column density. OCR errors accumulate in dense tabular content (commodity prices, shipping notices, legal advertisements), in small-type filler, and in damaged pages. Articles with high OCR error rates produce token streams dominated by character noise rather than recognizable words, which degrades dictionary matching. Two filters are applied in sequence.

Filter 1: OCR quality thresholds. Retain articles with `ocr_quality_proc` ≥ 95 and `nonchar_proc` ≤ 10. These thresholds are calibrated to retain articles with near-clean character recognition while discarding price tables and boilerplate whose high non-alphabetic character share (dollar signs, fractions, column separators) carries no issue signal and adds noise to every text measure.

Filter 2: Minimum length. Retain articles with at least 20 words after OCR correction, a fixed minimum-word-count threshold applied during corpus preparation. Very short articles—single-sentence fillers, dateline notices, classified advertisements—carry no reliable topical signal and disproportionately pass the OCR quality thresholds despite lacking substantive content.

Together, the two filters classify 26.1 percent of Union-state articles as low quality, retaining 4,497,209 clean Union-state articles in that earlier version. (In that version, scoring was run on the full clean corpus of 6,194,997 articles across all indexed states, with the Union-state and non-Confederate restrictions applied afterward, when the panel was assembled; the corpus used in the paper is scored the same way.) The OCR metrics

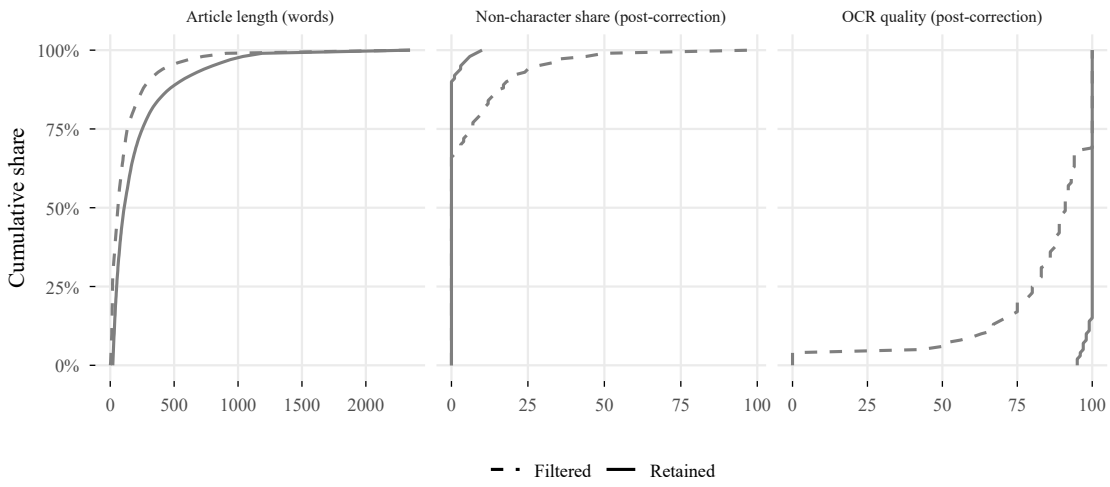


Figure B.1: Empirical CDFs of article length, non-alphabetic character share, and OCR quality, from left to right, for retained and filtered articles. Retained articles (solid lines) pass both quality filters; filtered articles (dashed lines) fail at least one. The filter thresholds are 20 words, a non-alphabetic share of 10 percent, and an OCR quality score of 95; they are not drawn in the figure. Character measures are computed after OCR correction. Source: AmericanStories dataset (1860–1865), Union states, earlier version of the corpus.

confirm that the filtered articles are qualitatively different: clean articles have a mean `ocr_quality_proc` of 99.60 versus 84.30 for low-quality articles, a mean `nonchar_proc` of 0.38 versus 5.29, and a mean length of 208 words versus 120 words. Low-quality articles are systematically shorter and more garbled—the profile of commodity listings, advertisements, and column filler rather than editorial and political content.

Figure B.1 plots the empirical CDFs of the three quality metrics separately for retained and filtered articles. Retained articles sit at the clean end of the two character measures by construction, with OCR quality above 95 and almost no non-alphabetic characters. The filtered group is mixed. About 68 percent of filtered articles have OCR quality below the threshold, but the other 32 percent score near 100: these failed the length filter, not the character filters, and are short, cleanly recognized items such as datelines and notices. Article length overlaps heavily between the two groups, since the length filter removes only the shortest items. The filter therefore removes two different kinds of material, garbled text and very short text, rather than a single well-defined segment.

B.4 Filter Bias Diagnostics

A filter that removes 26 percent of articles is consequential only if the removed articles are systematically different from retained articles in ways that correlate with the identification strategy. The relevant threat is not that the filter removes articles in general, but that it removes articles differentially by topic *within the same newspaper over time*—the variation that identifies the event-study estimates. Permanent between-newspaper

Table B.2: Chi-Square Tests of OCR Filter Homogeneity

Dimension	Groups	χ^2	d.f.	p-value
Year	6	1629.7	5	<0.001
State	23	81418.8	22	<0.001
Partisan affiliation	8	35569.5	7	<0.001
Page number	16	49878.2	15	<0.001
Newspaper (LCCN)	235	218813.3	234	<0.001

Note: Each row reports a chi-square test of the null hypothesis that the OCR filter rate is constant across categories of the listed dimension. The test is conducted on the Union-states corpus (6,084,017 articles). All tests are statistically significant at conventional levels given the sample size of approximately six million observations; the economically relevant quantity is the range of filter rates across categories, reported in the text, not the p-value.

differences in filter rates are absorbed by the newspaper fixed effects, so a newspaper that loses 30 percent of its articles to filtering still contributes valid coverage shares as long as the filter does not change systematically around elections.

Table B.2 reports chi-square tests of filter-rate homogeneity across five dimensions: year, state, partisan affiliation, page number, and newspaper LCCN.

All five tests reject homogeneity at conventional levels. With 6,084,017 articles, statistical significance is uninformative—even a one-tenth of one percentage point difference in filter rates would produce a highly significant chi-square statistic. The economically relevant quantity is the range across categories, which governs the potential distortion to the treatment measure.

Temporal variation. Filter rates vary from 25.0 percent in 1860 to 27.0 percent in 1864—a range of 2.1 percentage points over 5 years. There is no discontinuity around the issuance of the Emancipation Proclamation in September 1862 or any other political event in the study period. A drift of that size over 5 years is unlikely to generate spurious variation in coverage shares that coincides with the pre-election windows, which are measured at the monthly rather than annual level.

State variation. Filter rates range from 16.8 percent (Rhode Island) to 36.3 percent (New York). These differences are absorbed by newspaper fixed effects: a New York paper that systematically loses 36 percent of its articles to filtering still contributes valid coverage shares so long as the within-paper filter rate does not vary with the election calendar. Figure B.2 plots state-level filter rates.

Partisan variation. Table B.3 and Figure B.3 report filter rates by partisan classification. The range across the four main party groups (Republican and Peace, War, and Confederate Democratic) is 9.5 percentage points: War Democratic papers have the lowest filter rate (20.0 percent) and Confederate Democratic papers

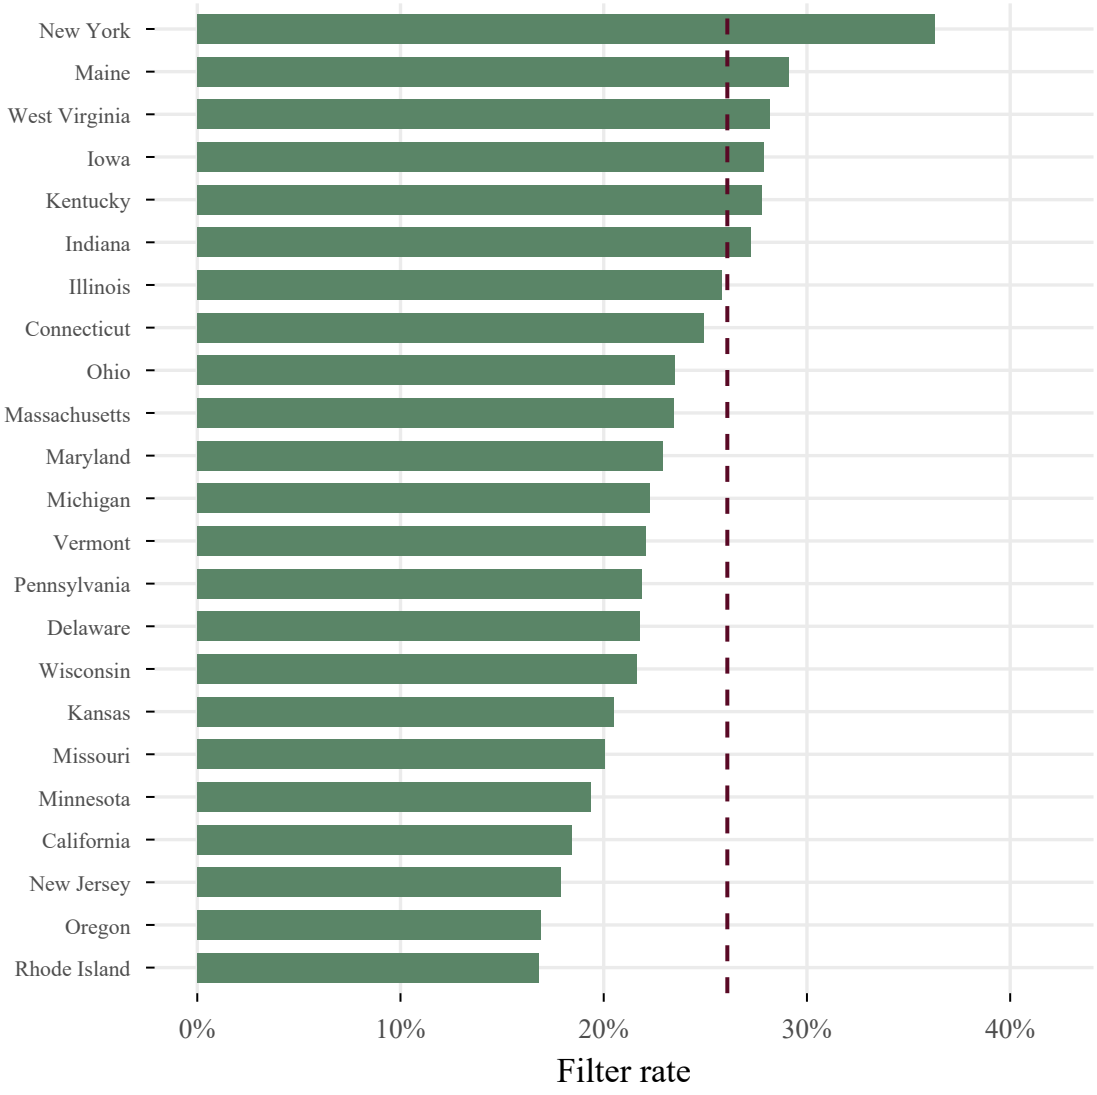


Figure B.2: OCR filter rate by Union state (share of articles classified as low quality). Rates reflect all articles in the Union-states corpus, 1860–1865. The overall filter rate of 26.1 percent is marked by a vertical dashed line. Source: AmericanStories dataset.

Table B.3: OCR Filter Rate by Partisan Classification

Partisan classification	Articles	Filter rate
Conditional Confederate	10,119	57.5%
Independent	136,953	41.6%
Confederate Democratic	665,222	29.5%
Republican	3,616,497	25.9%
Peace Democratic	1,234,945	24.8%
War Democratic	231,813	20.0%
Union Army	3,354	19.9%
Not recorded	185,114	19.9%

Note: Filter rate is the share of articles classified as low quality within each partisan group. Groups are sorted by filter rate, descending. Partisan classification here is the metadata's summary alignment label, assigned at the LCCN level; it is an older, coarser classification than the election-by-election ticket coding used in the paper. The overall rate is 26.1 percent. Source: AmericanStories dataset (1860–1865), Union states.

have the highest (29.5 percent). Confederate Democratic is a label in the metadata's summary alignment field. Among the three largest groups by article count (Republican, Peace Democratic, and Confederate Democratic) the range is narrower, 4.6 points. More importantly, within-city comparisons hold city and month fixed, so partisan differences in filter rates are absorbed so long as Democratic and Republican papers in the same city face similar OCR infrastructure—which they do, given that state-level digitization quality, not editorial content, drives the between-newspaper variation documented above.

Page variation. Filter rates are lowest on page 2 (21.6 percent) and highest on page 11 (42.6 percent). This gradient is the expected pattern: front and second pages carry news and editorial content printed in standard column widths, while inner and back pages accumulate advertisements, shipping tables, legal notices, and compressed filler set in small type. The filter is doing exactly what it should—removing the high-nonchar content concentrated on interior pages.

Newspaper-level variation. The chi-square test across the 235 Union-state LCCNs with enough articles to test produces a test statistic of 218,813 with 234 degrees of freedom. Per-newspaper filter rates vary widely, but this variation is mechanical: newspapers with older, less-preserved print runs have worse OCR. This is absorbed by the newspaper fixed effect by construction. Figure B.4 plots the distribution of per-newspaper filter rates.

Temporal variation within newspapers. The identification-relevant threat is not cross-sectional heterogeneity in filter rates but temporal variation within the same newspaper correlated with the election calendar. Figure B.5 plots the corpus-wide filter rate by year. The absence of a systematic pre-election spike

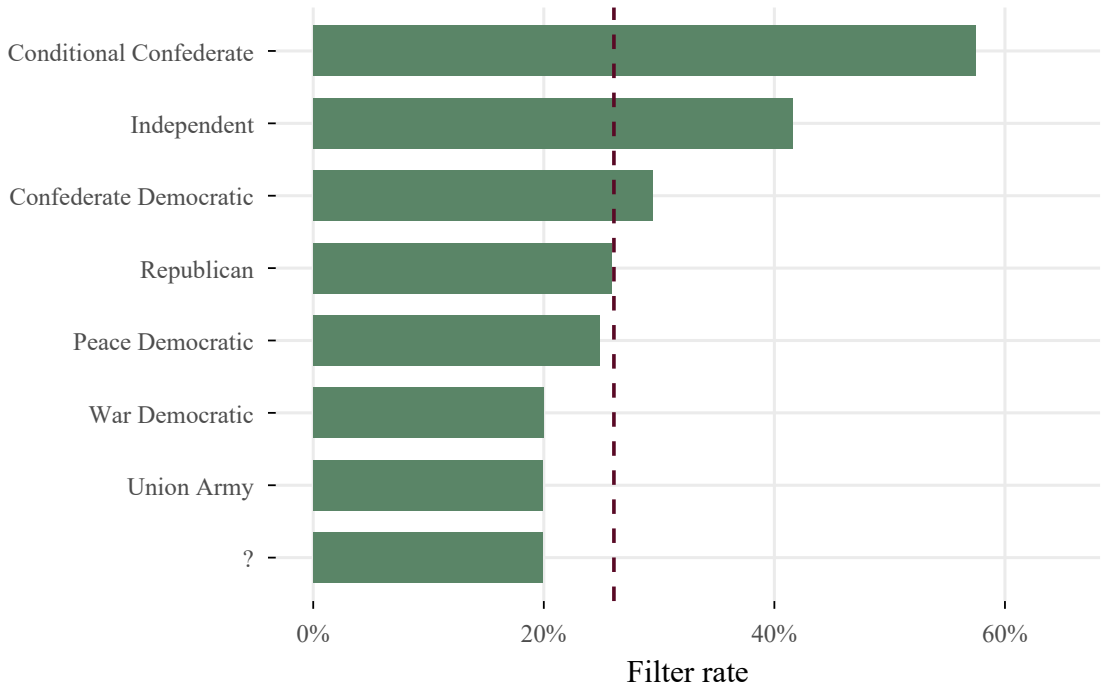


Figure B.3: OCR filter rate by partisan classification. The overall filter rate of 26.1 percent is marked by a vertical dashed line. Partisan classification is the metadata's summary alignment label, assigned at the LCCN level. Source: AmericanStories dataset (1860–1865), Union states.

in any of the three election years (1860, 1862, 1864) suggests the filter does not selectively remove articles in the months immediately before elections.

Taken together, the diagnostics suggest that OCR filtering introduces cross-sectional heterogeneity in sample composition—between states, between partisan groups, between interior and exterior pages—but not the within-newspaper, election-timed variation that would generate spurious signals in the coverage shares. The cross-sectional differences are absorbed by the newspaper fixed effects in the estimation equations.

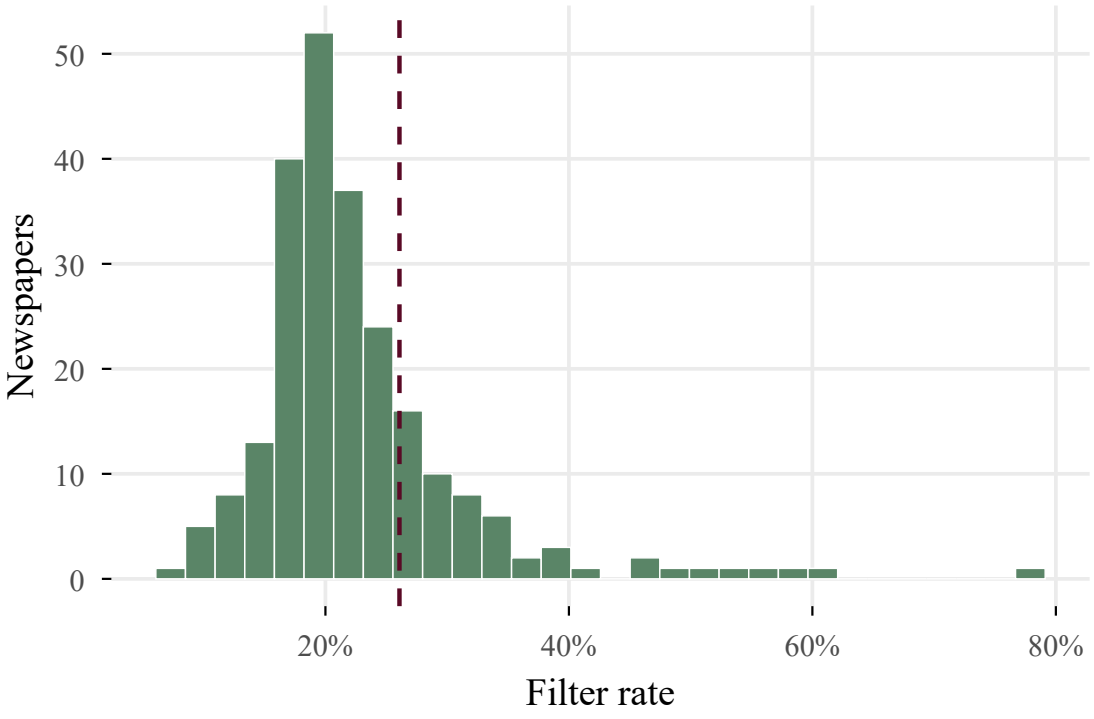


Figure B.4: Distribution of OCR filter rates across newspaper LCCNs. Each observation is one newspaper title (LCCN). The vertical dashed line marks the corpus-wide filter rate of 26.1 percent. Newspapers with few total articles (fewer than 100 in the full 1860–1865 period) are included in this distribution but contribute low weight to the analysis. Source: AmericanStories dataset (1860–1865), Union states.



Figure B.5: OCR filter rate by year, 1860–1865. The overall filter rate of 26.1 percent is marked by a dashed horizontal line. Election years are 1860, 1862, and 1864. The absence of a systematic spike in election years suggests the filter does not selectively remove articles in the pre-election period. Source: AmericanStories dataset, Union states.

Table C.1: Event-study coefficients, 1860 congressional cycle, under two reference choices

	Dictionary article share		Dictionary word share	
Months to election	Both sides	Pre-election only	Both sides	Pre-election only
−6	+0.0087 (0.0024)	+0.0045 (0.0041)	+0.0232 (0.0053)	+0.0144 (0.0095)
−5	+0.0085 (0.0022)	+0.0012 (0.0062)	+0.0244 (0.0059)	+0.0088 (0.0210)
−4	+0.0079 (0.0027)	−0.0029 (0.0074)	+0.0213 (0.0060)	−0.0026 (0.0260)
−3	+0.0104 (0.0034)	−0.0031 (0.0093)	+0.0265 (0.0077)	−0.0040 (0.0318)
−2	+0.0173 (0.0043)	+0.0005 (0.0098)	+0.0453 (0.0107)	+0.0073 (0.0359)
−1	+0.0146 (0.0046)	−0.0047 (0.0113)	+0.0299 (0.0101)	−0.0143 (0.0403)
0	+0.0138 (0.0044)	−0.0058 (0.0104)	+0.0268 (0.0109)	−0.0178 (0.0424)
+1	+0.0034 (0.0019)	−0.0183 (0.0118)	+0.0021 (0.0072)	−0.0480 (0.0491)
+2	+0.0078 (0.0029)	−0.0156 (0.0145)	+0.0113 (0.0113)	−0.0424 (0.0575)
+3	+0.0111 (0.0031)	−0.0135 (0.0149)	+0.0262 (0.0115)	−0.0294 (0.0597)
+4	+0.0072 (0.0019)	−0.0192 (0.0152)	+0.0158 (0.0079)	−0.0445 (0.0586)
+5	+0.0018 (0.0021)	−0.0267 (0.0168)	+0.0104 (0.0077)	−0.0532 (0.0637)
+6	−0.0017 (0.0021)	−0.0333 (0.0169)	+0.0004 (0.0057)	−0.0695 (0.0661)

Note: Event-study coefficients from the paper's Design 1, 1860 congressional cycle, all papers, with newspaper and calendar-month fixed effects. "Both sides" uses as the reference all months at least seven months from the paper's congressional election; "Pre-election only" uses only such months before the election, which the corpus, beginning in January 1860, limits to the first one to four months of 1860 for the states that voted in the autumn of 1860 and which fall in the 1860 presidential campaign for the states that voted in 1861. Standard errors clustered two-way on newspaper and election date in parentheses; the 1860 cycle has 11 election-date clusters, 6 of them votes held in 1861, so these analytic standard errors are approximate. Source: AmericanStories corpus, 1860–1865; author's calculations.

C Additional Results

C.1 Full Event-Study Coefficients

Table C.1 reports the monthly event-study coefficients for all 1860-cycle papers, on the dictionary article and word shares, under two references: the far months on both sides of the vote ($|k| \geq 7$), the paper's main reference, and the far pre-election months alone ($k \leq -7$). Against both sides, coverage sits above the reference in every month before the vote. Against the far pre-election months alone, the months before the vote are close to zero and the months after it fall below. The two columns differ by close to a straight line in event time, the drift that a one-sided reference cannot pin down within a single cycle (the paper's Design 1); the same problem makes the paper's grouped one-sided estimate sensitive to how the months near the vote are grouped. The shape across the vote is similar in both columns, and the break is about 21 percent smaller against the far pre-election months alone.

C.2 Grouped Pre-Election Windows

The grouped-window specification gives the same shape for the 1860 cycle. Against the far months on both sides, the dictionary article share is +0.0090 four to six months before the vote, +0.0142 one to three months before, +0.0144 in the election month, and +0.0066 and +0.0016 one to three and four to six months after.

Table C.2: Pre-election slavery coverage by cycle and party, congressional elections

Cycle	Party	Dictionary (articles)	Dictionary (words)	Date clusters
1860	Republican	+0.0117 [0.203]	+0.0280 [0.219]	7
1860	Democratic	+0.0102	+0.0165	5
1862	Republican	−0.0026 [0.461]	−0.0015 [0.801]	9
1862	Democratic	+0.0211	+0.0544	3
1864	Republican	−0.0036 [0.051]	−0.0075 [0.039]	9
1864	Democratic	+0.0086	+0.0276	4

Note: Coefficient on the one-to-three-month pre-election window, by cycle and party, against each paper's months at least seven months from its congressional election, with newspaper and calendar-month fixed effects. Brackets: exact wild-cluster-bootstrap p-value on election-date clusters. Cells with five or fewer election-date clusters report the point estimate only, because the bootstrap cannot resolve p-values below $2/2^G$ with G clusters. Constitutional Union papers are not assigned to either party. Source: AmericanStories corpus, 1860–1865; author's calculations.

The word share runs +0.0234, +0.0336, +0.0300, +0.0121, and +0.0057 over the same windows. Coverage is highest in the three months before the vote and the election month itself, about 37 percent above its far-month level of 3.9 percent of articles, and drops once the vote has passed. A Poisson model of the count of slavery articles, with $\log(\text{words}_{jt})$ as a control, puts the 1860 pre-election rise at an incidence-rate ratio of 1.23, a 23 percent rise in the count, significant on the analytic two-way standard error. Relative to 1860, the response is smaller in 1862 and 1864 (ratios 0.91 and 0.93, or 9 and 7 percent lower), though neither difference is distinguishable from zero.

C.3 Full Party-by-Cycle Table

Table C.2 reports every cycle-by-party cell behind the party comparison in the paper. Cells with three to five election dates—every Democratic cell—have too few dates for informative inference under the exact wild cluster bootstrap, whose smallest attainable p-value at those counts is between 0.0625 and 0.25; their point estimates are reported without a p-value. Among the Republican cells, only 1860 shows a pre-election rise.

Reprioritization in 1860. The paper asks whether 1860 Republican papers reprioritized their agenda toward slavery or added slavery on top of a larger political section, and finds that the evidence leans toward the second. Their campaign coverage rose significantly before the vote, by 3.8 percentage points (+0.038), about three times as much as their slavery coverage (1.2 points). The relative measure R_{jt} of Equation A.2 and the ratio measure are both indistinguishable from zero for these papers (Section D.3).

C.4 Papers with a Rival in Town

The paper reports a pre-specified test of whether the 1860 break was larger for papers with a competitor in the same town, and finds it uninformative. This section gives the definition and the inference behind that

verdict. A paper has a rival if another newspaper firm of either party in the same city, keyed on city and state, appears in the 1860-cycle sample. Rivals are counted by firm, with title chains and editions merged. That rule was adopted after the first estimates, to correct a coding error that counted a paper's own successor title as its rival. The paper's break specification is estimated with every event-time indicator interacted with the rival flag.

The minimum detectable difference, 0.0159 on the article share, uses the analytic standard error, which is smaller than the bootstrap implies, so the difference the bootstrap could detect is larger than 0.0159. On the word share the difference has the same sign as on the article share and is also indistinguishable from zero (-0.0147). Across the six-issue family of the paper's multiple-testing check and the campaign topic, no difference survives a Benjamini–Hochberg correction. The partisan version of the prediction, a rival of the opposite party, cannot be tested: the 14 papers with such a rival lie in only 2 election-date clusters.

The papers with a rival span only 5 of the 11 election dates. Randomization inference over the election calendar tests the sharper null of no election-timing effect in either group, not equal breaks, so it is not used to judge the difference. Competition is not randomly assigned, so the comparison is between kinds of market, not an effect of a rival. And papers whose rivals the archive never digitized are counted as alone, which pulls the difference toward zero.

C.5 What the Campaign Displaced

The paper reports that output was flat before elections in every cycle. This section reports what moved inside that fixed output. It identifies a pattern, not a cause: output is itself a publisher's choice, and nothing here separates a publisher who could not expand from one who judged expansion not worth its cost. The paper's conclusion discusses what an exogenous shifter of newsprint capacity would add.

Papers had to make room for the campaign somewhere. The design asks where, in three steps, using the grouped-window version of the paper's event study, which collapses the months into windows of one to three and four to six months before and after the election against the far months on both sides; pre_{jt} marks the window one to three months before the vote.

First, does output expand before an election? I estimate the grouped-window specification with $\log(n_{jt})$, $\log(\text{words}_{jt})$, and mean page count as outcomes. A coefficient near zero means the paper did not grow, so any topic's gain must be another's loss. It does not establish that the paper *could* not have grown; the data section of the paper separates the measured quantity from the capacity interpretation.

Second, what moves? I re-estimate the grouped-window specification with each issue dictionary as the outcome—slavery, the tariff and land policy, finance, immigration, labor, temperance, the draft, war news,

campaign coverage, and others—and with the dictionary's topics for the ordinary non-political content of a newspaper: commercial news, agriculture, local crime, weather, obituaries, legal notices, advertisements, and the like. If the campaign crowded out coverage in proportion to how much space a topic already held, regressing each non-campaign topic's pre-election change on its far-month share should give a negative slope, and residuals from that line identify topics that beat the proportional benchmark. Because this benchmark uses pre-election levels, it inherits the dependence of the pre-election level on the reference months.

Third, is slavery's gain specific to 1860? I estimate the interaction

$$y_{jt} = \beta_{1860}\,\mathrm{pre}_{jt} + \sum_{c\in\{1862,1864\}} \phi_c\,(\mathrm{pre}_{jt} \times \mathbf{1}[\mathrm{cycle} = c]) + \alpha_j + \lambda_t + \varepsilon_{jt} \tag{C.1}$$

with 1860 the omitted cycle, so that β_{1860} is the pre-election response in 1860 and ϕ_c is the difference between cycle c and 1860. Estimating the contrast directly, rather than comparing subsample point estimates, is what allows it to carry a standard error.

Papers did not grow before elections. In 1860, article counts in the three months before the vote were 0.050 log points lower than in the far months (about 4.8 percent), word counts 0.011 lower (about 1.1 percent), and pages per issue flat at about 3.9; none of the three is distinguishable from zero. The same holds in 1862 and 1864. In 1862 article counts are 0.014 log points higher and word counts 0.002 higher, both indistinguishable from zero. In 1864 article counts are 0.033 higher and word counts 0.031 higher, again indistinguishable from zero, and so is the direct 1864-versus-1860 contrast (+0.049 on articles, +0.043 on words). Every campaign surge had to come out of something.

The estimates in the rest of this section are pre-election levels and share their dependence on the reference months, which the results section of the paper documents. In 1860 the campaign surge was large. Campaign coverage—the candidate, ticket, and returns material that dominates late-campaign newspapers—rose significantly, by +0.043 on the dictionary measure in the three months before the vote, far more than anything else. That is 4.3 percentage points on a far-month level of 7.0 percent, a rise of 62 percent. On the article share, slavery and the Copperhead and states'-rights topic, whose vocabulary in 1860 is that of secession and disunion, were the two largest *issue* gainers, nearly tied and both significant (+0.014 and +0.013); on the word share the Copperhead topic leads (+0.038 against +0.034). Against their far-month levels the two gains differ more: slavery's is 37 percent of a 3.9 percent baseline, the Copperhead topic's 19 percent of a 6.9 percent baseline. The tariff and land-policy topic rose significantly, by +0.005. That is only 0.5 of a percentage point, but the topic's far-month level was 0.5 percent of articles, so the rise is 99 percent of it. Labor rose +0.003, an estimate indistinguishable from zero. Figure C.1 shows the full ranking.

The validation subsection of the paper's data section reports how often these topics' hits are on topic.

The war topic is a separate case before Fort Sumter: with no war to report, its military vocabulary can only catch incidental uses, and in 1860 its pre-election movement (−0.024, or 17 percent of a far-month level of 14.0 percent) is not distinguishable from zero. The war row is the weakest in the comparison for a second reason. Calendar-month fixed effects absorb the national course of the war but not where it was fought, and the cohorts that voted after Fort Sumter did so in wartime while the 1860 cohorts did not.

What else the campaign displaced is unclear. No category of a paper's ordinary content fell significantly before the vote. Advertisements fell −0.015, 7 percent of their far-month level of 19.7 percent of items, and agricultural, crime, shipping, and weather items fell by smaller amounts; none of these falls is distinguishable from zero. Government business rose significantly (+0.015, 15 percent of its far-month level). The campaign displaced something, since the sheet did not grow, but no single category of routine content carries the displacement.

Slavery also beats the benchmark a proportional squeeze would predict. If the campaign took space from other topics in proportion to how much each already held, larger topics would give up more. Across the non-campaign issues in 1860 they do: regressing each issue's pre-election change on its far-month share gives a slope of −0.137 (SE 0.070, $R^2 = 0.33$). Two issues sit well above that line, the Copperhead and states'-rights topic (residual +0.016) and slavery (+0.013); no other issue's residual exceeds +0.002. The benchmark predicted roughly no change for slavery (+0.001) and a decline for the Copperhead topic (−0.003); both rose far above it.

The pre-election rise in slavery coverage is specific to 1860. In Equation C.1 it is +0.010 in 1860 and smaller in 1862 by 0.009, 98 percent of the 1860 rise, and in 1864 by 0.007, 71 percent of it. On the word share the contrasts are 0.020 and 0.016. The 1862 contrast is significant on both measures and the 1864 contrast marginal on both. Measured in words, campaign coverage gains on both measures in 1860, +0.058 of words and +0.043 of articles, both significant, while the war topic's pre-election decline is indistinguishable from zero on either.

C.6 County Vote Shares

The county design asks whether counties whose newspapers ran more pre-election slavery-related coverage voted differently. The unit is the county × election, the outcome is Republican vote share, and the treatment C_{ge} is slavery-related coverage averaged over the six months before the election across the newspapers indexed to the county:

$$\text{RepShare}_{ge} = \psi\, C_{ge} + \mathbf{X}_g'\eta_e + \omega_g + \tau_e + \varepsilon_{ge}, \tag{C.2}$$

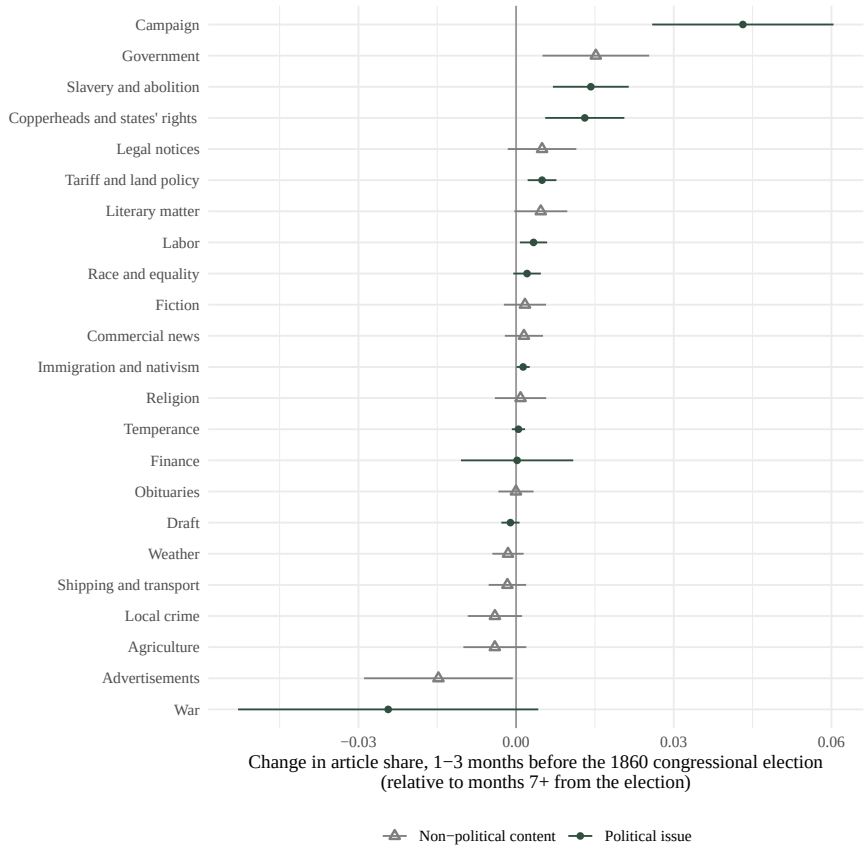


Figure C.1: Pre-election movement by issue, 1860 congressional cycle. Each point is a dictionary's change in article share in the one to three months before the 1860 congressional election, relative to the same newspaper's months at least seven months from it, from the grouped-window version of the paper's event study with newspaper and calendar-month fixed effects; bars are 95 percent confidence intervals from standard errors clustered two-way on newspaper and election date, approximate with 11 election-date clusters; the text reports bootstrap p-values. Filled dark circles are political issues; open grey triangles are non-political content. Source: AmericanStories corpus (Chronicling America), 1860–1865; author's calculations.

for county g and election e, with county fixed effects ω_g, election fixed effects τ_e, and 1860 census controls $\mathbf{X}_g$ interacted with election dummies. Standard errors are clustered by county. Presidential elections use the Republican or National Union share of the total vote; congressional elections use the Republican share of the two-party vote, with uncontested races removed.

Counties with more pre-election slavery coverage voted somewhat more Republican. With census controls, $\hat{\psi} = +2.57$ (SE 0.83), significant at the 1 percent level, across 313 county-elections in 105 counties. Both variables are shares, so each percentage point of pre-election slavery coverage goes with about 2.6 points of Republican vote share. Within a county, coverage varies across elections with a standard deviation of about 1.3 percentage points once county and election effects are removed, so a one-standard-deviation difference goes with about 3.4 points of vote share. Without controls the estimate is +2.68, also significant at the 1 percent level. Split by contest type, the association holds on presidential elections alone (+1.25, SE 0.60, significant; 152 county-elections in 76 counties) and not on congressional elections alone (−2.52, indistinguishable from zero; 53 county-elections in 22 counties).

The pooled estimate mixes contest types in a particular way. Because the county fixed effect absorbs everything constant about a county, the pooled regression compares each county's outcomes across elections, and 101 of the 105 counties are observed in both presidential and congressional contests. Much of the pooled comparison is therefore between a county's presidential and its congressional outcomes, which differ in the vote-share definition, in contestation, and in the candidates on offer, not only in coverage. The congressional-only sample is small because few counties appear in more than one congressional election with usable returns: 23 counties match in 1860 and 17 in 1864, against 100 in 1862.

It is an association. The identifying assumption—that within-county variation in pre-election coverage across elections is uncorrelated with potential vote shares—is demanding and untestable here. The treatment is continuous with no untreated units, a setting in which the parameters of a two-way fixed-effects regression can be hard to interpret (Callaway et al., 2024). And a differential county-level response to the Emancipation Proclamation is not absorbed by the election fixed effect. Nothing here identifies an effect of coverage on votes.

D Robustness Checks

D.1 The Retired Embedding Measure

An earlier version of the coverage measure built topical emphasis from sentence-transformer embeddings (mpnet) reduced with UMAP and clustered with HDBSCAN in the BERTopic framework, scoring each article by cosine similarity to a set of slavery-related topic centroids. On this OCR'd corpus that measure carried essentially no article-level topical signal. On a 10 percent sample, its median area under the ROC curve for picking out articles that contain a topic's keywords was 0.50, no better than chance. Articles containing explicit war vocabulary scored no higher on the war dimension than all other articles (Cohen's d of $+0.001$), and articles naming Antietam, Gettysburg, or Fort Sumter were no different (Cohen's d between -0.015 and $+0.004$). The monthly means were flat, and the density-based clustering produced no event-coherent topics. The keyword dictionary used in the paper replaces it. Its war series peaks at Antietam and Gettysburg; its slavery series does not single out the Emancipation Proclamations (validation subsection of the data section of the paper).

D.2 The Emancipation Proclamation

The Emancipation Proclamation is the one documented slavery event in the study period, and the slavery topic, which carries the emancipation vocabulary, barely registers it. On the paper-demeaned slavery series it ranks the month of the final Proclamation, January 1863, 12th of the 72 study months, at 1.30 times the series mean, and the month of the preliminary Proclamation, September 1862, 34th (0.99 times the mean). That matters for the 1862 cycle, because the preliminary Proclamation of September 22, 1862 fell inside its pre-election window, so a Proclamation-driven spike could inflate that cycle's coefficient on the slavery measure. There is no such spike in September, and the robustness section of the paper shows that the 1862 cycle carries no pre-election rise on its own, with an estimate of -0.0001, indistinguishable from zero.

D.3 Alternative Coverage Measures

The event-study results in the paper are reported for the dictionary share of articles s_{slav} and its word-weighted counterpart. They agree in sign on the break and on the 1860-cycle level against the two-sided reference. The normalized relative measure R_{jt} behaves differently and should not be read as tracking s_{slav}. Pooled across all papers and cycles it falls significantly before elections, by -0.62 within-paper standard deviations, because campaign coverage expands faster than slavery coverage does. In the 1860 cycle it is indistinguishable from zero for the press as a whole (-0.31) and for Republican papers (-0.43) under the bootstrap. Among

1860 Republican papers the scale-free ratio $R_{jt}^{\text{ratio}} = s_{slav}/(s_{slav} + s_{pol})$ is also indistinguishable from zero (-0.006). Slavery rose with the campaign in 1860, not ahead of it.

D.4 The Earlier Slavery Dictionary

An earlier version of the slavery dictionary counted an article on a single matching term, and so caught more slavery articles than the two-entry rule the paper uses. Three results differ between the two. First, on the earlier version the 1860 break was about 27 percent larger relative to its baseline, and it survived dropping the October 9 states. On the articles that both versions flag, the break is as large as it was on the earlier version, so the current measure attenuates the break rather than contradicting it. Second, the Republican break across the 1860 vote was significant on the earlier version; on the paper's measure it is not. Third, within cities the Democratic-minus-Republican gap in the slavery article share was negative and far from significant on the earlier version, where on the paper's measure Democratic papers ran marginally more slavery coverage than their neighbours. The Republican break and the within-city topical gap therefore depend on how slavery coverage is measured.

D.5 Time-Index Sign Convention

The analysis panel stores event time as the calendar month minus the month of the attributed congressional election, so negative values precede the election and zero is the month containing election day. The panel's construction verifies that $k = -1$ is the calendar month before the election month. Earlier panels used a 30.44-day bin with the opposite storage sign; no estimate in the paper uses them.

References

Broder, A. Z. (1997). On the resemblance and containment of documents. *Proceedings of the Compression and Complexity of Sequences 1997*, 21–29. https://doi.org/10.1109/SEQUEN.1997.666900

Callaway, B., Goodman-Bacon, A., & Sant'Anna, P. H. C. (2024). *Difference-in-differences with a continuous treatment* (NBER Working Paper No. 32117). National Bureau of Economic Research.

Clubb, J. M., Flanigan, W. H., & Zingale, N. H. (1987). *Electoral Data for Counties in the United States: Presidential and Congressional Races, 1840-1972: Version 1.* ICPSR - Interuniversity Consortium for Political and Social Research. https://doi.org/10.3886/ICPSR08611.V1

Dell, M., Carlson, J., Bryan, T., Silcock, E., Arora, A., Shen, Z., D'Amico-Wong, L., Le, Q., Querubin, P., & Heldring, L. (2023). American stories: A large-scale structured text dataset of historical U.S. newspapers. *Advances in Neural Information Processing Systems 36 (NeurIPS 2023) Datasets and Benchmarks Track*. https://arxiv.org/abs/2308.12477

Schroeder, J., Van Riper, D., Manson, S., Knowles, K., Kugler, T., Roberts, F., & Ruggles, S. (2025). *National Historical Geographic Information System: Version 20.0.* Minneapolis, MN: IPUMS. https://doi.org/10.18128/D050.V20.0

Smith, D. A., Cordell, R., & Mullen, A. (2015). Computational methods for uncovering reprinted texts in antebellum newspapers. *American Literary History*, *27*(3), E1–E15. https://doi.org/10.1093/alh/ajv029