Book Bans in American Libraries: Impact of Politics on Inclusive Content Consumption

Uttara M. Ananthakrishnan, Naveen Basavaraj, Sabari Rajan Karmegam, Ananya Sen and Michael D. Smith*

January 10, 2025

*The author names are ordered alphabetically by last names. The authors are grateful to David Christensen from The Seattle Public Library for facilitating access to library anonymized circulation data, which made several analyses possible. The authors are also grateful to Ari Lightman for generously providing access to Twitter data.

Abstract

Banning of books has become increasingly prevalent and politically polarizing in the United States. While the primary goal of these bans is to restrict access to books, conversations about the bans have garnered attention on a wider scale. This increased attention to bans can either have a chilling effect or can influence consumers to read the banned books. In this study, we use a novel, large-scale dataset of US library book circulations and evaluate the impact of high-profile book bans on the consumption of banned books. Using a staggered difference-in-differences design, we find that the circulations of banned books increased by 12% on average compared to comparable non-banned titles after the ban. We also find that banning a book in a state leads to increased circulation in states without bans. We show that the increase in consumption is driven by books from lesser-known authors, suggesting that new and unknown authors stand to gain from the increasing consumer support. Additionally, our results demonstrate that books with higher visibility on social media following the ban see an increase in consumption, suggesting a pivotal role played by social media. Using patron-level data from the Seattle Public Library that includes the borrower's age, we provide suggestive evidence that the increase in readership in the aggregate data is driven, in part, by children reading a book more once it is banned. Using data on campaign emails sent to potential donors subscribed to politicians' mailing lists, we show a significant increase in mentions of book ban-related topics in fundraising emails sent by Republican candidates. We also provide suggestive evidence on the impact of the rhetoric around these events on donations received by politicians.

1 Introduction

Book bans or challenges in public schools and libraries in the US have been recurring events over the past several decades (Foerstel 2002). However, since 2021, there has been a significant push from state officials, elected representatives, individual parents, community members, and advocacy groups to remove books construed to deal with sensitive topics related to race and gender from public and school libraries. While such challenges are not new, the number of book bans in a short period has made these events a subject of national debate due to their politically polarizing nature. For example, the American Library Association (ALA) released a statement indicating an "alarming" increase in book ban campaigns around the US.[1] PEN America finds that books have been banned at an "unprecedented" rate just between 2021 and 2022 in 138 school districts. Critics of book bans argue that they attempt to censor information and ideas and violate the First Amendment rights of students and library patrons. However, proponents of these bans refer to the broad discretion afforded to school boards by the Supreme Court and the rights of parents to voice concerns about the suitability of themes dealing with race and gender identity for children under a certain age.

The potential impact of such polarizing events could create a strategic dilemma for various stakeholders, including authors, publishers, and politicians. However, little is known about the impact of contentious rules and legislation designed to restrict access to goods at the local level on the consumption of these goods on a larger scale. We study this phenomenon in the context of book bans instituted by local school districts and city public libraries. Our focus will be on book titles that are relatively high-profile due to them being banned repeatedly. Although the decision to ban books is made at a 'micro-geographic' local level, these events receive extensive attention on social media and media coverage at the state and national level outside of those particular areas. A priori, the impact of book bans on the consumption of such content associated with diversity and inclusion topics is ambiguous.

[1] https://www.ala.org/news/2022/04/ala-submits-comments-opposing-book-banning-house-oversight-committee-hearing

On the one hand, books banned in select locations garner state-wide and national attention, which can increase the readership of these books (a phenomenon popularly known as the "Streisand effect" (Jansen and Martin 2015)). On the other hand, the sensitive nature of these book bans might result in a chilling effect, deterring consumption as parents and teachers opt to avoid controversy in a highly polarized environment. The impact of 'hyper-local' decisions on consumption in the aggregate can significantly impact the content strategy of authors and publishers, as well as the rhetoric used by politicians.

In this paper, using novel large-scale data on US library circulation, we investigate the impact of book bans on the circulation of the top-25 most banned book titles (each banned more than five times), as defined by PEN America and ALA, at school libraries and public libraries, which is our key measure of book demand. To identify the causal effect of book bans, we use the large variation in the timing of book bans across different states in a difference-in-differences event study design. Furthermore, given the widespread nature of such ban events, we study the spillovers of these bans across several dimensions. We use a variety of identification strategies, such as defining one single ban event across all books, 25 ban events (one for each title), and analyzing 127 titles $\times$ state events to analyze the impact of ban events on library circulations. These top-25 titles appear in over 85% of ban events in our sample and are most likely to be 'solo-banned', leading to increased salience on news and social media. We also use data from Amazon Sales Rank to validate our key baseline result.

Across a series of specifications, using different sources of identifying variation, we find that book ban events increase the circulation of banned books by 12% compared to a set of control books. We also find spillovers of the impact of bans with an increase in readership in non-focal states. More specifically, even in states without bans, the circulation of banned books increased by 11.2% compared to control books. Interestingly, social media plays a significant role in amplifying this impact. Books that attract high visibility on Twitter enjoy significantly higher circulation after the event than books that do not. Our effects are driven

entirely by books written by less famous authors, suggesting that these events shine a light on the works of lesser-known writers. Moreover, we find that this increase in consumption is prevalent in both Blue and Red states and spills over to states that do not see any bans.[2] That noted, we find that the focal state that experiences the book bans sees a larger increase in circulation relative to other states. To dig deeper into the mechanism, we use patron-level data that includes the age of the individual from the Seattle Public Library to show that after book bans are initiated in other states, there is an increase in readership of those books by children under the age of 18.[3] We validate our baseline results using Amazon sales rank data. We find a statistically and economically significant decline in sales rank, implying an increase in Amazon sales after the ban events.

In general, this increase in content consumption is related to inclusion themes, as the topics of these books often deal with issues of race, gender, and LGBTQIA. The critical question here, however, is why politicians bring such issues to the forefront, (potentially) inadvertently increasing the consumption of these books. To further analyze this issue, we explore the potential to incorporate these issues as part of political rhetoric. Using emails from a set of 245 political accounts (candidates, PACs, etc.), we analyze campaign emails to small-dollar donors and identify a marked rise in mentions of book bans, especially in communications from Republican politicians. We provide correlational evidence to suggest that transforming book bans into a political issue or debate — what we specifically define here as 'politicization' — tends to increase the amount of donations received by Republican House candidates relative to Democratic House candidates, but only in conservative states.

Our findings have strategic and managerial implications for a variety of stakeholders. A priori, high-profile book bans could discourage authors and publishers of books associated with diverse and inclusive content from creating such content due to potential chilling and reputational effects. This can have large downstream implications, such as the reduction in

[2]'Red' states voted Republican, and 'Blue' states Democratic, in the 2020 presidential election.

[3]This analysis comes with the caveat that we were only able to get data on children's reading behavior from one city. State laws often prohibit storing the reading data of patrons.

the supply of representative and diverse content available to the general public. However, our results suggest that there may be a broad audience for content beyond the local public library or school district level, where contentious rules target such books. Therefore, in the aggregate, such contentious rules could unexpectedly lead to an increase in the consumption of such content, at least in the short run. Strategically for an author, it is important to generate chatter on social media and news coverage in the aftermath of book bans to increase readership. Moreover, in terms of political strategy, politicians need to realize that while such rhetoric can potentially lead to more donations, it might also inadvertently amplify the content they seek to suppress. This unexpected outcome highlights the power of social media and public discourse in shaping book consumption patterns. In light of these findings, it becomes crucial that stakeholders consider the broader implications of their actions and strategies.

We contribute to several strands of Marketing and Economics literature. First, we contribute to the broad literature around "political consumerism" which is the act of consumers trying to influence business practices using their purchasing power or "voting with their wallets" (Stolle et al. 2005, Stolle and Micheletti 2013). Schoenmueller et al. (2023) show how brand preferences can predict political preferences and how increasingly polarized brand preferences are correlated with the trend in political polarization. Studies that aim to identify the causal impact of political polarization driving consumption are relatively new. In this emerging stream, Liaukonytė et al. (2023) study the effect of the CEO of a brand (Goya) taking a political stance by praising President Donald Trump to find a (temporary) increase in sales by 22%. Wang and Lu (2020) show that when it was accidentally disclosed that a company donated to a Republican politician, sales of that good increased in Republican counties. Bursztyn et al. (2022) demonstrate how politically polarized debates around the Affordable Care Act ("ObamaCare") lead to lower participation in more Republican counties.[4] Mumma (2022) and Goncalves et al. (2024) take a descriptive approach in providing

[4]There are related papers which show how political preferences can shape decisions in different settings credit ratings and online lending behavior Kempf and Tsoutsoura 2021, Wang and Overby 2022.

the landscape of book bans and the role of politics in determining the availability of children's books in school libraries. We believe that these papers provide the foundation for our study to take the next step in analyzing causal effects. We complement this literature by providing causal evidence on how consumption of books increases in the aggregate in the aftermath of (high-profile) book bans, in the process shedding light on a context where formal restrictions exist on access to certain products at the local level or legislations are proposed around sensitive social issues. Such conditions, in turn, create a polarized environment that could impact the consumption of the good on a broader scale. In addition, we highlight a feedback loop that identifies the impact of polarizing political rhetoric through such issues in campaign emails.

Next, we contribute to the marketing literature that studies how the decision to censor a particular product leads to more scrutiny and explosion of publicity for this product. This effect, popularly dubbed the "Streisand Effect", has been exacerbated in the age of social media (Jansen and Martin 2015, Mach 2022). Studies in this literature have examined the impact of publicity, both positive and negative, as a form of advertising. For example, (Berger et al. 2010) study online reviews for books and concludes that any publicity is good for relatively lesser-known products as reviews, including negative ones, increase product awareness and lead to increased readership. Similarly, Ahluwalia et al. (2000), Ackerberg (2001) study the different mechanisms behind the potential role of negative publicity in informing consumers' purchase decisions. More recently, Shipilov et al. (2019) found that both positive and negative media coverage influences the adoption of governance practices by firms. In our paper, we find evidence to suggest that banning books provides informational value about these books on the state and national stage. In particular, we find that book bans lead to greater coverage of these books on social media. This informational effect leads to increased consumption not only in the focal states where some public or school libraries banned these books but also in states where these books were not banned. Our results shed light on the fact that censorship in a politically polarized environment provides informational

value on the books being banned, leading to higher consumption of the restricted product. These results come with an important caveat. Our results do not speak to how consumption would be impacted if such bans took place at a more aggregate level, for example, if there was a federal ban on any particular book.

Finally, our paper relates to the studies examining the effectiveness of political marketing and campaigns. For example, Gordon and Hartmann (2013, 2016) study the different factors that influence the level of advertising in political campaigns. Studies have also analyzed how political parties leverage the media and creative elements to communicate with the electorate (Soberman and Sadoulet 2007). More recently, Petrova et al. (2021) study how social media aids new politicians to communicate with their constituency effectively and helps mitigate the incumbency gap with their competitors. Further, Fujiwara et al. (2021) quantify the role of Twitter in the 2016 US Presidential Elections, highlighting that Twitter lowered Republican vote share. However, the impact of using contentious policies, rules, or legislation as a potential tool to raise funds, with social media playing a moderating role, is relatively understudied. We find that the conservative stance, which prompts such book bans, seems to lose out on the focal dimension of readership of those books more generally, at least in the short run. On the other hand, we provide suggestive evidence that (some) Republican politicians gain donations, which may provide them with an electoral advantage.

2 Background and Data

2.1 Background: Libraries and Book Bans in the United States

Public and school libraries play an important role in the dissemination of information for a large majority of adults and children in the United States. Recent research documents the positive effects of public libraries on educational outcomes (Karger 2021, Gilpin et al. 2021). Public libraries in the United States originated in the 19th century, financed by Andrew Carnegie (Berkes and Nencka 2021). In 2019, the United States had 16,548 public

libraries and accounted for a large majority of book consumption in the country (Reimers and Waldfogel 2022). More than half of public schools in the United States have a full-time librarian, and often schoolchildren have to go to the school library once a week as part of their regular schedule (Gavigan et al. 2010). These institutional details highlight the key role of public and school libraries in allowing adults and children to access books.[5]

The concept of banning books in classrooms, schools, and public libraries dates back several decades in US history.[6] While proponents of book bans propose limiting the circulation of certain books to protect children from potentially "inappropriate" content or "abusive" language,[7] opponents usually see them as censorship events that violate students' First Amendment rights.[8] Attempts to restrict access to books have faced several legal challenges, with the Supreme Court weighing in on several occasions.[9]

More recently, calls to ban books have increased significantly across the United States. According to a recent report by the American Library Association (ALA), there were 729 unique attempts to ban books in public and school libraries in 2021 and 1,269 attempts to ban books in 2022.[10] Book bans have become part of the national conversation and cultural debates, fueled by increasing concern about parental rights.[11] We adopt the definition provided by PEN America which is that "a book ban occurs when an objection to the content of a specific book or type of book leads to that volume being withdrawn either fully or partially from availability, or when a blanket prohibition or absolute restriction is placed on a particular title". Private individuals, government officials, or organizations can initiate the

[5]Mumma (2022) also provides some initial descriptive evidence on the availability of different types of books, including those dealing with topics of diversity, in different libraries.

[6]https://www.ala.org/advocacy/bbooks/frequentlychallengedbooks/classics

[7]https://www.ala.org/advocacy/bbooks/banned-books-qa

[8]https://njsbf.org/2023/02/15/does-banning-books-violate-the-first-amendment/

[9]In Island Trees School District v. Pico (457 U.S. 853, 1982), the Supreme Court ruled that school boards cannot remove books from their libraries simply because they are considered offensive. However, in Hazelwood School District v. Kuhlmeier (484 U.S. 260, 1988) and Morse v. Frederick (551 U.S. 393, 2007), the Court ruled that the First Amendment does not protect obscene materials in schools. These two sets of rulings provide significant challenges in interpreting the appropriateness of book bans in schools and public libraries.

[10]https://www.ala.org/news/press-releases/2023/03/record-book-bans-2022

[11]https://www.washingtonpost.com/education/2021/10/28/parental-rights-in-schools-untenable/

book ban process. We constructed some statistics based on publicly available PEN America data and additional disaggregated data we acquired from the American Library Association. Based on this data, parents initiate 34% of the book bans, while pressure groups (e.g. Moms for Liberty) initiate 31%. Administrators of libraries and other patrons account for 12% and 9.6%, with a long tail of elected officials, teachers, librarians, etc. School districts typically do not announce book bans to the public. However, many of the ban requests made during school board meetings are typically covered by local newspapers. During the review process, the book may be temporarily removed from the shelves. The result of the process can lead to the book being retained, restricted to a subset of patrons, or removed completely from collections. If a book is banned, it can last a few months to an indefinite period until the bans themselves are challenged or reversed, and in some cases, they last until the book is reviewed and potentially reinstated. Some books are permanently banned if the decision is not reversed.

Book bans in our dataset, and more generally, happen at a small geographic level, such as a public library or a school district. Since book bans happen at the public library or school district level, states that are both Republican (e.g., Texas) and Democratic (e.g., New York) experience book bans. We find that there can be a significant overlap in the books banned by school and public libraries, with 45% of the books being banned by both. In 90% of the cases, books are first banned in school libraries. Our circulation data are aggregated at the title state month, implying that we cannot identify what happens at the library level. However, we will be able to identify the overall effects at more aggregated levels. The titles that are banned often deal with diversity and inclusive topics related to Race, Gender, and Sexuality. Some examples include Out of Darkness, This Book is Gay, Gender Queer, etc. The complete list of banned titles in our data, along with genre and author information is available in Table A1 in the Appendix.

Book bans have also become politically polarizing and play a central role in discussions on social media, attracting arguments for both parental rights and free speech. Book bans

have also received endorsements from politicians and are often used in electoral campaigns.[12] Hence, a priori, the impact of book bans on a wider readership could go either way. Given the polarizing nature of these issues, it could have a chilling effect when parents, teachers, and librarians refrain from advocating for such books. With potentially controversial topics, there might be an unwillingness to encourage reading and discussions around them, as shown in the case of the Internet search and potential government intervention (Marthews and Tucker 2017). On the other hand, attention to these bans can create awareness and increase readership more broadly. Theoretically, this is in line with the Streisand Effect. The Streisand Effect is named after an attempt by singer Barbara Streisand to limit the viewership of a picture of her house that paradoxically led to an increase in views due to the attention on the issue.

2.2 Data

2.2.1 Book Circulation Data

We obtain book circulation data from a large library content and services supplier to major public and academic libraries in the U.S. The supplier provides physical and digital books, library maintenance software, and analytics solutions to libraries. The data cover circulations from January 2021 to June 2022 across 38 states in the US. The data is aggregated at the Title-ISBN-Month-State level. We aggregate the data to create a balanced panel of circulation data at the Title-Month-State level. Given the Non-Disclosure Agreement with our partner vendor, we cannot report absolute levels of book circulations.

We primarily focus on the book titles that PEN America and ALA identified as the most commonly banned books in the US between July 2021 and June 2022. Hence, apriori, these banned books are potentially more high-profile than other banned books that are not banned as often. The PEN list overlaps significantly with the ALA's list of banned books in public

[12] https://www.politico.com/news/2022/09/22/house-and-senate-democrats-prepare-resolutions-to-oppose-local-book-bans-00058200

libraries.[13] For each book title, we obtain the month of the ban event and the state where the ban originated based on these lists. It is important to note that all the book bans we observed in the data were restricted to a particular school district or public library.

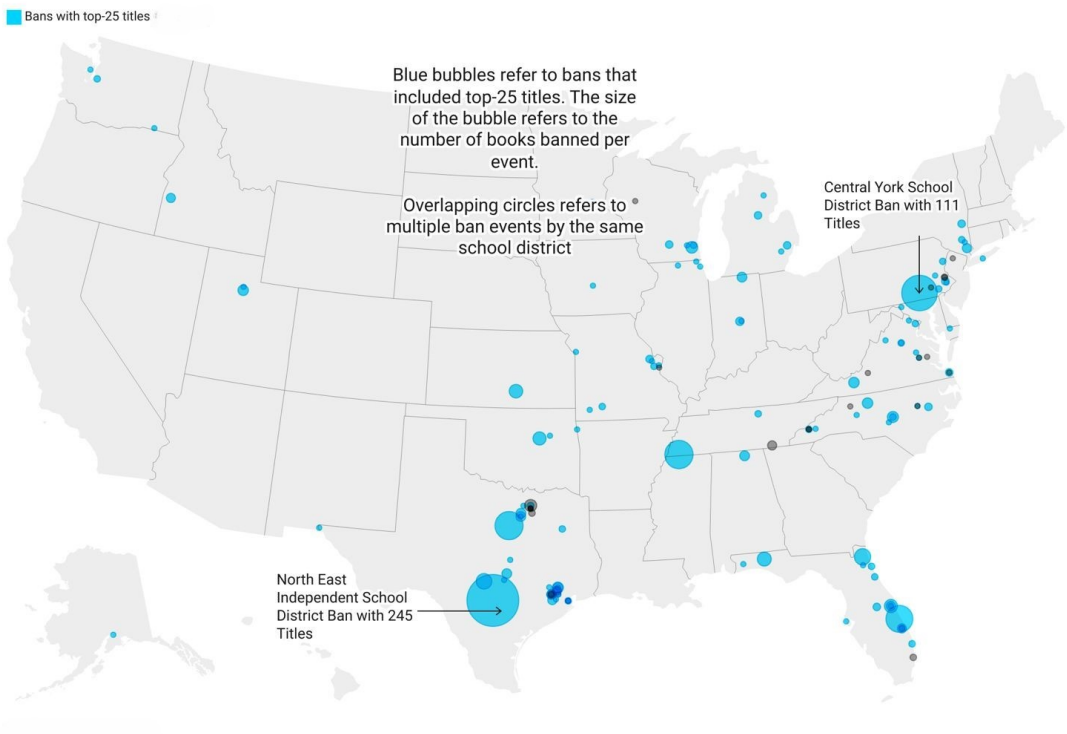
To limit the role of confounders and unobservables, we do not consider banned books published before 1 January 2010, as many have a long history of being banned and have been long-standing TV shows or movies. We also remove banned books whose circulations were impacted by the Black Lives Matter movement unrelated to the book bans - for example, following the murder of George Floyd, circulations of "How to Be an Antiracist" book surged, and the book eventually became a top seller on Amazon.com. This exercise leaves us with a sample of 468 books banned in 24 states. We have circulation data for 38 states, including the 24 states that experienced bans.

We aim to identify the aggregate market-level impact of book bans, for which the salience of these ban events due to social and traditional media coverage is a necessary (but not sufficient) condition for consumers to respond at scale. To do so, and define our analysis's scope, we focus on the top-25 most banned books in the ALA and PEN America lists. These 25 books based on ALA and PEN lists were significantly more salient for several reasons despite 468 banned books in 939 district-title events. These 25 titles are not outliers but are, in fact, the rule rather than the exception. In particular, our sample period saw 'mega ban' events where eight or more (90th percentile) books were banned together, with the median number of books banned in any book ban event being 1.5. We find that 80% of the banned books in the overall set of 468 books were banned only as a part of these mega ban events, while our sample of 25 were banned in 85% of all bans. To illustrate this, in Figures 1 and 2, we plot all the bans in the form of a bubble chart, shaded in gray, with the size of the bubble representing the number of unique titles banned. We can spot the mega bans (with eight or more books banned) as the larger bubbles. In Figure 1, we distinguish the events that our top-25 books were part of, shaded in blue. In Figure 2, we mark the events that our non-

[13] See here for the ALA list `https://www.ala.org/advocacy/bbooks/frequentlychallengedbooks/top10` and here `https://pen.org/index-of-school-book-bans-2022/` for the PEN list.

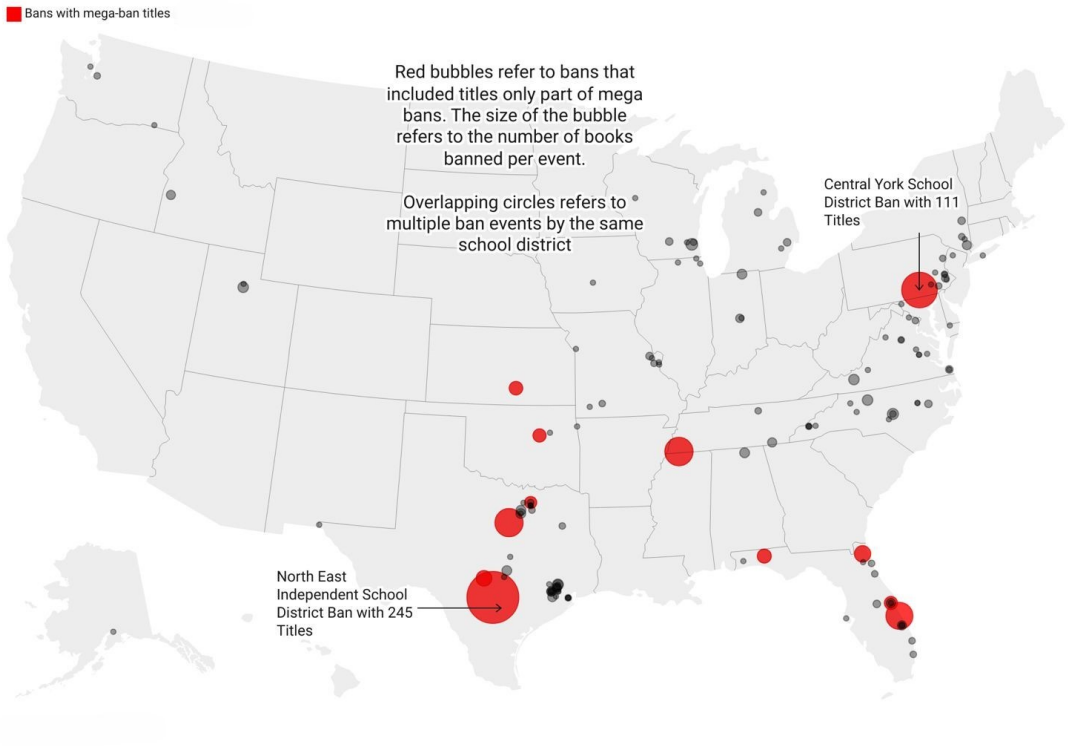
top-25 books were part of in red. We can immediately observe that our 25 books were part of a significant number of bans across geographies and time, whereas the non-top-25 books were almost only part of the larger mega bans. This plot is also useful in highlighting other key attributes that impact the attention received by any particular banned book. These top-25 titles are significantly more likely to be 'solo-banned' – that is, they were banned alone and not with any other books, are more likely to be banned with fewer other books on average, and are likely to be banned simultaneously across states. This indicates that the top-25 titles in our sample do not need to share attention with other books and receive a lot of scrutiny in the media about the nature of the book, the content, and the reason why they are requested to be pulled off the shelves.

Figure 1: Ban Coverage - Top-25 Titles



Note: The figure shows the ban events where the top-25 titles are banned.

Figure 2: Ban Coverage - Mega Ban Titles



Note: The figure shows the ban events where titles that were banned only as part of mega ban events are banned.

Next, we directly test whether our sample of 25 books received more attention relative to other banned books as a result of the patterns highlighted above. In Table 1, we report descriptive results that show that our sample got significantly more news (column (1)) and social media (column (3)) coverage relative to the rest. Additionally, even in months with mega-ban events when several books were banned together, the 25 titles were likely to get significantly more coverage than the rest of the sample (columns (2) and (4)). Using Goodreads data, we find that users were 46 times more likely to put a book from our top-25 sample into a 'banned-book shelf' than the rest of the banned books, which indicates that consumers are more aware of the top-25 books as being banned. This is rather intuitive since these top-25 books are part of 85% of all bans in the U.S, across different states, which

might lead to higher awareness about these top-25 books as banned compared to the non-top-25 books. This analysis suggests that our sample of book titles is more likely to be in the spotlight and become salient in the context of book bans happening around the country, which is important for market-level effects on consumption.

Table 1: Impact of Ban Events on Coverage

Dependent Variables:	Log(News Article Count)		Log(Tweet Count)	
	All Months	Mega Ban Months	All Months	Mega Ban Months
	(1)	(2)	(3)	(4)
Variables				
Top-25	0.406***	0.484***	0.856***	0.961***
	(0.136)	(0.162)	(0.202)	(0.225)
Year Month Fixed Effect	Yes	Yes	Yes	Yes
R^2	0.07423	0.08157	0.08415	0.05637
Observations	4,680	3,276	4,680	3,276

Note: Robust standard errors clustered at the level of the title in parentheses. The unit of observation is title-month. The dependent variable is log of news article count or tweet count per title per month. Top-25 is an indicator variable which denotes if the title is a top-25 title or not. Columns (1) and (3) include all the months after October 2021 (the month when the first significant ban event occurred). Columns (2) and (4) include only those months after October 2021 which had a mega ban event.
$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

The books that we considered in our analysis were mainly in the fiction and non-fiction genres, covering topics or themes such as LGBTQ, Race, Diversity, etc., as identified by our data partner. Our Goodreads shelving data confirmed this, as seen in column (2) of Table 2. Next, for each of these titles, we identify the months when the title was banned in a particular school district or public library using our banned book lists.[14] We also use a set of books to construct our control group(s). We elaborate on those steps after discussing the empirical framework.

[14] `https://pen.org/report/banned-usa-growing-movement-to-censor-books-in-schools/`

Table 2: Shelving on Goodreads

Dependent Variables:	Ban-Related (1)	LGBTQ-Related (2)	School-Related (3)
Variables			
Top-25	184.5***	741.6***	871.8**
	(7.93)	(283.8)	(366.8)
(Intercept)	4.54**	512.3***	813.0***
	(1.87)	(67.1)	(86.8)
R^2	0.54901	0.01511	0.01253
Observations	447	447	447

Note: The unit of observation is title. The dependent variable is the number of times a book is added to a particular shelf on goodreads. Top-25 is an indicator variable which denotes if the title is a top-25 title or not. Column (1) denotes ban-related shelves, column (2) denotes lgbtq-related shelves, and column (3) denotes school-related shelves.
$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

2.2.2 Goodreads Reviews Data

As an alternative data source for social chatter, we use Goodreads, a popular cataloging website owned by Amazon with over 90 million members. We use this popular book review platform to identify reviews and numerical ratings for each book in our sample. We also extract the review text (and thus, the length of the review), the date the review was written, and the reviewer details for the period between January 2021 and June 2022. Goodreads also has a feature in which the date the user starts to read a book and the date that the user finishes a book are recorded for most reviews. In fact, some of this is automatically logged for books read on Kindle. This also allows to measure if there is a change in the time taken to read the books before and after the ban. Finally, we obtained details on the shelving or tagging of books by readers (categories such as banned books, LGBTQ, etc) to shed light on how people perceive the topics of the book.

2.2.3 Seattle Public Library Checkout Data

To study how children consume banned books, we use de-identified patron level checkout data from the Seattle Public Library. The Seattle Public Library (SPL) is among the top-25 public libraries in America, with a collection of over 2.4 Million. SPL has more than 17 million patron visits in the 27 branches with an annual circulation of more than 12 million books.[15] Our data are available at the patron title date level for the period between 01/01/2021 and 12/31/2022 and contain information on the age group of the patrons.[16] SPL classifies patrons into different age groups, such as "Adults", "Seniors", "Young Adults", and "Juniors", depending on the birth date provided at the time of joining SPL. Public schools in Seattle also allow schoolchildren to use their school ID to access SPL books. In this case, the age is coded as "Seattle Public School". This does not include teachers because they are coded as "Seattle Public School Teachers." We focus our analysis on Seattle and King County patrons and those whom the Seattle Public Library classifies as young adults (born between 2005 and 2009), juniors (born between 2010 and 2021), or from Seattle Public Schools (SPS checkouts).

To provide some descriptive evidence that young readers are potentially consuming these books, based on the literature (Zhu et al. 2023, Agarwal and Sen 2022), we collected additional Goodreads school shelf or tag data to identify a book's relevance to school-related reading or educational purposes. In column (3) of Table 2, we find that the top-25 titles in our main banned sample are shelved by Goodreads readers on school-related shelves (such as "K-8", "pre-K" reading, "read with kiddos", etc.) significantly more than the rest of the banned books in our sample. This provides us with some descriptive evidence of our top-25 titles' readership among young readers.

[15] https://www.spl.org/about-us/library-impact/2019-impact-report/2019-statistics

[16] We acknowledge the tremendous help from David Christensen at the Seattle Public Library and Seattle Open data by providing the data.

2.2.4 Political Campaign Emails Data

To identify the issues political candidates highlight in their campaign communication strategies, we use political email data from the Democratic and Republican party candidates. Political committees and candidates send regular campaign emails to individuals who subscribe to their mailing lists, highlighting the key issues of their political strategies. We use a dataset created by Derek Willis, who teaches journalism at the University of Maryland, by subscribing to mailing lists of various political candidates from across party lines.[17] Although the list of candidates is not exhaustive, the data can provide insight into strategies politicians use to communicate with their potential donors and voters, similar to other recent papers in the literature (Mathur et al. 2023). Within these emails, we focus on emails that refer to book bans or book ban-related keywords.[18] Overall, our data contain emails from 245 unique email IDs during the period between 01/01/2021 and 12/31/2022. Some of the most prominent accounts belonged to candidates such as Ron DeSantis, Jim Jordan, and Vivek Ramaswamy and PACs such as the 1776 Project and the American Federation for Children. We found that 39 states were represented in our emails, with all 24 states that experienced bans as part of the 39. This implies that the emails' text matched the locations where bans took place.

2.2.5 Political Contributions Data

We identify individual-level contributions to all Republican and Democratic parties from the Federal Election Commission's political donation data. We focus on all donations to candidates during the 2021-2022 election cycle.[19] The data contain information on each donor's home state, donation amount, and the party (also referred to as committee) the

[17] `https://createsend.com/t/t-97F63A7D578A8F0B2540EF23F30FEDED`

[18] The complete list of keywords used are the following - "book ban", "book censor", "censorship", "restricted access", "banned book", "intellectual freedom", "freedom of speech", "content control", "challenged books", "parental rights", "woke agenda", "radical left", "liberal indoctrination", "pornography", "critical race theory"

[19] The FEC tracks donations to political candidate committees if the total election cycle-to-date contribution amount is above $200 - `https://www.fec.gov/data/browse-data/?tab=bulk-data`

candidate is affiliated with. We also use FEC committee candidate links to match a unique candidate with contributions to their committee.

2.2.6 Twitter Data

In our analysis, we also use social media data, focusing on Twitter. We use Brandwatch, a popular social media analytics platform[20] to fetch all the geotagged tweets from the US that contain the name of banned or control titles and their respective author names in the text of the tweet. For each tweet, we obtain the text, URL, Twitter user ID, follower count, location, and total impressions for each tweet.[21]

To create a measure of social media visibility, we identify tweets that mention titles and their respective authors between January 2021 and July 2022 and classify titles with greater than the total median impressions on Twitter as "High Visibility". Additionally, to construct an alternate measure of author experience, we count all the tweets that mention the banned book authors' names between January 2012 and December 2020. We then split the sample of authors into more experienced and less experienced authors based on the median count of tweets.

3 Empirical Framework and Control Group

3.1 Empirical Framework

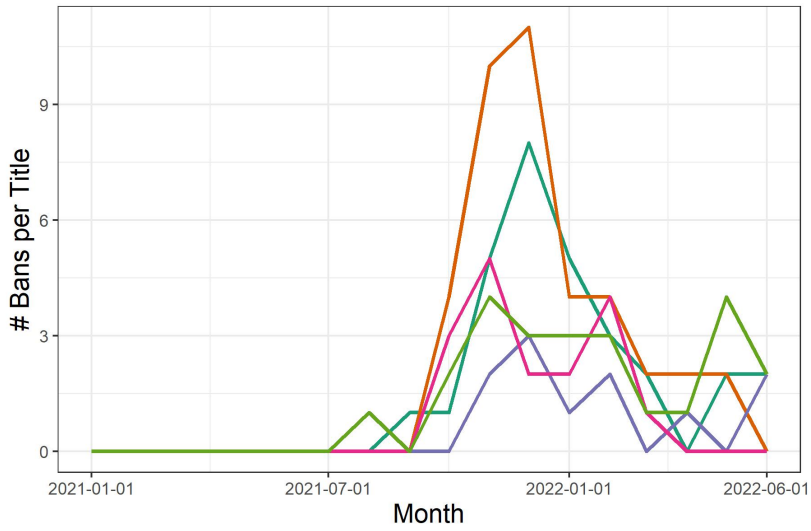
Although different books are banned at different times, the bans for each title continue to cascade after the first time a title is banned in any state, as shown in Figure 3 with the five most banned book titles.

Moreover, once a title gets banned in one state, local news media and social networks highlight ban events across the focal state as well as non-focal states that do not experience

[20] `https://www.gartner.com/reviews/market/social-monitoring-and-analytics`

[21] Impressions correspond to the sum of views, likes, and retweets for the tweet.

Figure 3: Number of Bans per Title



Note: The figure shows the total number of ban events per title per calendar month for the five most banned books in our sample. Each line represents a title.

the bans, which can affect the demand for these titles in other states. Therefore, we define our treatment at the title level, and once a title is treated (banned in the first state), it remains treated for the rest of our analysis.[22] Our identification relies on different titles being banned at different times, providing us with a staggered Difference in Differences (DiD) framework.

We estimate the DiD model using our main specification shown in the following equation:

$$y_{ijt} = \alpha_i + \gamma_j + \tau_t + \beta \times PostBan_{it} \times Banned_i + \epsilon_{ijt} \tag{1}$$

where $PostBan_{it}$ is the event variable which takes a value of 1 if the month t is following the first time title i is banned in any state and 0 otherwise.[23] For all control titles, the variable takes a value of 0 for all months. $Banned_i$ indicates whether the title i is in the list of books eventually banned. α_i, γ_j, and τ_t capture fixed effects for title, state, and month, respectively. The coefficient of interest β captures the average impact of banning a title

[22] We test alternative event definitions to ensure the robustness of our results.

[23] Similar to Agarwal and Sen (2022), there is generally a lag between when a ban is initiated and when it becomes known to the general public. Moreover, many ban events happen towards the second half of the month. We test the robustness of our results in relation to this definition later in the paper.

in any state on circulations. We use a four-month pre-period before the ban event and a six-month post-period window.[24] The standard errors are clustered at the state level.

We verify the parallel trends assumption using the event study specification.

$$y_{ijt} = \alpha_i + \gamma_j + \tau_t + \sum_{m=-4}^{m=6} \beta_m \times 1[t - e_i = m] \times Banned_i + \epsilon_{ijt} \tag{2}$$

The indicator variable $1[t - e_i = m]$ takes a value of 1 if the calendar month t is m months apart from the ban event e_i. The baseline period is -1, which means that the average change in the circulations of banned titles compared to the circulations of control titles is measured relative to the difference in circulations one month before the ban. We use an array of estimators suggested in recent literature to account for potentially heterogeneous treatment effects in staggered DiD settings to verify the robustness of the Two Way Fixed Effects (TWFE) estimates (Callaway and Sant'Anna 2021, Sun and Abraham 2021, Borusyak et al. 2021)

It is pertinent to note that our broad empirical approach follows the intuition of studies in the literature that focus on the impact of salient events on consumption behavior (Liaukonytė et al. 2023, Wang and Lu 2020, Agarwal and Sen 2022). We also broaden the approach relative to these studies, which often analyze a single salient event. In particular, we use a range of event definitions to ensure robust results across the board. While our conservative estimation approach uses the first ban event of a title as our main treatment event, we also analyze 127 title-state bans as events in our robustness analysis that significantly increase the number of events. At the other end of the spectrum, we also use the timing of the very first book ban event in our data as the sole event leading to a single event analysis, as in the studies above. Both of these approaches have their pros and cons. The drawback of using a single event is that of external validity. The drawback of having a large number of events is the potential for spillovers and overlapping time windows across events. Hence, by focusing

[24] We also test the robustness of the results to a different time window.

on 25 title level events, we believe we have taken a middle path while being grounded in both external validity and econometric logic.

3.2 Control Group

We outline the logic behind identifying relevant books that might serve as an appropriate control group. From our partner, we acquired additional circulation data for books that could be part of the control group. Our first concern is to limit the spillovers of the treatment to the control books (Agarwal and Sen 2022). We also do not include books in the control group that have the same BISAC codes as banned books. The BISAC Subject Codes are standard categories used by publishing companies, booksellers, and libraries. BISAC Subject Codes represent a sub-categorization of genres. For example, FIC049000 refers to books with African-American themes, and FIC022030 refers to mystery novels, even though both BISACs fall into the fiction genre. Therefore, when selecting a control group, we excluded books with the BISAC code FIC049000 but included books with the BISAC code FIC022030. We ensure that the circulations of the control books are comparable by selecting control books whose six-month pre-ban period mean circulation is similar to that of the banned books. This process results in a set of 349 titles that form the control group, which does not contain the same themes or topics as the banned books (i.e., race and gender).

The construction of such a control group ensures that there is limited contamination bias. However, whether this serves as the right counterfactual could be questioned despite similar circulation levels. Hence, we use various other samples of control titles to check the validity of our results. First, we restrict the current sample of control books to those belonging only to those genres that banned books belong to (i.e., fiction and non-fiction) and only to those genres that banned books do not belong to (i.e., business and economics), respectively. Second, we perform a check by automating the selection of control books using a propensity score matching approach (PSM). Finally, we also use books of the same BISAC Subject

Codes as those of the banned books as a control group as part of the additional robustness exercises below. We elaborate on these results and other checks in the analyses below.

4 Baseline Results and Heterogeneity

4.1 Baseline Results

First, we provide some model-free evidence of the impact of book bans on the circulation of treated and control group books. When we compare the normalized average monthly circulation during the pre-ban and the post-ban period between control and the top-25 banned books, we observe that the average circulation of banned books has increased by 14.9% after bans compared to an increase in average circulation by just 0.8% for control books. However, the average effect might be confounded due to seasonality or state-specific factors, among others. Therefore, we employ a Difference-in-Difference (DiD) model to provide a causal estimate of the impact of book ban events on circulations.

In the main results, we first begin by analyzing the entire set of 468 banned books in our data. In column (1) of Table 3, we find a small, marginally statistically significant effect of book bans on readership for the entire sample. We then split the sample into the top-25 titles (column (2)) and the rest (column (3)) to find that the top-25 titles experience a strong positive effect of bans while the rest experience precisely estimated null effects.[25] In column (4), we also present the results for the whole sample but with the number of bans as a continuous interaction variable. As can be seen, the impact on circulation increases with the number of times a book is banned. Given that our main sample of the top-25 banned titles from PEN America and ALA drive our main effects, we focus on this particular sample following the strategy of Agarwal and Sen (2022).

[25]It is pertinent to note that given the aggregate nature of our data, we are not able to identify the impact of book bans at a more micro-geographic level. Indeed, the non-top-25 titles might see a significant effect at a smaller scale. We discuss this issue more in Section 6.

Table 3: All Titles: Impact of Ban Events on Circulation

Dependent Variable:	Log(Circulation)			
	All Titles	Top-25 Titles	Non-Top-25 Titles	Impact Per Ban
	(1)	(2)	(3)	(4)
Variables				
PostBan × Banned	0.016*	0.113***	0.011	-0.009
	(0.009)	(0.017)	(0.009)	(0.009)
PostBan × Banned × Intensity of Ban				0.018***
				(0.002)
State Fixed Effect	Yes	Yes	Yes	Yes
Title Fixed Effect	Yes	Yes	Yes	Yes
Year Month Fixed Effect	Yes	Yes	Yes	Yes
R^2	0.62819	0.57637	0.62452	0.62831
Observations	421,382	249,660	411,122	421,382

Note: Robust standard errors clustered at the level of the state in parentheses. The unit of observation is state-title-month. The dependent variable is the log of circulation. Columns (1) and (4) use all the banned books, column (2) uses only the top-25 most banned books, column (3) uses all the banned books except the top-25. Column (4) uses monthly cumulative count of bans as interaction variable.
$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

We present the TWFE estimates for the top-25 sample in Table 4. Column (1) replicates the result from column (2) of Table 4. On average, based on column (1), the circulation of banned books increases by about 12% (exp(0.113) - 1) compared to the circulation of control books during the 6-month window following the first ban. Column (2) shows a positive and significant average effect, corresponding to the outcomes for banned states - specifically, the 23 states that enacted at least one book ban in the period from July 2021 to June 2022. Interestingly, we see a significant increase in readership for states that never experienced a ban (column (3)).

Note that any library circulation change in percentage or absolute terms is a conservative lower bound when we attempt to understand the economic significance of these estimates. This is because both physical and electronic circulation is limited by existing inventory, limiting how high demand can go. This issue is similar in spirit, for example, to online

Table 4: Impact of Ban Events on Circulation

Dependent Variable:	Log(Circulation)						
	All States	Banned States	Non-banned States	Author Experience	Twitter Visibility	Blue States	Red States
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Variables							
PostBan × Banned	0.113***	0.118***	0.106***	-0.005	0.015	0.119***	0.107***
	(0.017)	(0.022)	(0.029)	(0.020)	(0.021)	(0.023)	(0.026)
PostBan × Banned × Less-experienced Author				0.229***			
				(0.031)			
PostBan × Banned × High Visibility					0.206***		
					(0.036)		
State Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Title Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year Month Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.57637	0.55644	0.61870	0.57647	0.57645	0.52521	0.64069
Observations	249,660	151,110	98,550	249,660	249,660	131,400	118,260

Note: Robust standard errors clustered at the level of the state in parentheses. The unit of observation is state-title-month. The dependent variable is log of circulation. PostBan is an indicator variable which denotes if the time period is after the month following the first time the title is banned in any state. PostBan is 0 for all control titles. Banned indicates whether the title belongs to the banned or control group. Column (1) includes all the states, column (2) includes only those states with at least one ban event, column (3) includes only those states without any ban events, column (6) includes only Red states, and column (7) includes only Blue states. Column (4) includes author experience as interaction variable and column (5) includes Twitter visibility as interaction variable.
$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

rental marketplaces for clothes (Kan et al. 2024) and short-term stays (Zhang et al. 2022). Moreover, these books borrowed from school and public libraries often cater to a broader community, such as a classroom, as noted in Agarwal and Sen (2022). Unfortunately, we cannot report units for the circulation data since it is proprietary information. We can say that post-ban, the increase in circulation itself puts the title in the 90th percentile of most circulated books in our broader dataset. Additionally, based on some publicly available estimates, we can say that the increase in circulation over the six months after the book ban surpasses the lifetime sales of the average book published in the US. We also interpret the estimates' economic significance differently by normalizing circulation numbers by the pre-treatment category level mean and the pre-treatment title level means. The result shown in

column (3) of Table A4 indicates that, on average, the increase in circulation in the post-ban period is 0.32 times (or 32% of) the mean circulation of all the books in their respective categories. Results are shown in column (4) of Table A4. Our results indicate that, on average, the increase in circulation in the post-ban period is 0.46 times (or 46% of) the mean circulation of a title in the pre-ban period.

Figure 4 shows the event study estimates using the TWFE estimator and the three recently developed estimators by Borusyak et al. (2021), Sun and Abraham (2021), and Callaway and Sant'Anna (2021). There are a few clear takeaways. First, we find no significant changes between the treated and control groups in the pre-period, verifying the parallel trends assumption. The parallel trends assumption holds not only for the TWFE estimator but also for three other popular estimators that explicitly account for the staggered nature of the treatment. Second, there is a statistically significant jump in the circulation numbers of banned titles relative to control books after banning a title in any state. Overall, these results give us confidence in the credibility of our results. We present the average estimates for these three alternative methods in columns (1)-(3) of Table 5.[26] In column (4) of Table 5, we also replicate our main estimate with a population-weighted regression, using state-level weights, to find qualitatively and quantitatively similar results.

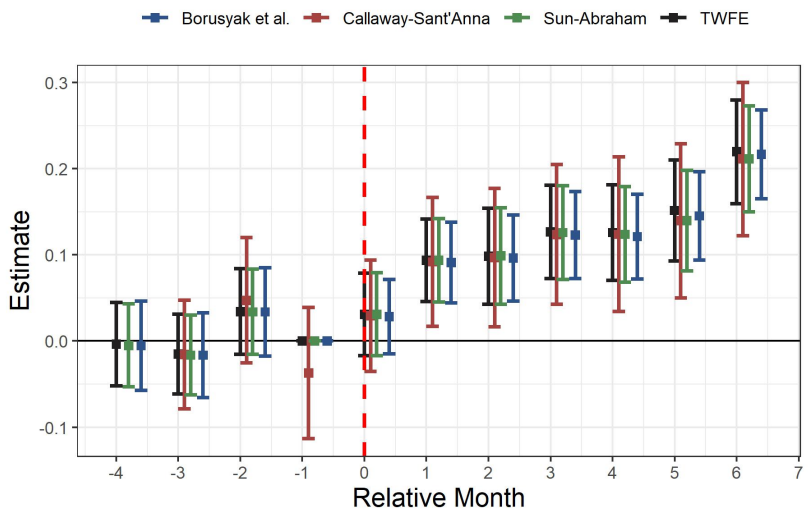
Overall, these results imply that these book bans targeted at individual public libraries or school districts lead to a Streisand effect more generally rather than a potential chilling effect. A caveat is that we do not observe circulation at the school or public library level. Hence, we cannot shed light on the reading behavior in the locations that felt the direct consequences of the ban.

4.2 Robustness Checks

In this section, we explore whether our results are robust on several different dimensions.

[26] Using the definition of an event as the same month when the book ban was initiated leads to qualitatively and quantitatively similar results as seen in column (1) of Table 6.

Figure 4: Event Study Estimates



Note: The figure shows event study estimates using four different estimation methods. The dependent variable is the log of circulation. The unit of observation is at the title-state-month level. 95% confidence intervals are shown.

4.2.1 Alternative Identification and Spillovers

First, our main specification considers a title to be treated as soon as it is banned in any state. However, it is possible that the first significant ban event in our data had spillovers across all eventually banned book titles and states. Therefore, we use October 2021 as the month of treatment for this specification when the first large wave of book bans occurred.[27] In other words, we consider that all the books in our treated category were banned in October 2021, regardless of the state-level bans on individual effects. In column (5) of Table 5, we compare the monthly circulations of the 25 banned titles before and after October 2021 with those of control titles, finding results qualitatively similar to our baseline specification. Next, on the other hand, it is also possible that each state's ban event of each title is truly independent. This condition implies that there are no spillovers across states, i.e., banning a title in State A would not affect readership in State B, where the title is not yet banned. For each state, a book acts as a control book for other banned books until it gets banned in that particular

[27]Panel a of Figure A3 in the Appendix shows that the number of ban events increased substantially after October 2021.

state, and we also use books that never get banned in any state as part of the control group. The result, shown in column (6) of Table 5 is qualitatively similar but higher in magnitude than the results from our main identification specification shown in column (1) of Table 4, with a 34% effect size. This implies that the focal state sees a much stronger effect even though there is an economically meaningful spillover onto non-banned states.

Table 5: Additional Estimators: Impact of Ban Events on Circulation

Dependent Variable:	Log(Circulation)					
	Callaway Sant'Anna (1)	Sun & Abraham (2)	Borusyak et al. (3)	Population Weighted TWFE (4)	Complete Spillover TWFE (5)	No Spillover TWFE (6)
Variables						
PostBan × Banned	0.114***	0.115***	0.115***	0.149***	0.084***	0.300***
	(0.022)	(0.020)	(0.017)	(0.022)	(0.020)	(0.069)
State Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
Title Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
Year Month Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
R^2		0.57673		0.65015	0.59632	0.54692
Observations	249,660	249,660	249,660	249,660	156,750	146,133

Note: Robust standard errors clustered at the level of the state in parentheses. The unit of observation is state-title-month. The dependent variable is the log of circulation. Column (1) uses Callaway Sant'Anna, column (2) uses Sun and Abraham, and column (3) uses Borusyak et al. instead of the TWFE estimator. Column (4) shows average estimates using TWFE estimator with state population weights. Columns (5) and (6) show average estimates for alternate event identification assumptions.
$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

As discussed above, spillovers could be from one banned book to another. In a different setting, Borah and Tellis (2016) show that negative chatter about a product can have a spillover effect and impact the consumption of related products. We use two different measures to test the spillover of book bans on the circulation of related books. First, we test whether there are spillovers to books of similar genres as defined by the same BISAC codes of the banned books. Next, we look at spillovers into other books written by the same author who wrote the banned book. As can be seen from columns (1)-(3) of Table A2 in

the Appendix, we do not observe spillover effects of book bans on circulations of books with similar genres or other books authored by authors of banned books.

4.2.2 Alternative Control Groups

To verify the validity of our results to the selection of alternate control groups, we perform additional checks using multiple different selections of control book titles. First, we restrict our original control group books to only those genres (such as fiction and non-fiction) in which banned books appear. Second, we limit the original control group books to only those genres in which banned books do not appear (such as business and economics). This exercise ensures that our results are not influenced by similarity or dissimilarity to the control group books. Columns (5) and (6) in Table 6 as well as Figure 5 show that the event study estimates are similar to the estimates that use the original control group. Third, instead of using a control group of books with unrelated BISAC Subject Codes, we use books with the same BISAC Subject Codes as those of the banned books to create the control group. This exercise ensures that we compare banned book circulation to the circulation of similar but not-banned books. Additionally, we also use the matching approach to choose control titles with the same BISAC Subject Codes as those of the banned books. Results shown in Columns (7) and (8) of Table 6 illustrate the robustness of the treatment effect to these additional subsets of control group titles.

We provide further evidence of the robustness of our results to the choice of control group books in Table A3 in the Appendix, where we replicate the whole of Table 4, and Figure A4 in the Appendix that replicates Figure 4 using propensity score matching. Instead of limiting the choice set of control books ourselves, we use a matching technique to identify a group of control books from the entire sample of books that are not of the same BISAC Subject Codes as those of the banned books. Using propensity score matching, we match the six-month pre-ban period mean circulation and the genre of control books to that of the banned books and identify a set of 25 control books. This method ensures that the control

Table 6: Additional Robustness: Impact of Ban Events on Circulation

Dependent Variable:	Log(Circulation)							
	Alternate First Ban Month	Alternate Author Experience	Banned Red States	Non-banned Red States	Same Genre Controls	Different Genre Controls	Banned BISAC Controls	Matched Banned BISAC Controls
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Variables								
PostBan × Banned	0.088***	0.030	0.121***	0.071	0.128***	0.082***	0.120***	0.095***
	(0.018)	(0.025)	(0.029)	(0.060)	(0.018)	(0.014)	(0.019)	(0.017)
PostBan × Banned × Less-experienced Author		0.161***						
		(0.031)						
State Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Title Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year Month Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.57618	0.57642	0.64999	0.43352	0.61063	0.64579	0.58258	0.73088
Observations	249,812	249,660	85,410	32,850	101,232	41,724	434,340	34,884

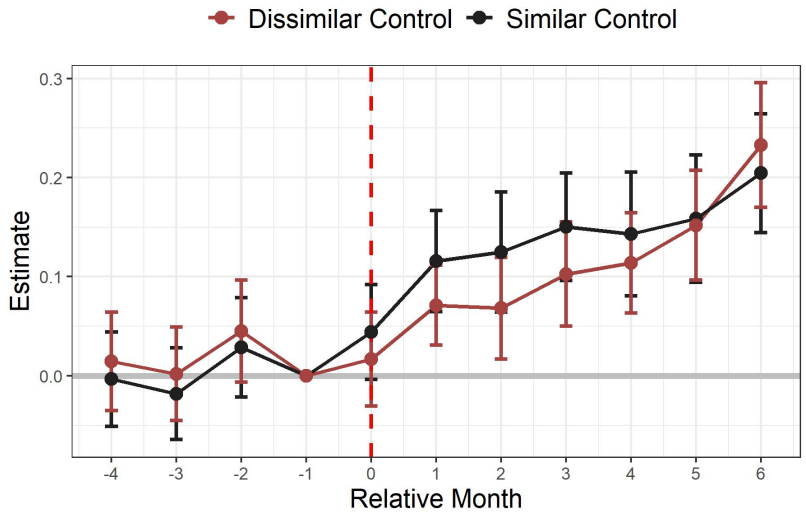
Note: Robust standard errors clustered at the level of the state in parentheses. The unit of observation is state-title-month. The dependent variable is the log of circulation. Column (1) uses the month of first ban as the event instead of the month following the first ban. Column (2) includes alternate measure of author experience as an interaction variable. Columns (3) - (4) show the average treatment effect in Red states with at least one ban event and no ban events respectively. Columns (5) and (6) show the average treatment effect with control book titles with a subset of genres. Columns (7) - (8) show the average treatment effect with control book titles that have the same BISAC Subject Codes as that of banned books. Column (8) uses with control book titles obtained by Propensity Score Matching technique.
$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

group is similar to that of the banned books in terms of pre-period circulation and genre. Our results show that the choice of control titles has a minimal effect on the impact of bans on banned book circulation.

4.2.3 Additional Robustness Checks

We perform several additional robustness checks to verify the validity of our results. To ensure that our results are robust to alternative functional forms of the dependent variable, we reestimate the model using inverse hyperbolic sine transformation of circulations as the dependent variable to find the results qualitatively similar, as seen in column (1) of Table A4 in the Appendix. We find similar results in column (2) when using a negative binomial model specification instead. We also repeat the analysis using quarter- and quarter-state fixed effects to account for the potential increase in state-level reading trends driven by

Figure 5: Impact of Book Bans on Circulation - Alternate Control Books



Note: The figure shows event study estimates using the TWFE estimator with the log of circulation as the dependent variable. It plots the estimates along with the 95% confidence intervals using control group books similar to banned books (in black) and control group books dissimilar (in red) to banned books separately.

organic growth of interest in reading banned books, and we find no change in our estimates (column (3)). To ensure no selection into treatment, we conduct a falsification analysis where we randomly assign the treatment date to each banned title between January and June 2021 to find null results (column (4)). Finally, to verify that the choice of pre-period length does not influence our results, we repeat the analysis using a pre-period of 6 months instead of 4 months to find qualitatively similar results (column (5)).[28] To ensure that our results are not driven by any one particular title, we carried out two sets of leave-one-out analyses. In the first one, we replicate the baseline average effect by leaving out one title at a time. Figure A6 in the Appendix shows that average effects remain statistically significant (at the 5% level) and similar in magnitude to our full sample estimate. Next, we also carry out the same leave-one-out analysis but for the heterogeneity analysis and report the triple interaction coefficients in Figure A7 in the Appendix to find that the results are qualitatively and quantitatively similar to the estimates when all books are included.

[28] We provide the corresponding event study in the Appendix as Figure A5.

4.2.4 Sample Extension and Treatment Intensity

We test the robustness of our results by extending the sample of book bans in Table A5 in the Appendix. Our primary sample consisted of 25 titles, each banned more than five times and representing about 30% of book bans in states. In columns (1) and (3), we extend the sample to include books that are banned at least four times (41 titles accounting for 38% of the bans as shown in panel a of Figure A8). In column (1), based on the main specification, we find that the effect of bans is still positive and statistically significant at the 5% level. The magnitude is smaller than the baseline, which is intuitive given that the impact of these bans is mediated through the publicity effect. In column (3), we use state-title as the ban event to find similar positive and statistically significant effects, but with magnitudes at 20% that are lower than when we analyze the 25 banned titles. We extend this sample further in columns (2) and (4) by looking at 71 titles that are banned at least three times, accounting for 46% of all ban events, as shown in panel b of Figure A8 in the Appendix. Similar to the previous analysis, we find that the effects remain positive and statistically significant at the 5% level but with smaller magnitudes. This analysis suggests that the treatment intensity, captured by the number of times a book is banned, significantly affects readership. In column (5), we capture this directly by examining how circulation is affected by each additional ban to find that each additional ban increases the readership by 2%. In column (6), we use the entire sample of 468 banned books available in our data to measure the compounding effect of bans. The interaction variable '>= Five Bans' indicates whether the book has been banned five or more times. The results indicate that the impact of bans is statistically significant only for the top-25 books that are part of 85% of the ban events and have higher salience on media. Overall, these results suggest that our estimates are robust and that our approach is with minimal loss of generality in the overall implications for stakeholders. This will also be evidenced by the heterogeneity analysis that follows. Moreover, as mentioned earlier, focusing on salient events is in line with the general approach taken in the literature (Liaukonytė et al. 2023, Wang and Lu 2020).

4.3 Heterogeneous Impact of Book Bans

We look at several dimensions of heterogeneity to understand the potential mechanisms driving the baseline results above.

4.3.1 Author Experience

First, we try to understand whether the average effects vary with the author's experience. Studies in the literature (Petrova et al. 2021) suggest that media coverage could have an information or persuasive effect. In particular, Petrova et al. (2021) show that (social) media helps new players on the market because of the information effect. To construct a measure of experience, we use the total number of books the author wrote before the book ban, excluding the banned book. We categorize the authors in our treated sample as more experienced and less experienced authors based on the median count of books and estimate Equation 1 with author experience as an interaction variable. The results, shown in column (4) of Table 4, reveal that titles by less experienced authors drive the impact of book bans on circulation. Figure 6a) shows the event study estimates using sub-sample analyses. We see no impact of book bans on books written by more experienced authors, whereas there is a significant increase in the circulation of books by less experienced authors following bans. We also use social media visibility as an alternate measure of author experience. We count all tweets that mention the names of banned authors between January 2012 and December 2020. We then split the sample of authors into more-experienced and less-experienced authors based on the tweet counts. Our results, shown in column (2) in Table 6, are consistent with the measure of the author's experience from Goodreads data.

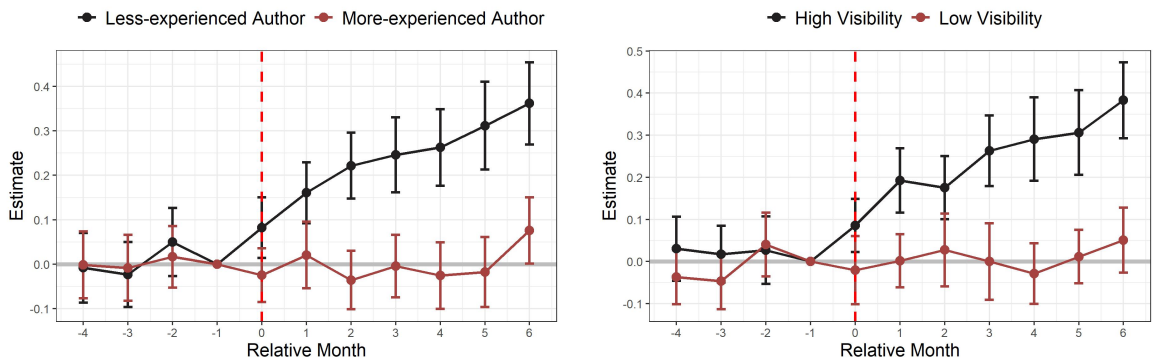
4.3.2 Social Media Visibility

Social media has played a key role in cultural debates and in driving local and wider national conversations around book bans. To understand the potential role of Twitter in explaining our findings, we repeat the main analyses using visibility on Twitter as an interaction vari-

able. Results in column (5) of Table 4 show that the effect of book bans is driven significantly by titles with high visibility (titles with more than the median number of impressions) on Twitter. Figure 6b) shows the estimate of the event study using sub-sample analyses. Books with low Twitter visibility did not see an impact of ban events, while books with high Twitter visibility saw a significant increase in circulation following bans.

These results on social media chatter, along with those on author experience, suggest that conversations about inexperienced authors can provide information and impact consumer decisions. The fact that even within these 25 books that were banned more than 300 times, the only books that see an increase in circulation are those that are discussed on social media highlights the critical role of new media in the context of polarizing issues. It highlights that the critical mechanism that leads to any impact of such book bans is the spread of information; otherwise, it has a null effect on consumer behavior.

Figure 6: Heterogeneous Impact of Book Bans on Circulation



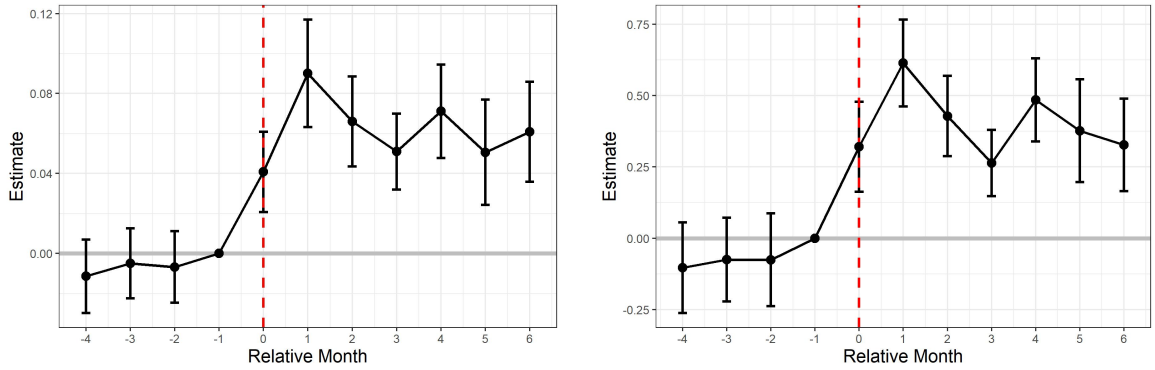
a) By Author Experience b) By Twitter Visibility

Note: The figure in panel a) shows event study estimates using the TWFE estimator. It plots the estimates along with the 95% confidence intervals for banned book titles by less-experienced (in black) and more-experienced authors (in red) separately. The figure in panel b) plots the estimates along with the 95% confidence intervals for banned book titles with high (in black) and low Twitter visibility (in red) separately.

Next, we carry out a check to rule out reverse causality with social media chatter leading to book bans in the first place. In particular, we estimate the same regression model as in Equation 1 but use Twitter engagement (tweet counts and impressions) as the dependent variable instead of circulation numbers. The results in Figure 7 demonstrate two clear points.

First, there is no significant pre-trend in social media chatter or engagement between the banned and the control books prior to the ban. Second, after the ban, Twitter chatter and engagement with banned books increased significantly relative to the control books. Given the lack of pre-trends, it is unlikely that increased social media mentions leading to the book bans.[29]

Figure 7: Event Study Estimates for Twitter Chatter



a) Tweet Counts b) Tweet Impressions

Note: The figure in panel a) shows the TWFE event study estimates with log of tweet counts as the dependent variable. The figure in panel b) shows TWFE event study estimates with log of tweet impressions as the dependent variable. 95% confidence intervals are shown.

To dig further into the mechanism behind the increase in circulation, we analyzed the news coverage of each of the 142 district-month level book ban events to identify whether a national or local news outlet first covered the story. To do so, we obtained the top 100 articles on a Google News search with the school district name, ban timing, and state name. We manually verified and classified all the news stories as local or national outlets/organizations (such as PEN America, ALA, ACLU, or NCAC). In only five cases was the first article about the ban event written by a national news outlet (e.g. CNN, WSJ, NYT, USA Today

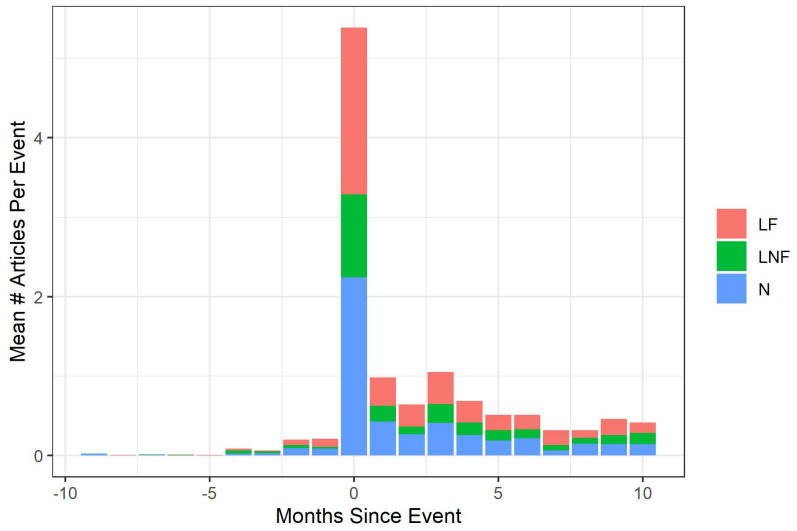
[29]There can be a quicker change in the volume of online chatter associated with book bans relative to the change in reading behavior, and thus, it is possible that we can see social media engagement increasing immediately after the ban, whereas it takes time to show up in library circulations. PEN America database only has the month of the ban but not the exact date. We hand-collected data on the exact ban dates, extracting information from online blogs, parent forums, and media outlets for about 85% of book titles in our main list. While a noisy measure, it provides a starting point for us. We then manually inspected tweets within seven days before and after the ban date. We find that there were no ban-related tweets in the 7 days leading up to the ban, and they occurred only after the ban, with only a handful of book bans being discussed online in the first seven days.

etc.). In all other cases, the story was *first* covered by a local news outlet. There was also a distinct difference in how the local and national news reported the stories. Specifically, local news outlets focused on the *particular district or library* where the ban was enacted and interviewed local students, parents, and educators. However, national news outlets talked about book bans more generally as a rising national trend and mentioned the bans by combining several bans as examples of these trends. We show the news coverage of bans descriptively month-on-month around the time of the ban in Figure 8 to see a spike in news coverage after a ban was in place, with the bulk of the coverage driven by local news outlets. We break this coverage down by differentiating between local news outlets of the focal state where the ban occurred (for example, a local newspaper covering a county in Pennsylvania talking about a book ban in that county it serves), national news coverage, and non-focal local news outlets (these refer to local newspapers that refer to book bans happening outside their state). Figure 8 also shows that local news outlets in the focal state are more likely to cover these events than other local outlets and are equally likely relative to national news outlets. These numbers highlight the emphasis provided by local outlets in the focal state and align with the narrative of the initial buzz being generated locally but then being picked up more generally, leading to a larger effect in the focal state but also positive spillovers onto other states that didn't see a ban. This finding is also intuitive because book bans originate in local school board meetings and are very much a local issue. Further, this points to the fact that book bans become interesting nationally only when enough of them appear as a rising trend worth reporting.[30]

Finally, to provide additional evidence on the role of social media chatter, we also analyze data from Goodreads. We find that after ban events, there is an increase in the number of reviews written, the number of reviews mentioning book bans explicitly, and the average rating of the banned books relative to the control. We provide details of these results in Table A6 in the Appendix.

[30]This evidence is suggestive since these different pieces of coverage and chatter co-occur; we cannot quantitatively separate the relative importance of each channel.

Figure 8: Number of News Articles



Note: The figure shows the mean number of news articles covering a ban event each month relative to the month of the event. The news articles are classified as local newspapers covering a focal state event (LF), local newspapers covering a non-focal state event (LNF), and national newspapers covering the event (N).

4.3.3 Blue vs. Red States

Books bans have become a prominent part of politically polarizing 'culture wars'. Politicians have actively used these issues as part of election campaigns. Therefore, there could be a political dimension to who reacts to these book bans. Using the electoral outcomes of the 2020 presidential election, we classify each of the 38 states as Blue (if the majority vote in the state was for Joe Biden in the 2020 elections) or Red (if the majority vote in the state was for Donald Trump in the 2020 elections). Overall, our data has 20 Blue states and 18 Red states. Given that our circulation data is at the title-state-month level, we use state-level measures of political preferences for this analysis. Columns (6) and (7) in Table 4 show that the impact of book bans can be seen both in the Blue and Red states, respectively. The average effect is 11% higher in the Blue states. In columns (3) and (4) of Table 6, we show that for the Red states, the circulations of banned books have increased only in the states with book bans. However, among Blue states, the circulation of banned books has increased in states with and without bans.

4.4 Impact on Children's Reading Behavior

Our primary data do not distinguish reading behavior between adults and children. To provide evidence that the increased consumption might also include children reading more banned books, we use de-identified patron-level checkout data from the Seattle Public Library (SPL). This data provides unique insight into children's reading behavior in the aftermath of book bans in different parts of the United States.

We aggregate the data at the patron-title-month level in our analyses. Out of the 25 banned books from our sample, we matched 24 titles from the SPL checkout data. To create a control group of titles, we remove books dealing with subjects similar to the banned titles and whose pre-ban circulation trends are similar to those of the banned books. Overall, we have 24 banned titles and 76 control titles. Our consultation with the SPL revealed that parents would often check out books on behalf of their children. We use only patrons who are identified as children, which ensures that our results are conservative.

We estimate the causal impact of book bans on children's reading behavior using the staggered DiD model shown in Equation 1 where y_{it} refers to the circulation count of title i at month t. Since we use data that only pertain to the SPL, we drop the subscript j. Column (1) in Table 7 shows that, on average, children read banned books about 19% more than the control books following book ban events. These magnitudes are in the ballpark of the magnitudes we find across different specifications in the aggregate analysis. We further classify the patrons as inclusive content readers and inclusive content non-readers based on whether they read any books dealing with race and/or gender identity in the pre-ban period (between Jan 2021 and Jun 2021). Our results, shown in columns (2) and (3), indicate that the increase in banned book circulation is driven by patrons who have not yet been exposed to inclusive content. Our results suggest that book bans expose "new" readers among children to inclusive subject matter in libraries, furthering our belief that book bans end up increasing the salience of the books that they aim to ban. We also use an alternate

identification specification, based on the single event difference in differences, to highlight the robustness of our results, as shown in column (4).

Overall, these results suggest that children began to read more of these banned books after they were banned in different parts of the United States. This analysis sheds some light on the mechanisms behind the aggregate results presented previously. These results come with the caveat that the data are based on patron-level data from one city's public library system. However, we believe that these data provide unique insight into the reading behavior of children that are hard to come by otherwise.[31]

Table 7: Impact of Ban Events on Circulation

Dependent Variable:	Log(Circulation)			
	All Readers	Inclusive Readers	Inclusive Non-Readers	Alternate Identification
	(1)	(2)	(3)	(4)
Variables				
PostBan × Banned	0.189***	0.061	0.166***	0.279***
	(0.059)	(0.055)	(0.054)	(0.066)
Title Fixed Effect	Yes	Yes	Yes	Yes
Year Month Fixed Effect	Yes	Yes	Yes	Yes
R^2	0.46983	0.40276	0.34480	0.51134
Observations	2,088	2,088	2,088	1,100

Note: Heteroskedasticity-robust standard errors in parentheses. The unit of observation is title-month. The dependent variable is log of circulation. PostBan is an indicator variable which denotes if the time period is after the month following the first time the title is banned in any state. PostBan is 0 for all control titles. Column (1) includes all the readers, column (2) includes only those readers who checked-out inclusive titles, and column (3) includes only those readers who did not check-out inclusive titles. Column (4) shows average estimates for an alternate event identification assumption.
$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

[31] Additionally, as mentioned previously, in Table 2 (column (3)), we show that the 25 titles in our sample are shelved by Goodreads readers on school-related shelves significantly more than the rest of the banned books. Unfortunately, we cannot carry out the same analysis using event studies because we don't have time stamps for the shelving action on Goodreads.

4.5 Alternate Measure of Consumption: Amazon Sales Rank Data

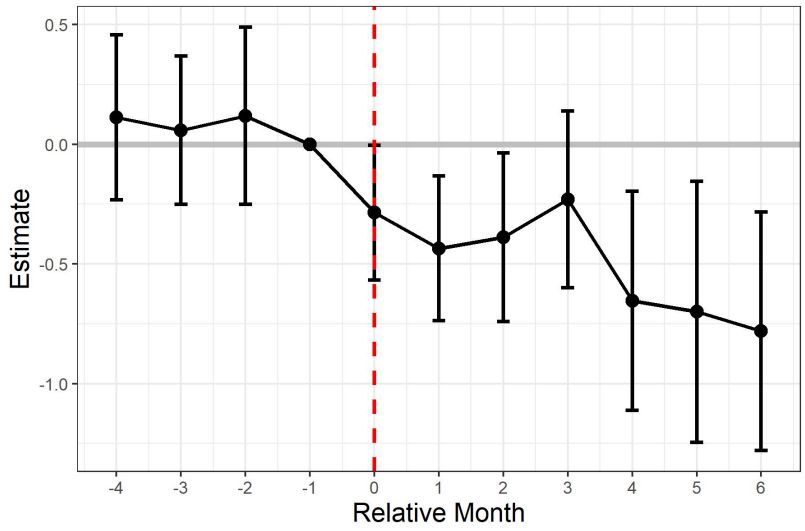
We use Amazon sales rank data from Keepa for each book every month as an alternate book consumption measure to add credence to our results. Such sales rank data has been widely used in the literature previously (Reimers and Waldfogel 2021) and from Keepa in particular, for example by Socal et al. (2024). We run the same difference in differences estimation but with log sales rank as the outcome variable. The results reported in Figure 9 show parallel trends in the pre-period and then a reduction in the sales rank of the treated books relative to the control group. There is a 41% (exp(-0.53) - 1) decline in sales rank, implying a significant increase in sales. While it is not straightforward to interpret these sales rank changes because of potential non-linearities in the sales and sales rank link, we can benchmark against other sales rank estimates in the literature. For example, Reimers and Waldfogel (2021) find that a book's Amazon sales rank improves by 78% in the five days after a New York Times review is published. We believe that our magnitudes are economically significant relative to this benchmark. We do not have sales data on quantities to convert these sales-rank changes into unit changes. However, we can use some publicly available tools[32] to convert Amazon sales rank changes into absolute quantities. Using this, we find that the number of copies sold per month would have increased from 90 to 360 on Amazon alone. We report these magnitudes with the caveat that we do not know the underlying model this company uses to translate sales rank to units.

5 Book Bans and Political Rhetoric

The above analysis shows that book bans targeted at micro-geographic levels can lead to media attention on those particular books and, in turn, increase readership. We also find that this readership can increase in both Democratic and Republican states. As mentioned previously, politicians have used the recent surge in book ban attempts to gain an electoral

[32] `https://bookbeam.io/amazon-book-sales-calculator/`

Figure 9: Impact of Book Bans on Amazon Sales Rank



Note: The figure shows the TWFE event study estimates with log of Amazon Sales Rank as the variable of interest. 95% confidence intervals are shown.

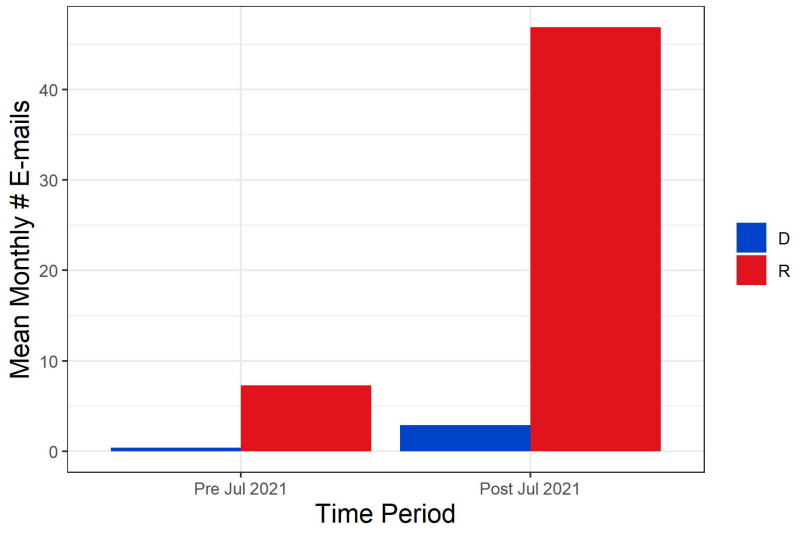
advantage over their competitors.[33] In this Section, we attempt to dig deeper into the political dimension of these bans and quantify whether politicians gain in any way from these issues. In particular, similar to Petrova et al. (2021), we analyze how these book bans feature in political rhetoric and might be associated with political donations received by politicians.

First, we examine how book bans feature in political candidates' communication strategies. We do this by analyzing the text of emails sent to potential donors who subscribe to campaigns of political candidates and committees. We identify book ban-related emails using keywords such as "banned books", "book challenges", "book censorship", etc. Figure 10 shows that before the first book bans in our sample, a small number of emails from Republican candidates mentioned any book ban-related topic. After the first book bans in our sample, there was a significant increase in emails sent out by Republican candidates and committees referring to book bans.[34] Furthermore, as shown in the figure, this increase is seen mainly for Republican candidates and committees rather than their Democratic coun-

[33] https://www.washingtonpost.com/education/2022/02/10/book-bans-maus-bluest-eye/

[34] The first book ban event took place in July 2021 in Wisconsin and Indiana.

Figure 10: Mean Monthly Emails containing Ban Related Keywords



Note: The figure shows the average number of campaign emails sent per month by Democratic and Republican party candidates before and after the first book ban event.

terparts. Overall, more than 90% of the emails were from politicians in the Republican party, even though bans occurred in the Red and Blue states. A caveat in interpreting these results is that we have a selected sample of 245 political accounts, and hence, generalizations should be made with caution.

Next, we want to understand whether politicians enjoy any political benefit from these book bans. In particular, we want to examine whether such events benefit Democrats or Republicans, if at all. To identify the potential association of book ban events with political donations, we compare political contributions to Democratic and Republican House candidates before and after the first book ban in each state to which the candidate belongs. In line with Petrova et al. (2021), our analysis focuses on small donations potentially influenced by these communications, specifically examining contributions under $500 each.

The results shown in Table 8 provide a nuanced picture. In column (1), we see no differential association of ban events with donations received by Republican politicians relative to Democrats. In column (2), focusing on the Blue States (based on their vote for Joe Biden), we find that this null effect in the difference between the Republican and Democratic do-

nations persists. Column (3) provides suggestive evidence in line with a concerted political strategy. In particular, in the Red States, Republican candidates get a significant increase in political donations relative to their Democrat counterparts. Economically, the estimates suggest a 30% increase in donations under $500 for Republican politicians relative to Democrat candidates. In absolute terms, the average Republican gets $800 (mean is $2668) more every month after book ban events.

Table 8: Impact of Ban Events on Political Donations

Dependent Variable:	Log(Donations)		
	All States	Blue States	Red States
	(1)	(2)	(3)
Variables			
PostBan × Republican	0.084	-0.155	0.305**
	(0.201)	(0.360)	(0.132)
State Fixed Effect	Yes	Yes	Yes
Candidate Fixed Effect	Yes	Yes	Yes
Year-Month Fixed Effect	Yes	Yes	Yes
State x Year-Month Fixed Effect	Yes	Yes	Yes
R^2	0.68125	0.67531	0.68087
Observations	16,405	8,012	8,393

Note: Robust standard errors clustered at the level of the state in parentheses. The unit of observation is candidate-month. The dependent variable is log of the donation amount. Column (1) uses all the banned states, columns (2) and (3) use only Blue and Red banned states respectively.
$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

Overall, this analysis provides suggestive evidence for why some Republican politicians might favor book bans. While such book bans increase circulation, Republican politicians in the Red States that see bans seem to gain donations relative to Democratic candidates. A caveat is that this analysis does not entirely establish causality and should be seen as descriptive and as a first step in understanding the politics around book bans.

6 Discussion and Conclusions

Although book bans have been occurring at increased frequencies in the past couple of years, there is limited causal evidence of their impact. In this study, we present the first evidence of book ban events on book circulations, as well as their political implications. Using novel data on book circulation, we show that book bans increase the circulation of banned books by 12% compared to other books. This effect is persistent in states with bans and states without bans, though the effect is larger in focal states experiencing bans. Author experience and social media visibility play a significant role in driving this impact - titles by less-experienced authors and titles with high visibility on Twitter see significantly higher circulations following bans. Finally, Republican party candidates use these contentious book bans as part of their political rhetoric in emails to donors, which is associated with an increase in small donations, particularly in the Red States.

Our results have significant managerial and policy implications. Firstly, we highlight the impact of politically motivated censorship on consumers' consumption behavior. Despite the increasing incidence of book bans in schools and public libraries, bans create a "Streisand" effect on consumption, where politically motivated consumers increase the consumption of banned books. This finding has takeaways for content creators and publishers. This has long-term implications because micro-level bans may create a chilling effect among creators who produce diverse content, which might be detrimental to consumers of such content by furthering the underrepresentation of such content. Our results suggest that such bans at the micro-level may not hinder overall demand. We also highlight the importance of social media visibility in driving political conversations around sensitive issues. Stakeholders should be aware of the potential of social media to magnify political beliefs and thereby influence consumption decisions on a large scale. Finally, stakeholders should also be mindful of the feedback loop, where pursuing contentious policies can lead to more polarized rhetoric through campaign emails, leading to an increase in financial support for politicians, which can further incentivize pursuing contentious policies.

Our empirical results come with caveats regarding external validity. First, we focus on the 25 most banned titles that account for 127 titles × state events that are the most salient events. For other titles, we find a null result. Hence, the main analysis might only apply to high-profile titles that get significant attention on social and traditional media which might be a different treatment relative to less prominent events. Next, our political emails analysis is not based on comprehensive data, but on 245 accounts for which it is curated as a public good. Going forward, it could be important to analyze the digital messaging strategy of a more representative set of politicians in the wake of such book bans. Finally, we do not have comprehensive data on children's reading behavior because of privacy regulations and need to rely on data from one public library system along with inferred data from Goodreads.

More broadly, our research also opens avenues for further research. First, while we observe an increase in circulations after bans in the short term, the long-term implications remain uncertain. With increased political polarization, there is the possibility that librarians and educators, anticipating possible future bans, may choose not to stock or assign such books.[35] Furthermore, we do not measure the impact on beliefs and attitudes of adults and children that might evolve slowly and see a change in the longer term. Second, the book bans in our sample happen at the public library or school district level, and our circulation data is more aggregated. A more granular approach, perhaps within specific demographic segments, could yield insight into nuanced patterns of consumption and censorship. Indeed, if more granular data becomes available then it could provide insights into the impact of bans on less prominent book titles that aren't banned often and don't gain widespread coverage. Moreover, there is a possibility that policymakers could establish legal measures to institute more widespread book bans, consequently restricting the availability of contentious titles on a broader scale, posing a significant concern. Such an expansion of censorship efforts could dramatically alter the availability and accessibility of contentious titles, thus reshaping the

[35] https://www.technologyreview.com/2022/07/15/1055959/book-bans-social-media-harassment/

landscape of public discourse and knowledge dissemination. We hope that future research can use our results as a basis for further exploration of these open issues.

7 Funding and Competing Interests

The authors declare that they do not have any funding from or competing interests with any of the stakeholders discussed in the paper.

References

Ackerberg, Daniel A. 2001. Empirically distinguishing informative and prestige effects of advertising. *RAND Journal of Economics* 316–333.

Agarwal, Saharsh, Ananya Sen. 2022. Antiracist curriculum and digital platforms: Evidence from black lives matter. *Management Science* **68**(4) 2932–2948.

Ahluwalia, Rohini, Robert E Burnkrant, H Rao Unnava. 2000. Consumer response to negative publicity: The moderating role of commitment. *Journal of marketing research* **37**(2) 203–214.

Archak, Nikolay, Anindya Ghose, Panagiotis G Ipeirotis. 2011. Deriving the pricing power of product features by mining consumer reviews. *Management science* **57**(8) 1485–1509.

Berger, Jonah, Alan T Sorensen, Scott J Rasmussen. 2010. Positive effects of negative publicity: When negative reviews increase sales. *Marketing science* **29**(5) 815–827.

Berkes, Enrico, Peter Nencka. 2021. Knowledge access: The effects of carnegie libraries on innovation. *Available at SSRN 3629299* .

Borah, Abhishek, Gerard J Tellis. 2016. Halo (spillover) effects in social media: do product recalls of one brand hurt or help rival brands? *Journal of Marketing Research* **53**(2) 143–160.

Borusyak, Kirill, Xavier Jaravel, Jann Spiess. 2021. Revisiting event study designs: Robust and efficient estimation. *arXiv preprint arXiv:2108.12419* .

Bursztyn, Leonardo, Jonathan T Kolstad, Aakaash Rao, Pietro Tebaldi, Noam Yuchtman. 2022. Political adverse selection. Tech. rep., National Bureau of Economic Research.

Callaway, Brantly, Pedro HC Sant'Anna. 2021. Difference-in-differences with multiple time periods. *Journal of Econometrics* **225**(2) 200–230.

Chevalier, Judith A, Dina Mayzlin. 2006. The effect of word of mouth on sales: Online book reviews. *Journal of marketing research* **43**(3) 345–354.

Foerstel, Herbert N. 2002. *Banned in the USA: A reference guide to book censorship in schools and public libraries*. Greenwood Publishing Group.

Fujiwara, Thomas, Karsten Müller, Carlo Schwarz. 2021. The effect of social media on elections: Evidence from the united states. Tech. rep., National Bureau of Economic Research.

Gavigan, Karen, Shana Pribesh, Gail Dickinson. 2010. Fixed or flexible schedule? schedule impacts and school library circulation. *Library & Information Science Research* **32**(2) 131–137.

Gilpin, Gregory, Ezra Karger, Peter Nencka. 2021. The returns to public library investment gregory gilpin, ezra karger, and peter nencka .

Goncalves, Marcelo SO, Isabelle Langrock, Jack LaViolette, Katie Spoon. 2024. Book bans in political context: Evidence from us schools. *PNAS nexus* **3**(6).

Gordon, Brett R, Wesley R Hartmann. 2013. Advertising effects in presidential elections. *Marketing Science* **32**(1) 19–35.

Gordon, Brett R, Wesley R Hartmann. 2016. Advertising competition in presidential elections. *Quantitative Marketing and Economics* **14** 1–40.

Jansen, Sue Curry, Brian Martin. 2015. The streisand effect and censorship backfire .

Kan, Yu, Ye Liu, Uttara M Ananthakrishnan, Yong Tan. 2024. One size fits all? informational accessibility and inclusivity in online platforms. *Informational Accessibility and Inclusivity in Online Platforms (May 28, 2024)* .

Karger, Ezra. 2021. The long-run effect of public libraries on children: Evidence from the early 1900s. *SocArXiv. Doi* **10**.

Kempf, Elisabeth, Margarita Tsoutsoura. 2021. Partisan professionals: Evidence from credit rating analysts. *The journal of finance* **76**(6) 2805–2856.

Kovács, Balázs, Amanda J Sharkey. 2014. The paradox of publicity: How awards can negatively affect the evaluation of quality. *Administrative science quarterly* **59**(1) 1–33.

Liaukonytė, Jūra, Anna Tuchman, Xinrong Zhu. 2023. Frontiers: spilling the beans on political consumerism: do social media boycotts and buycotts translate to real sales impact? *Marketing Science* **42**(1) 11–25.

Luca, Michael. 2016. Reviews, reputation, and revenue: The case of yelp. com. *Com (March 15, 2016). Harvard Business School NOM Unit Working Paper* (12-016).

Mach, Martin. 2022. Streisand effect in the context of the right to be forgotten. *European Studies* **9**(1) 110–121.

Marthews, Alex, Catherine E Tucker. 2017. Government surveillance and internet search behavior. *Available at SSRN 2412564* .

Mathur, Arunesh, Angelina Wang, Carsten Schwemmer, Maia Hamin, Brandon M Stewart, Arvind Narayanan. 2023. Manipulative tactics are the norm in political emails: Evidence from 300k emails from the 2020 us election cycle. *Big Data & Society* **10**(1) 20539517221145371.

Mumma, Kirsten Slungaard. 2022. Politics and children's books: Evidence from school library collections.

Petrova, Maria, Ananya Sen, Pinar Yildirim. 2021. Social media and political contributions: The impact of new technology on political competition. *Management Science* **67**(5) 2997–3021.

Reimers, Imke, Joel Waldfogel. 2021. Digitization and pre-purchase information: the causal and welfare impacts of reviews and crowd ratings. *American Economic Review* **111**(6) 1944–1971.

Reimers, Imke C, Joel Waldfogel. 2022. The first sale doctrine and the digital challenge to public libraries. Tech. rep., National Bureau of Economic Research.

Schoenmueller, Verena, Oded Netzer, Florian Stahl. 2023. Frontiers: Polarized america: From political polarization to preference polarization. *Marketing Science* **42**(1) 48–60.

Shipilov, Andrew V, Henrich R Greve, Timothy J Rowley. 2019. Is all publicity good publicity? the impact of direct and indirect media pressure on the adoption of governance practices. *Strategic Management Journal* **40**(9) 1368–1393.

Soberman, David, Loic Sadoulet. 2007. Campaign spending limits and political advertising. *Management Science* **53**(10) 1521–1532.

Socal, Mariana P, Yuchen Wang, Elizabeth Plummer, Gerard F Anderson, Ge Bai. 2024. Spending on dual over-the-counter and prescription drugs in the medicare part d program. *JAMA* .

Stolle, Dietlind, Marc Hooghe, Michele Micheletti. 2005. Politics in the supermarket: Political consumerism as a form of political participation. *International political science review* **26**(3) 245–269.

Stolle, Dietlind, Michele Micheletti. 2013. *Political consumerism: Global responsibility in action*. Cambridge University Press.

Sun, Liyang, Sarah Abraham. 2021. Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics* **225**(2) 175–199.

Wang, Hongchang, Eric Overby. 2022. Do political differences inhibit market transactions? an investigation in the context of online lending. *Management Science* .

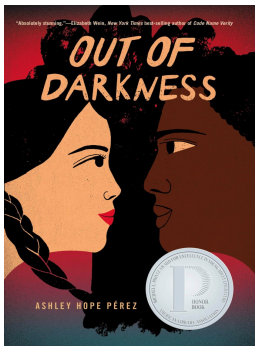
Wang, Kitty, Shijie Lu. 2020. Corporate political positioning and sales: Evidence from a natural experiment. *Available at SSRN 4084106* .

Zhang, Shunyuan, Dokyun Lee, Param Vir Singh, Kannan Srinivasan. 2022. What makes a good image? airbnb demand analytics leveraging interpretable image features. *Management Science* **68**(8) 5644–5666.

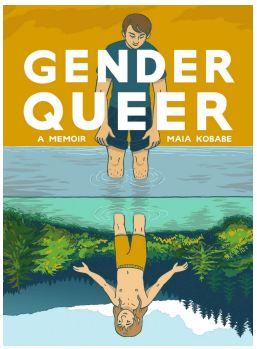
Zhu, Kai, Qiaoni Shi, Shrabastee Banerjee. 2023. Unintended consequences of platform monetization on digital cultural markets: Evidence from a natural experiment on goodreads. *Available at SSRN 4440563* .

Appendix

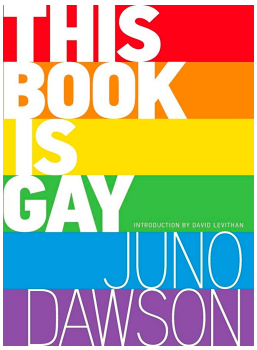
Figure A1: Sample of Banned Books



a) Dealing with Race



b) Dealing with Gender



c) Dealing with Sexuality

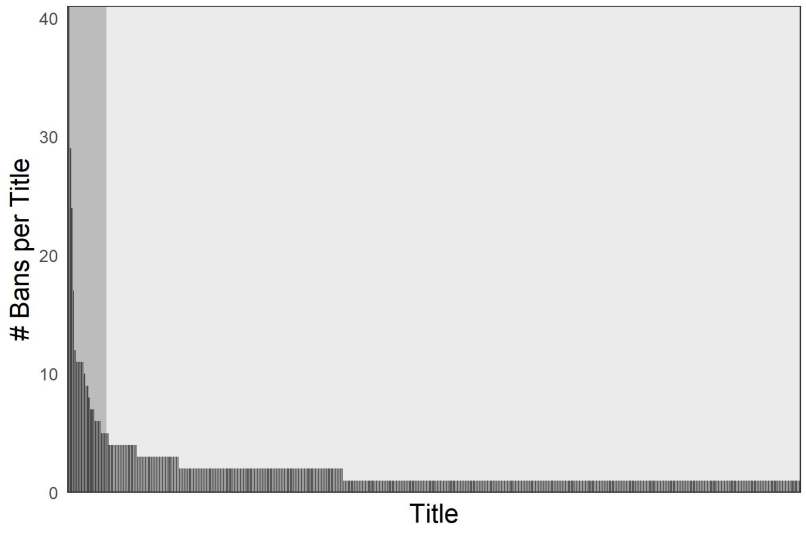
Note: The figure shows examples of banned books dealing with Race, Gender, and Sexuality respectively.

Table A1: List of Top-25 Banned Book Titles

Title	Genre	Primary Author	Author Gender	Author Race
All American Boys	Young Adult Fiction	Jason Reynolds	Male	Black
All Boys aren't Blue	Young Adult Non-Fiction	George M. Johnson	Non-Binary	Black
Almost Perfect	Young Adult Fiction	Brian Katcher	Male	White
Being Jazz	Young Adult Non-Fiction	Jazz Jennings	Transgender	White
Beyond Magenta	Young Adult Non-Fiction	Susan Kuklin	Female	White
Dear Martin	Young Adult Fiction	Nic Stone	Female	Black
Drama	Juvenile Fiction	Raina Telgemeier	Female	White
Flamer	Young Adult Fiction	Mike Curato	Male	Mixed-Race
Gender Queer	Comics & Graphic Novels	Maia Kobabe	Non-Binary	White
I Am Jazz	Juvenile Non-Fiction	Jazz Jennings	Transgender	White
l8r, g8r	Young Adult Fiction	Lauren Myracle	Female	White
Lawn Boy	Fiction	Jonathan Evison	Male	White
Melissa	Juvenile Fiction	Alex Gino	Genderqueer	White
Monday's Not Coming	Young Adult Fiction	Tiffany D. Jackson	Female	Black
More Happy Than Not	Young Adult Fiction	Adam Silvera	Male	Hispanic
Out of Darkness	Young Adult Fiction	Ashley Hope Perez	Female	White
Real Live Boyfriends	Young Adult Fiction	E. Lockhart	Female	White
The 57 Bus	Young Adult Non-Fiction	Dashka Slater	Female	White
The Breakaways	Juvenile Fiction	Cathy G. Johnson	Non-Binary	White
The Infinite Moment of Us	Young Adult Fiction	Lauren Myracle	Female	White
The Truth About Alice	Young Adult Fiction	Jennifer Mathieu	Female	Hispanic
This Book is Gay	Young Adult Non-Fiction	Juno Dawson	Trans Woman	White
This One Summer	Young Adult Fiction	Mariko Tamaki	Female	Asian
Two Boys Kissing	Young Adult Fiction	David Levithan	Male	Wihte
We Are the Ants	Young Adult Fiction	Shaun David Hutchinson	Male	White

Note: List of all the 25 banned book titles included in our study along with genre, name of the primary author, their gender and race.

Figure A2: Histogram of Number of Bans per Title

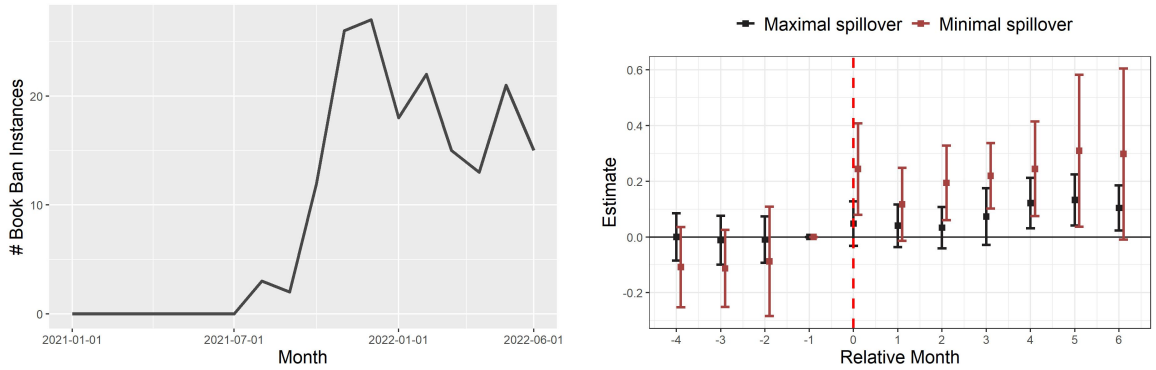


Note: The figure shows the distribution of the number of total ban events per title arranged by the decreasing count of ban events. The area marked represents the sample of banned books considered in our analyses.

Additional Details on Book Bans

Figure A2 demonstrates the number of times each title is banned in different school districts, with the X axis denoting each title and the Y axis denoting the number of bans. We observe that the distribution is heavily skewed to the left and has a significant long tail of titles that are banned only once. This indicates that some titles are banned more frequently than others. In fact, only 25 banned books contribute to more than 300 of the 940 challenges (30%) and are banned at least five times by different school districts. Books that are banned more than four times each contribute to 38% while books that are banned over three times contribute to 46.5% of the distribution. In summary, only 71 titles are banned at least three times. In our analysis, we focus primarily on the 25 books that face the most challenges from school districts and are shown in the shaded part in Figure A2. Our results remain qualitatively similar to this baseline if we consider all 71 titles that are banned at least three times.

Figure A3: Event Study Estimates using Alternate Event Definitions



a) Number of Unique Bans b) Event Study Estimates

Note: The figure in panel a) shows the total number of unique bans per calendar month. The figure in panel b) shows event study estimates using the TWFE estimator with the log of circulation as the dependent variable. It plots the estimates along with the 95% confidence intervals using two alternate ban event definitions - a single ban event for all the banned books occurring during October 2021 and a separate ban event for each title in each state the title gets banned.

Additional Measure of Social Chatter: Goodreads Data

We use Goodreads data as an additional measure of social chatter. On Goodreads, an avid community of readers interact by reading, shelving, sharing, and reviewing books. Activity on Goodreads pre vs. post book ban events could help add credence to our results. We run the regression from Equation 1, but since the Goodreads data do not have state-level information, the subscript j would not apply to this analysis. We use the total number of monthly reviews as an alternative measure of demand.[36] As shown in column (1) of Table A6, there is a positive and statistically significant increase in the number of reviews written for banned books. Next, we analyze the total number of reviews that refer to words related to book bans[37] to find that there is a positive and statistically significant (in column (2)) suggesting that the readers were influenced by news related to book ban. We also look at the average monthly rating for each banned book on Goodreads relative to the control group. In

[36]There is an established relationship between the number of reviews and book demand (Chevalier and Mayzlin 2006, Archak et al. 2011))

[37]We perform a keyword search on the reviews for the terms related to book bans such as "ban", "banned", "bans", "censorship", "censor", "library", "libraries", "profanity", "sexual", "classroom", "challenge", "challenges", "porn", "pornography" to identify reviews related to book bans.

Table A2: Impact of Ban Events on Circulation of Other Books

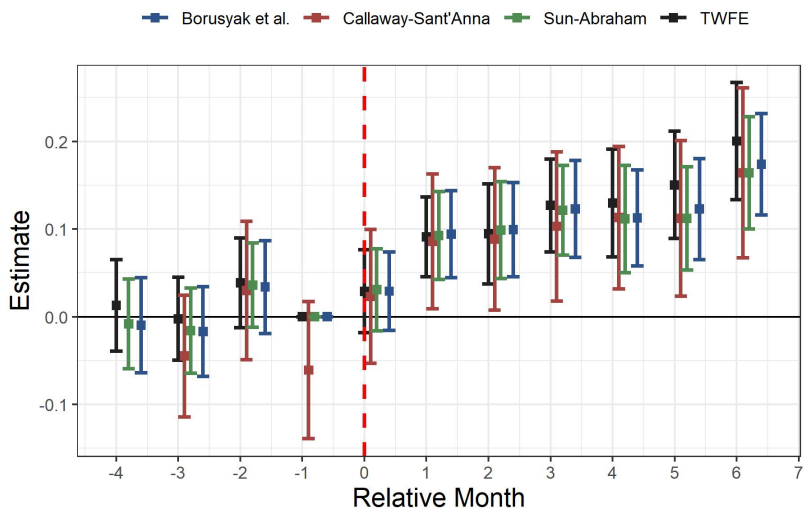
Dependent Variable:	Log(Circulation)		
	Genre	BISAC	Author
	(1)	(2)	(3)
Variables			
PostBan × Banned	-0.024	-0.0005	-0.056
	(0.060)	(0.091)	(0.035)
PostBan × Banned × SameGenre	-0.028		
	(0.089)		
PostBan × Banned × SameBISAC		-0.034	
		(0.110)	
PostBan × Banned × SameAuthor			0.021
			(0.055)
State Fixed Effect	Yes	Yes	Yes
Title Fixed Effect	Yes	Yes	Yes
Year Month Fixed Effect	Yes	Yes	Yes
R^2	0.60122	0.54188	0.60145
Observations	481,418	1,046,670	200,298

Note: Robust standard errors clustered at the level of the state in parentheses. The unit of observation is state-title-month. The dependent variable is the log of circulation. All the banned books are excluded from the sample. SameGenre indicates if a book is of the same genre as banned books. SameBISAC indicates if a book has the same BISAC as the banned book BISACs. SameAuthor indicates if a book is authored by an author of banned books.
$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

column (3), we find an increase in the ratings for the book post-ban. Based on an alternative dataset, these results increase confidence in the overall takeaways of the paper.

It might be that some of the reviews written were motivated in the contentiousness of the issue (for example, reviews written "in solidarity" with the author or completely against the books) — some reviewers are perhaps not reading the book, but are either rating and reviewing the books positively or negatively depending on their viewpoint of the ban. If this is the case, there should be less information in reviews that are written after the first ban. However, we find that there is no change in the number of words in reviews after the ban (b=2.07, p=0.82). There is also no change in the number of days it takes for the reviewers

Figure A4: Event Study Estimates using Matched Controls



Note: The figure shows event study estimates using four different estimation methods and control book titles derived using matching. The dependent variable is the log of circulation. The unit of observation is at the title-state-month level. 95% confidence intervals are shown.

to start and finish a book (mean is 8 days while median is 3, both pre and post ban). Both analyses provide descriptive evidence that the review information is not biased one way or another by readers who are not actually reading the book.

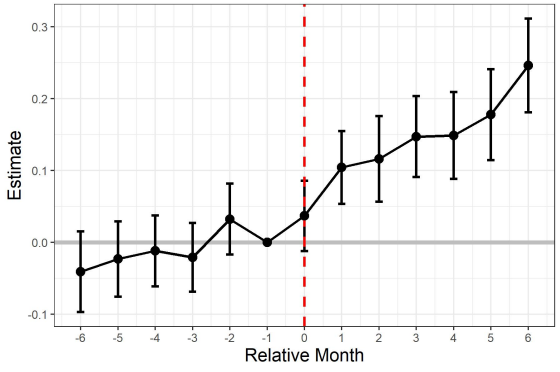
In terms of magnitudes, given that the number of reviews increased by 13%, with the mean number of reviews per banned book in a month was 47.2, indicating that six additional reviews were written. This effect is economically significant, as only a small fraction (less than 5% based on publicly available estimates) leave reviews on Goodreads and online platforms, more generally. Further, in column (3), we find that the average rating increases by 0.29 for the banned books after the ban event. This magnitude is economically meaningful according to the literature (Luca 2016). The change in average ratings could result from a change in the type of reader who reads the book in the aftermath of a book ban. As the composition of readers changes, the average star rating can also change, as has been pointed out in related studies by Zhu et al. (2023) and Kovács and Sharkey (2014).

Table A3: Matched Controls: Impact of Ban Events on Circulation

Dependent Variable:	Log(Circulation)						
	All States	Banned States	Non-banned States	Author Experience	Twitter Visibility	Blue States	Red States
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Variables							
PostBan × Banned	0.092***	0.103***	0.077***	-0.028	-0.008	0.096***	0.089***
	(0.016)	(0.023)	(0.022)	(0.022)	(0.021)	(0.020)	(0.027)
PostBan × Banned × Less-experienced Author				0.231***			
				(0.032)			
PostBan × Banned × High Visibility					0.211***		
					(0.037)		
State Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Title Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year Month Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.68425	0.67711	0.70149	0.68485	0.68475	0.70462	0.67899
Observations	28,044	16,974	11,070	28,044	28,044	14,760	13,284

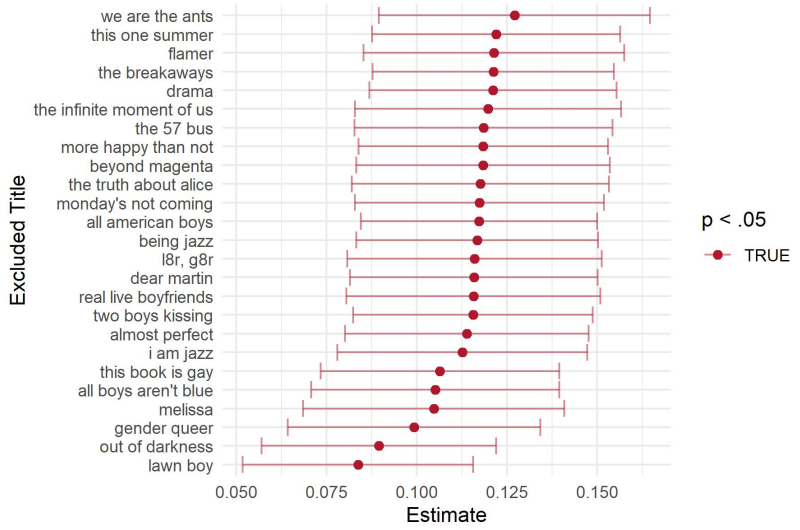
Note: Robust standard errors clustered at the level of the state in parentheses. The unit of observation is state-title-month. The dependent variable is log of circulation. PostBan is an indicator variable which denotes if the time period is after the month following the first time the title is banned in any state. PostBan is 0 for all control titles. Banned indicates whether the title belongs to the banned or control group. Column (1) includes all the states, column (2) includes only those states with at least one ban event, column (3) includes only those states without any ban events, column (6) includes only Red states, and column (7) includes only Blue states. Column (4) includes author experience as interaction variable and column (5) includes Twitter visibility as interaction variable.
$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

Figure A5: Event Study Estimates using Alternate Time Window



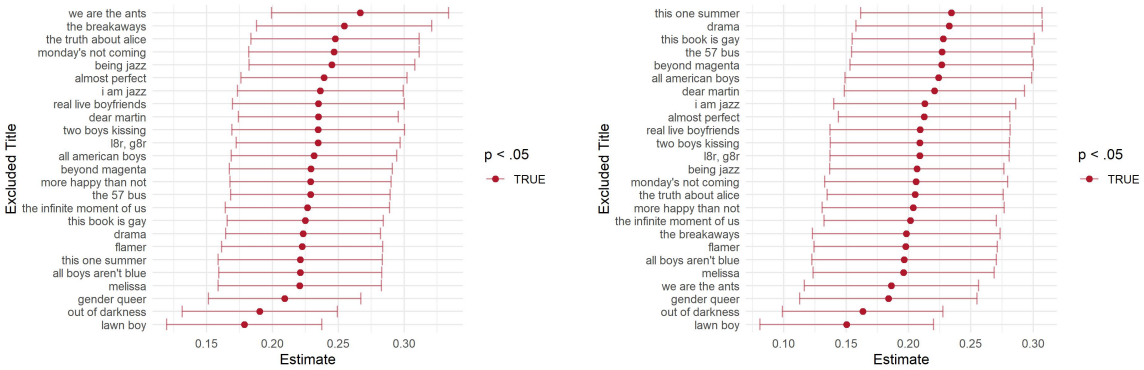
Note: The figure shows the event study estimates using the TWFE estimator with the log of circulation as the dependent variable. It plots the estimates along with the 95% confidence intervals using time window of 6 months before the first ban and 6 months after the first ban.

Figure A6: Leave One Out TWFE Estimates



Note: The figure shows estimates using leave one out TWFE estimates. Each dot represents the TWFE estimates after leaving the title on the y axis out of the analyses. 95% confidence intervals are shown.

Figure A7: Leave One Out TWFE Estimates - Heterogeneity



a) Author Experience b) Twitter Visibility

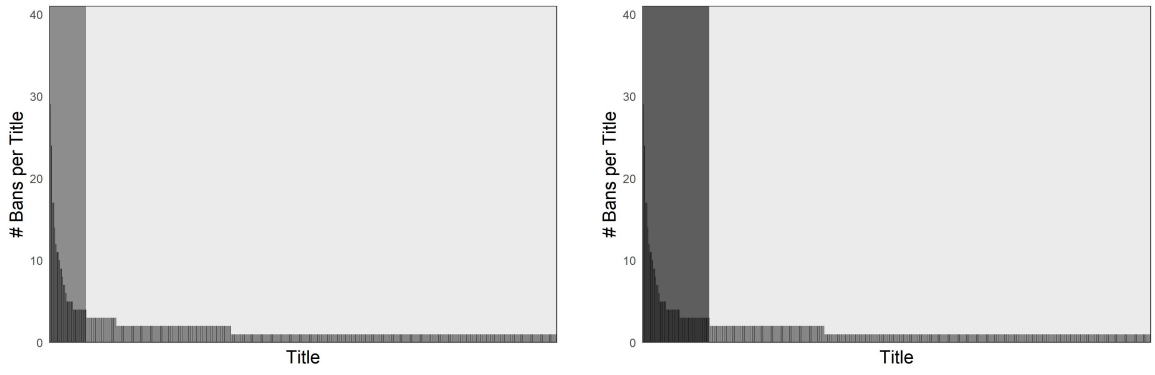
Note: The figure in panel a) shows the estimates from leave-one-out analyses moderated by author experience. The figure in panel b) shows the estimates from leave-one-out analyses moderated by twitter visibility. 95% confidence intervals are shown.

Table A4: Impact of Ban Events on Circulation

Dependent Variables:	Circulation		Lift		Log(Circulation)		
	IHS	Negative Binomial	Category Lift	Mean Lift	Organic Growth	False Treatment	Alternate Pre-period
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Variables							
PostBan × Banned	0.139***	0.211***	0.315**	0.462***	0.127***	0.023	0.123***
	(0.020)	(0.067)	(0.135)	(0.145)	(0.029)	(0.019)	(0.018)
State Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Title Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year Month Fixed Effect	Yes	Yes	Yes	Yes		Yes	Yes
Quarter-State Fixed Effect					Yes		
Quarter Fixed Effect					Yes		
R^2	0.57200		0.32475	0.30993	0.58110	0.53950	0.57821
Observations	249,660	243,090	249,660	249,660	249,660	249,850	251,560

Note: Robust standard errors clustered at the level of the state in parentheses. The unit of observation is state-title-month. The dependent variable is log of circulation. Columns (1) and (2) test for alternate model specification using inverse hyperbolic sine transformation and negative binomial models respectively. Columns (3) and (4) use category lift and mean lift in circulation as the dependent variable. Column (5) tests for organic growth assumption. Column (6) tests for falsification of treatment timing. Column (7) tests for the alternate pre-period window.
$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

Figure A8: Histogram of Number of Bans per Title



a) More than four bans b) More than three bans

Note: The figure shows the distribution of the number of total ban events per title arranged by the decreasing count of ban events. The marked area in panel a) shows books that are banned four or more times. The marked area in panel b) shows books that are banned three or more times.

Table A5: Additional Robustness: Impact of Ban Events on Circulation

Dependent Variable:	Log(Circulation)					
	At least Four Bans	At least Three Bans	At least Four Bans	At least Three Bans	Impact Per Ban	More vs. Less Bans
	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
PostBan × Banned	0.049**	0.031**	0.203**	0.155**	-0.011	0.010
	(0.018)	(0.014)	(0.072)	(0.068)	(0.018)	(0.009)
PostBan × Banned × Intensity of Ban					0.021***	
					(0.002)	
PostBan × Banned × >= Five Bans						0.097***
						(0.015)
State Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
Title Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
Year Month Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.59327	0.59922	0.55316	0.55587	0.57659	0.62823
Observations	257,602	268,242	146,712	147,214	249,660	421,382

Note: Robust standard errors clustered at the level of the state in parentheses. The unit of observation is state-title-month. The dependent variable is the log of circulation. Columns (1) and (3) use all the books that were banned at least four times. Columns (2) and (4) use all the books that were banned at least three times. Columns (1) - (2) show the average estimates using the main event identification assumption. Columns (3) and (4) show the average estimates using an alternate event identification assumption. Column (5) uses monthly cumulative count of bans as explanatory variable. Column (6) uses all the banned books with split of ban counts as interaction variable.
$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

Table A6: Impact of Ban Events on Goodreads reviews

Dependent Variables:	Reviews	Reviews with Bans	Ratings
	(1)	(2)	(3)
Variables			
PostBan × Banned	0.131**	0.183**	0.298**
	(0.054)	(0.073)	(0.116)
Title Fixed Effect	Yes	Yes	Yes
Year Month Fixed Effect	Yes	Yes	Yes
R^2	0.87292	0.79331	0.32797
Observations	2,738	2,738	2,738

Note: Heteroskedasticity-robust standard errors in parentheses. The unit of observation is title-month. Columns (1) and (2) use log count of reviews as the dependent variable. Column (1) includes all the reviews, and column (2) includes only those reviews with reference to book bans. Column (3) uses numerical rating as the dependent variable and includes all the reviews.
$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.