Supplementary Appendix for Pop Culture Censorship & Authoritarian Stability

March 10, 2020

Contents

A Text Analysis Methods **1**
- A.1 sIBP Estimation 1
- A.2 Tracking Themes Coding Scheme 3

B Correlations Between Terms and Censorship **4**

C sIBP Additional Information **5**
- C.1 sIBP Features 5
- C.2 sIBP Robustness Checks 5

D Tracking Themes Additional Information **7**
- D.1 Hand-Coded Content Measures 10
- D.2 Robustness Checks 11
- D.3 Education Censors and Voting 13

E Alternative Explanations **15**
- E.1 Distributor Self-Censorship 15
- E.2 Censor Preferences 20

F Political Content **20**
- F.1 Censorship Change 20
- F.2 Hand-Coded Political Content 21

Appendix A: Text Analysis Methods

A.1 sIBP Estimation

The supervised Indian Buffet Process (sIBP) is an automated method of text analysis designed to identify features (clusters of terms) from high-dimensional texts associated with a low-dimensional outcome. It takes as an input a document term matrix, in this case of film keywords. 50% of observations are then used to train a model, and 50% to test the model, to identify coefficient estimates and 95% confidence intervals. This avoids overfitting or p-hacking. The sIBP estimates the marginal effect of any one feature on the outcome (Hainmueller, Hopkins, and Yamamoto 2014), accounting for the fact that bigger values on one dimension often imply smaller values on another (Fong and Grimmer 2016). Specifically, the sIBP estimates the Average Marginal Component Specific Effect for factor k (AMCE$_k$). AMCE$_k$ is defined as the the difference in the outcome given that a feature is or is not present, averaged across all other features. More formally, following Fong and Grimmer (2016): "$AMCE_k = \int_{\mathbf{Z}_{-k}} \mathbb{E}[Y(Z_k = 1, \mathbf{Z}_{-k}) - Y(Z_k = 0, \mathbf{Z}_{-k})]m(\mathbf{Z}_{-k})d\mathbf{Z}_{-k}$, where $m(\mathbf{Z}_{-k})$ is some analyst-defined density on all elements but k of the treatment vector" (3). Following convention, sIBP chooses $m(\cdot)$ - the density of $\mathbf{Z}_{-k}$ across the population – as uniform. Other controls are added into models after identifying features and their scores across films.

Two parameters additionally affect results: α, which determines how common the text treatments are, and σ, which determines the degree to which the treatments must explain variation in the text. Since outcomes are sensitive to starting values, I search across several. The sIBP also takes as an input iterations, the number of starting values to attempt for each combination of α and σ, "to evaluate the output at several different local modes" (Fong and Grimmer 2016). I set this to eight. The number of features that the model outputs can also be chosen, and I set this at five. This tended to produce the most coherent output of features, but results, particularly for sexual content, are robust to alternatives.

Parameter values are selected based on a quantitative metric and qualitatively. For each set of parameter values, the sIBP calculates a coherence and exclusivity metric (Mimno et al. 2011; Roberts, Stewart, and Airoldi 2016, see Fong and Grimmer (2016), pages 5-6, for further details about estimation). Based on the most influential words for each feature, final parameter values are then chosen qualitatively, based on both performance on the exclusivity metric and substantive interpretability.

The sIBP requires four conditions to be met to be interpreted causally. Developed to identify the latent features most determinant of an outcome when text is randomly assigned, these are particularly challenging for observational data. The first assumption is that the response of an individual is determined only by the document they view. This is likely to be met in the pop culture censorship data: censors ban films because of the content of those movies. The second assumption requires that individuals not select into the text they receive, which here requires that censors should not be persistently more or less likely to watch certain classes of film. While partially accounted for through censor fixed effects, it remains an issue if censors consistently watched or were excluded from watching types of movies. CCC administrative practices reduce these concerns: councilors were assigned to parlors to watch films at different times depending on schedules, and multiple censors watched each film. To empirically explore this, I test whether any censors were significantly more or less likely to watch films rated as R or higher by the MPAA. 7 of 69 censors were significantly more likely to watch R-rated content, 4 of these by less than 10%; removing the films these censors reviewed does not impact results. An additional assumption, almost certainly met by the data, is that important features do not perfectly covary (e.g., that sex and gore may appear together but do not *always* appear together).

The last assumption is the most difficult to verify in this data: that any observed latent features (Z) are uncorrelated with any unobserved features (W). Otherwise, sIBP cannot marginalize over W. This is a strong assumption in the case of films, because movies are a complex media, and decisions over censorship may relate to visual imagery that cannot be fully captured when flattening films to keywords. That IMDb summaries and keywords are user-submitted makes this particularly difficult to evaluate. Qualitatively, we can validate that the content identified in the sIBP is linked to censorship using comments left on *expedientes*, which suggests those themes identified in the sIBP motivated censorship. Still, this assumption is a strong one to meet to call these results causal (Fong and Grimmer 2019).

A.2 Tracking Themes Coding Scheme

Table A1: Coding scheme.

	Politics
Revolution	References to real uprisings ("Russian Revolution") or collective action ("rebellion", "protest").
Communism	Mention of communism ("Marxism", "Trotsky") and communism-related ideas ("class warfare", "class consciousness").
Authoritarianism	References to dictatorships ("autocrat"), government oppression ("repression") and authoritarian right-wing ideologies ("fascism").
	Gore
Gore	Particularly violent imagery ("spurting blood", "beheading") or gore ("slasher", ' 'homicidal maniac"). Excludes lower levels of violence ("shot in the chest", "blood").
	Sexual Content
Nudity	Partial or complete nudity (e.g., "nude male", "breasts").
Sexual extremes	Sexual activities or preferences considered extreme, like fetishes (e.g. "bondage", "voyeurism").
Other	Terms referencing sexual content not included in categories above (e.g."panties," "sex comedy").
	Moral Taboos
Homosexuality or gender identity	References to homosexuality; references to transgenderism, transsexualism, or transvestitism.
Abortion	References to abortion.
Sexual assault	Nonconsensual sexual content ("rape", "sexual assault").
Sex work	References to sex work or sex workers ("prostitution", "brothel").
Suicide	References to suicide.
Drugs	References to drugs or drug use.
Cheating	Terms related to extramarital affairs (e.g., "mistress", "cheating on husband").
Cannibalism	Reference to cannibalism, zombies, vampires, and eating flesh.
Catholic themes	References to biblical characters or Catholic figures ("pope", "Christ").
Child abuse	References to abusing children.
Incest	References to incestuous relationships.
Necrophilia	References to necrophilia.
Bestiality	References to bestiality.

Appendix B: Correlations Between Terms and Censorship

Figure B2 shows a collection of selected keywords and their relationship to censorship, defined as the probability a movie is censored given that its IMDb page does and does not contain that keyword ($Pr(Censorship|X = 1) - Pr(Censorship|X = 0)$, where X is the keyword of interest). Figure B1 show the keywords with the largest positive score on this metric, meaning those keywords that are most related to censorship. Figure B2 shows selected terms of interest and their relationship to censorship. Both confirm that immoral content played a central role in censorship in Chile.

Figure B1: Most influential keywords.

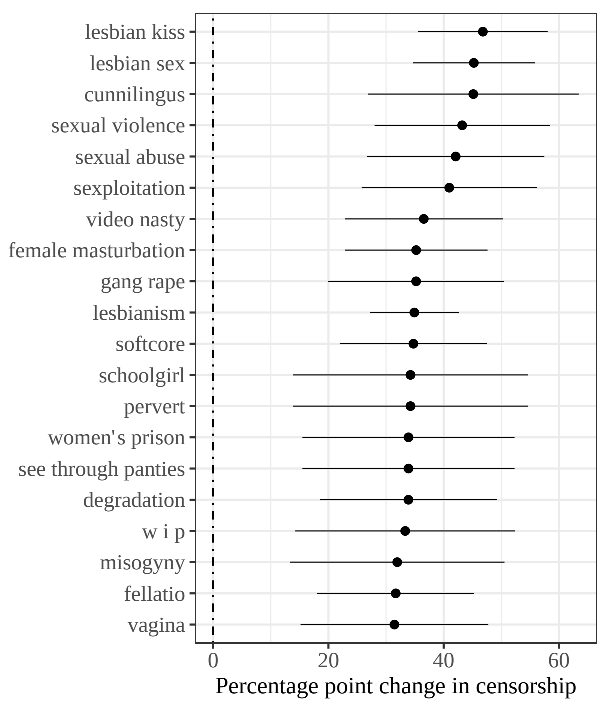

Figure B2: Selected terms and censorship.

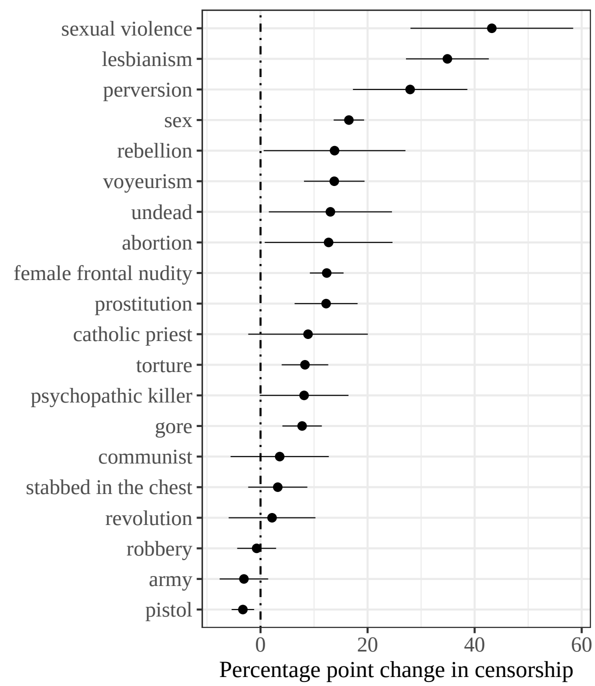

Appendix C: sIBP Additional Information

C.1 sIBP Features

Table B1: Top 20 feature keywords.

Feature	Keywords
1	maniac, homicidal maniac, psychopathic killer, psycho killer, fear, cruelty, knife murder, desire, serial murder, american horror, evil, human monster, slasher, psycho, terror, rage, rampage, mysterious man, video nasty, insanity
2	singing, apology, opening a door, train, scene during opening credits, visit, lying on bed, euphemism, shaking hands, map, man with glasses, applause, bomb, horse, singer, dancer, umbrella, sexual euphemism, telling someone to shut up, song
3	showdown, tough guy, action hero, blood splatter, hand to hand combat, hand to hand combat, fistfight, gunfight, warrior, one man army, one against many, brawl, disarming someone, bad guy, shootout, violence, hero, mixed martial arts, serial murder, anti-hero, martial arts
4	blood, death, violence, blood splatter, cult film, corpse, knife, fear, sadism, murder, gore, shot to death, cruelty, pistol, maniac, cigarette smoking, psychopath, subjective camera, revolver, evil
5	scantily clad female, leg spreading, female removes her clothes, panties pulled down, female pubic hair, lust, voyeurism, female removes her dress, lingerie, black panties, voyeur, female rear nudity, no panties, cleavage, white panties, blonde, sexual desire, panties, promiscuous woman, erotica

C.2 sIBP Robustness Checks

In order to ensure that the results of the sIBP analysis truly reflect important text treatments, Figure B3 shows the results with censor fixed effects, including only censors who watched at least five films. This increases confidence in results in three ways. First, it shows that missingness in the *expedientes* data does not significantly change results, reducing concerns of bias. Second, it reduces threats to non-random assignment, strengthening confidence in results. Third, it shows that idiosyncratic censor preferences are unlikely to be the central driver of findings. Results hold, though as in the body of the paper violence plays a more ambiguous role in censorship than sexual content.

Figure B3: sIBP features and censorship, with censor fixed effects.

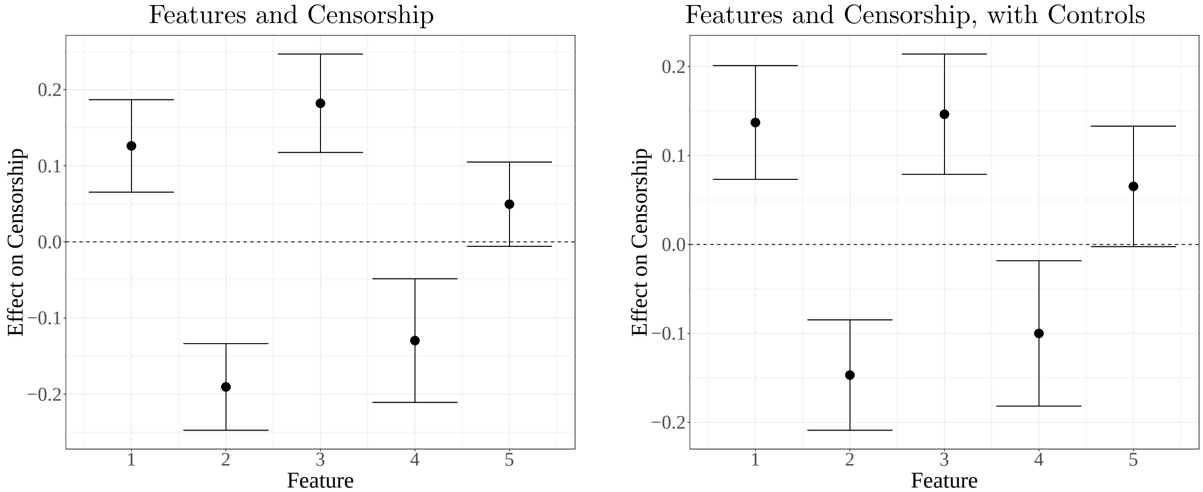

Feature				
1	2	3	4	5
psychopathic killer	action hero	sexual desire	euphemism	blood
homicidal maniac	shootout	sexual attraction	opening a door	blood splatter
psycho killer	tough guy	psycho killer	scene during opening credits	death
maniac	brawl	maniac	sexual euphemism	violence
psycho	pistol	psychopathic killer	apology	corpse

N =4,985. 95% confidence intervals. Censor, genre, year, number of keywords, and video controls.

Figure B4 subsets the data to include only movies that were rated as for 18 and over or higher. This reduces concerns that the sIBP results simply capture the content of movies suitable only for adults. Results hold, particularly for sexual content, though the effect of violent materials is more ambiguous and less robust. Still, particularly gory content appears correlated with censorship.

Figure B4: Keywords and censorship, 18+ only.

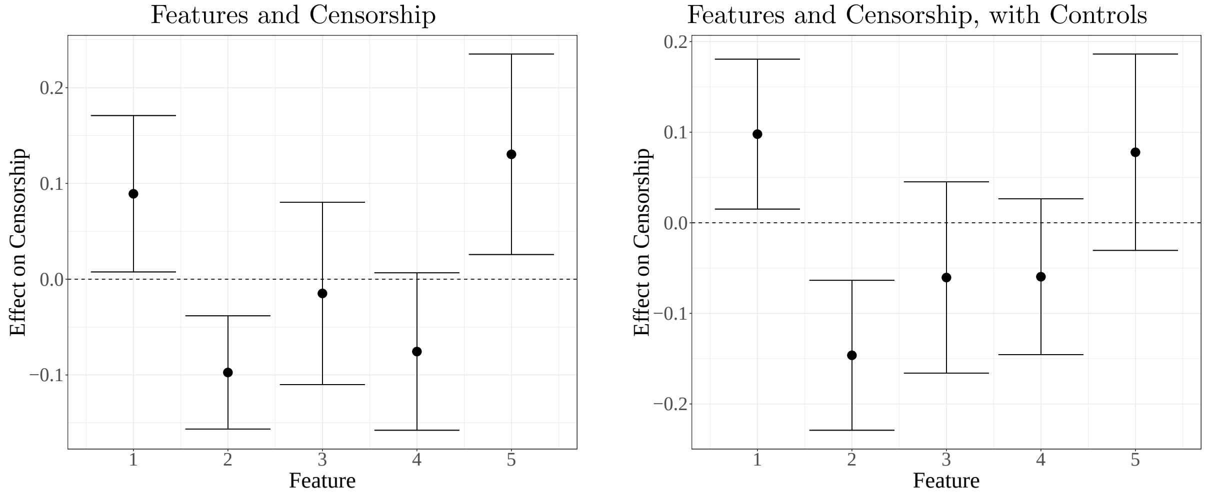

Feature				
1	2	3	4	5
female removes her clothes	scene during opening credits	blood	gunfight	evil
scantily clad female	camera	pistol	showdown	slasher
cleavage	photograph	blood splatter	tough guy	maniac
female removes her dress	telephone call	revolver	one man army	american horror
white panties	dancing	death	shootout	psycho killer

N =2,706. 95% confidence intervals. Genre, year, number of keywords, and video controls. Restricting analysis to movies rated as for 18 and over reduces concerns that results are driven by the difference between adult and children's movies.

Appendix D: Tracking Themes Additional Information

Figure D1 shows the percent of films reviewed with at least one keyword associated with a given category (right), and the rate of censorship for each category (left). Table D1 reports summary statistics for the major keyword variables. Table D2 shows the relationship between censorship the number of keywords associated with each major category, using a logistic regression with heteroskedasticity-robust standard errors clustered at the day of review. Table D3 shows the correlation between censorship and themes using only movies where censors' votes are known, controlling for censor fixed effects and clustering robust standard errors at the censor group. Table D4 replicates this, but clusters the standard errors at the censor group-date of review using multiway clustering. Table D5 shows the distribution of keywords variables, to improve the interpretation of substantive effects: overall five keywords is a large increase within the support of all categories. Among banned films, five keywords is in the 55th percentile of movies for sexual content, the 76th for moral taboos, the 93rd for gore, and the 96th for political topics.

Figure D1: Movie types, reviewed by year and banned.

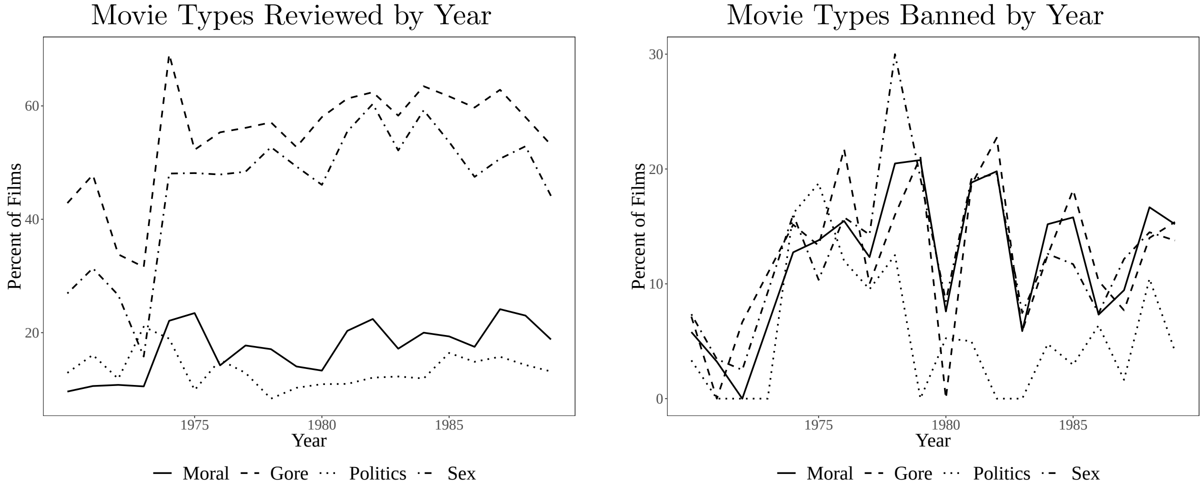

Table D1: Summary Statistics for Keyword Variables

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Keywords	6,361	51.723	63.536	1	9	68	553
Sexual Keywords	6,361	3.406	7.320	0	0	3	92
Moral Taboos Keywords	6,361	1.548	2.934	0	0	2	36
Gore Keywords	6,361	0.604	1.982	0	0	0	28

Table D2: Tracking Themes Models

	(1)	(2)	(3)	(4)
Politics	0.031			
	(0.029)			
Sexual Content		0.081***		
		(0.006)		
Gore			0.156***	
			(0.021)	
Moral Taboos				0.198***
				(0.019)
Constant	−1.682***	−1.721***	−1.615***	−1.809***
	(0.270)	(0.277)	(0.271)	(0.281)
Year FE	Y	Y	Y	Y
Controls	Y	Y	Y	Y

*p<0.1; **p<0.05; ***p<0.01

Table D3: Tracking Themes Models, Clustered at the Censor Group

	(1)	(2)	(3)	(4)
Politics	0.041			
	(0.031)			
Sexual Content		0.085^{***}		
		(0.007)		
Gore			0.195^{***}	
			(0.023)	
Moral Taboos				0.193^{***}
				(0.021)
Constant	-3.060^{***}	-3.138^{***}	-2.961^{***}	-1.846^{***}
	(0.728)	(0.756)	(0.730)	(0.250)
Year FE	Y	Y	Y	Y
Censor FE	Y	Y	Y	Y
Controls	Y	Y	Y	Y

*p<0.1; **p<0.05; ***p<0.01

Table D4: Tracking Themes Models, Clustered at Censor Group-Date

	(1)	(2)	(3)	(4)
Politics	0.041			
	(0.032)			
Sexual Content		0.085^{***}		
		(0.007)		
Gore			0.195^{***}	
			(0.024)	
Moral Taboos				0.193^{***}
				(0.021)
Constant	-3.060^{***}	-3.138^{***}	-2.961^{***}	-1.846^{***}
	(0.761)	(0.765)	(0.739)	(0.253)
Year FE	Y	Y	Y	Y
Censor FE	Y	Y	Y	Y
Controls	Y	Y	Y	Y

*p<0.1; **p<0.05; ***p<0.01

Table D5: Distribution of Keywords Variables

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Number of Keywords	6,361	51.723	63.536	1	9	68	553
Political Keywords	6,361	0.487	1.617	0	0	0	23
Sexual Content Keywords	6,361	3.406	7.320	0	0	3	92
Gore Keywords	6,361	0.604	1.982	0	0	0	28
Moral Taboo Keywords	6,361	1.548	2.934	0	0	2	36

D.1 Hand-Coded Content Measures

In order to ensure the validity of the tracking themes content measure, I hand-coded a random sample of 1,000 films for whether they prominently featured sex, gore, or immoral content in descriptions found online (sources included IMDb, Wikipedia, and review sites). I then used whether or not the film contained at least one keyword related to gore, sexual content, or moral taboos as a classifier; classification metrics show this performs fairly well (Table D6). Classification metrics used are F1, which balances precision and recall and is most useful for uneven distributions of 0s and 1s[1]; F2, which similarly balances precision and recall but places a higher weight on precision; accuracy, which again balances precision and recall but weights true negatives more highly; Cohen's Kappa, which measure how closely classification matches the truth, controlling for the accuracy of a random classifier (the levels reported here puts the classifier as reasonably well performing); sensitivity, which measures the rate at which true positives are identified; and specificity, which measures the rate at which true negatives are identified. I do the same for political content, but here I have classified the full sample of films for whether they contain a significant amount of content related to revolution, communism, or authoritarianism. While the coding takes into account the direction of these references – e.g. whether the theme is portrayed positively, negatively, or neutrally – here I simply measure whether the film had significant political content in its plot (Table D7). Overall, having at least one keyword associated with a given theme performs about as well overall as using a higher threshold.

[1]Precision is true positives over predicted positives; recall is true positives over actual positives.

Table D6: Classification of Immoral Films

	Classification Threshold (Number of Keywords)		
Metric	**>0**	**>1**	**>5**
F1	.74	.80	.83
F2	.73	.70	.66
Kappa	.48	.48	.44
Accuracy	.73	.77	.78
Sensitivity	.91	.84	.76
Specificity	.60	.68	.82

Table D7: Classification of Political Films

	Classification Threshold (Number of Keywords)		
Metric	**>0**	**>1**	**>5**
F1	.94	.96	.95
F2	.65	.65	.28
Kappa	.60	.62	.25
Accuracy	.90	.93	.92
Sensitivity	.98	.96	.91
Specificity	.51	.67	.84

D.2 Robustness Checks

This section demonstrates that results hold across several robustness checks: using unigrams instead of keywords, to show results are similar when drawing on plot summaries; controlling for censor fixed effects, demonstrating results are not due to individual censor preferences; considering bivariate correlations only, demonstrating the consistency of results; normalizing theme measures by the number of keywords in a film, to account for variation in the amount of user-submitted text; and including only films rated as for 18 or 21 and over, to ensure that results differentiate mature films from one another. All figures show substantive effects with 95% bootstrapped confidence intervals clustered at the day of review, with the exception of the censor fixed effects figure, which uses censor group (the set of censors that watched a film) as the cluster variable. The censor fixed effects additionally only includes censors who watched at least 20 films to minimize convergence issues when bootstrapping, but results hold when using standard robust clustered standard errors at the censor group level and including censors who watched at least five films. Except the bivariate relationships, all models other include genre, year, number of keyword, and video controls. Normalized models show the effect of increasing theme keywords by one, because larger standard errors otherwise makes visualization difficult, while the others show the effect of increasing theme keywords by five.

Figure D2: Themes and censorship, unigrams.

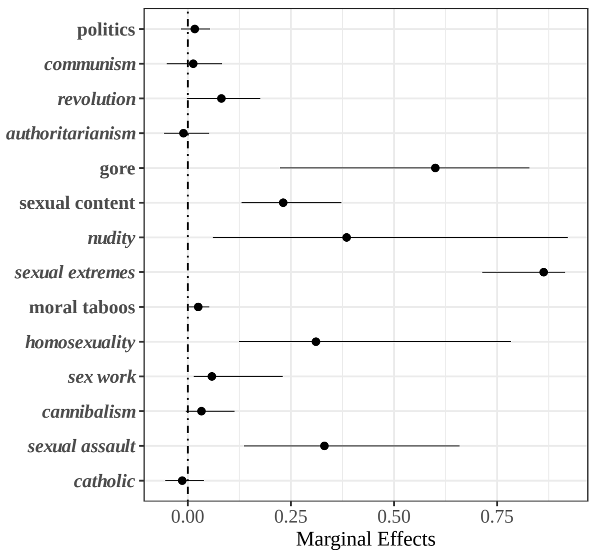

Figure D3: Censor controls.

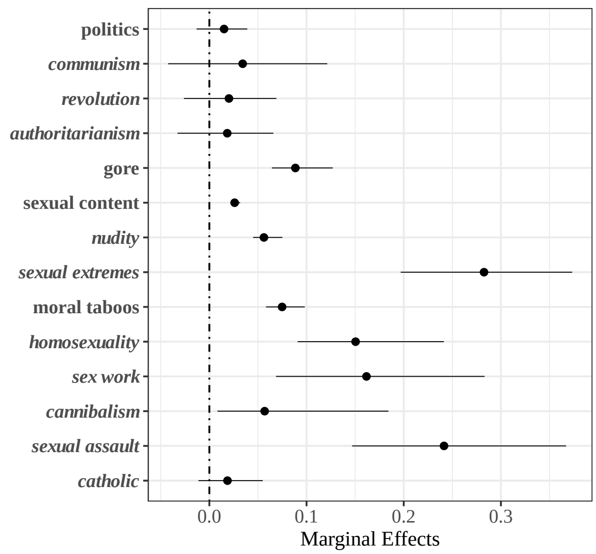

Figure D4: Bivariate relationships.

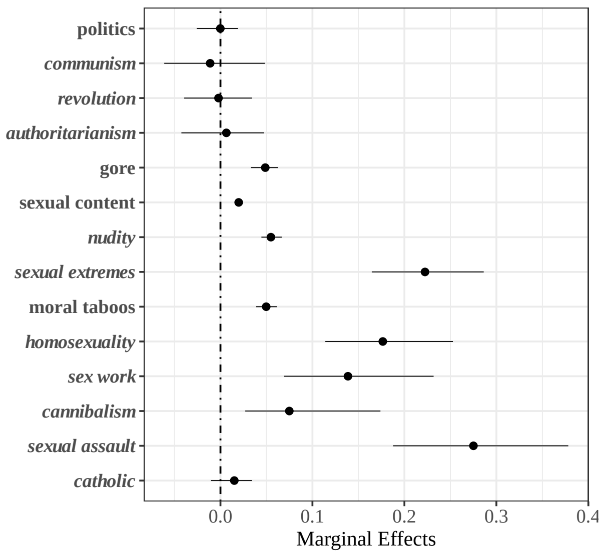

Figure D5: Normalized by keywords.

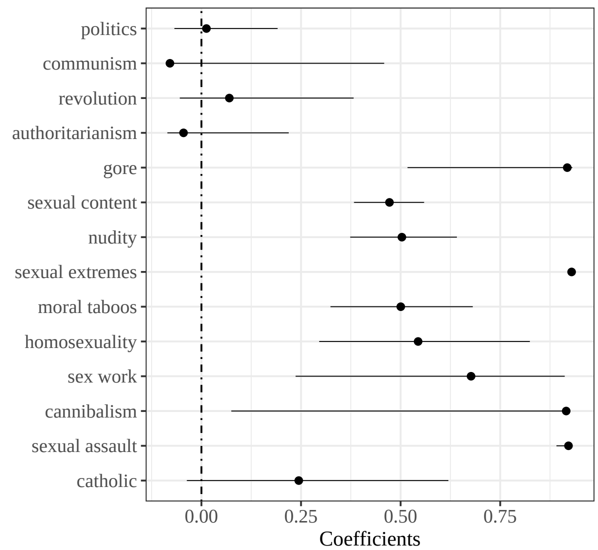

Figure D6: 18+ movies.

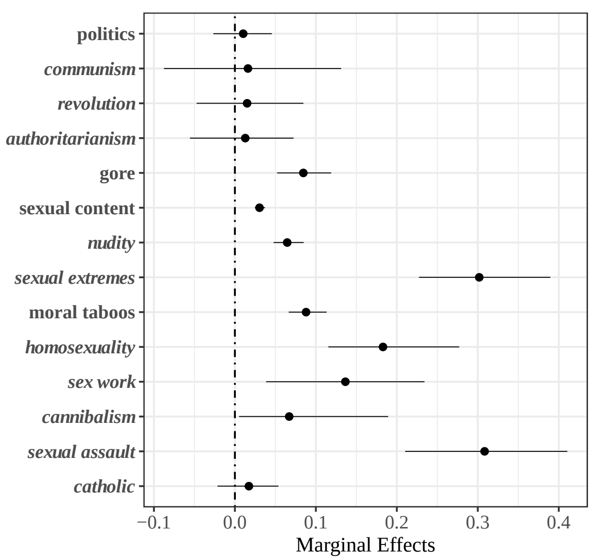

Figure D7: 21+ movies.

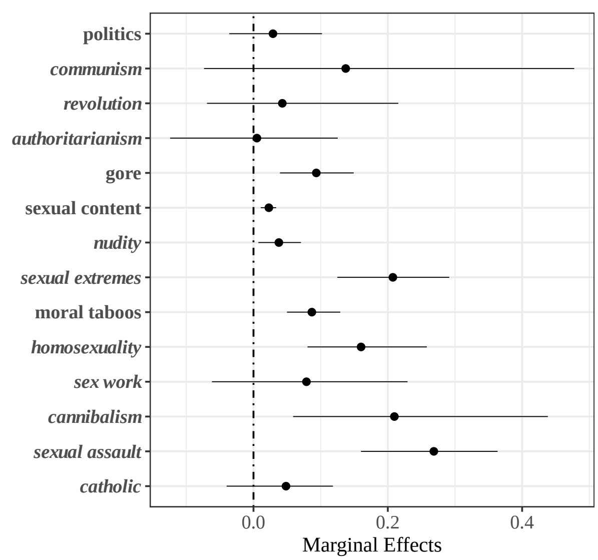

D.3 Education Censors and Voting

Education-appointed censors are expected to vote to restrict film content more aggressively, given their ties to Opus Dei. I first show that results hold when using robust standard errors clustered at the censor level through block bootstrap (Figure D8). While there are few clusters ($N = 20$), this provides greater confidence in results. I do not include the title fixed effects models, due to the computational intensity of bootstrapping these specifications. I additionally show that censors from the armed forces were more likely to vote to ban films with at least one political keyword, but were otherwise somewhat less likely to vote to prohibit films (Figure D9). The armed forces results require some caution, because the standard errors become larger when calculating confidence intervals through block bootstrap clustered at the censor level. Substantive effects remain, however, and – providing greater confidence in results – results are statistically significant when using the hand coded political metric (not reported).

Figure D8: Education-appointed censors and voting.

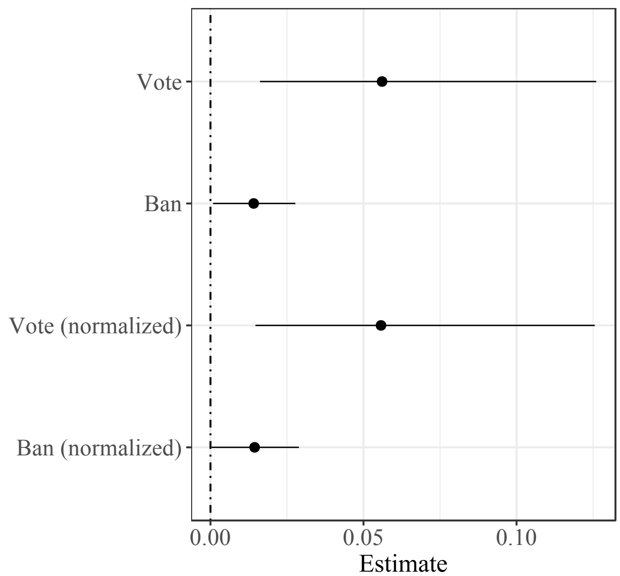

Correlation between education-appointed censors and voting. $N = 14,795$. OLS coefficients and 95% confidence intervals, block bootstrapped at the censor level (5,000 iterations). Results suggest that censors appointed by the MoE, the organization most linked to Opus Dei, are more likely to vote for immoral bans or restrictions.

Figure D9: Armed forces censors and voting.

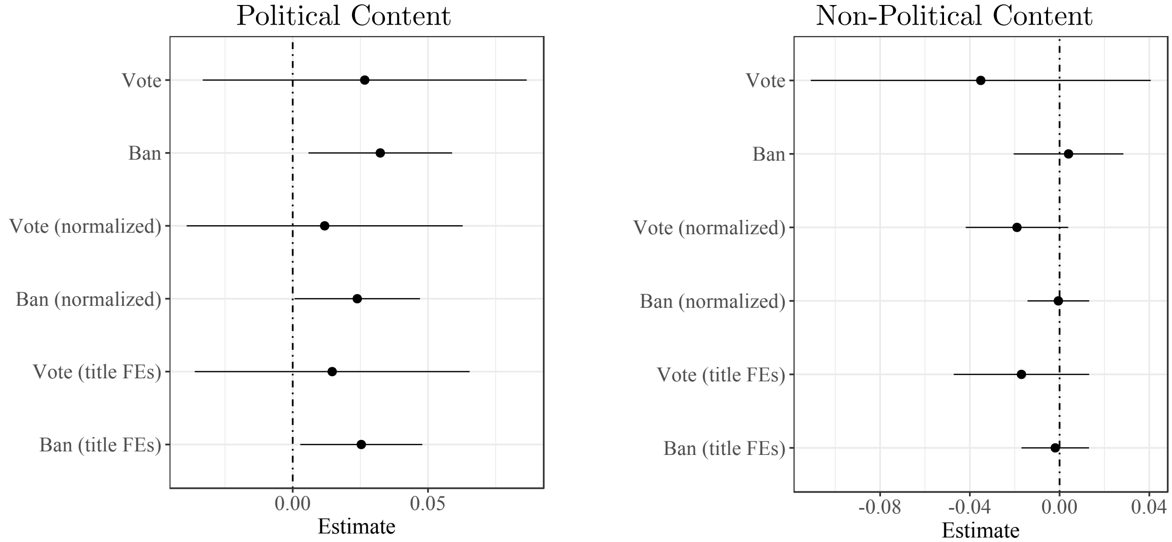

Correlation between armed forces censors and voting. $N = 2,436$ *(left) and* $N = 12,359$ *(Right). OLS coefficients and 90% robust confidence intervals. Results suggest that censors from the armed forces are more likely to vote for bans on political content, but not other films.*

Appendix E: Alternative Explanations

E.1 Distributor Self-Censorship

To explore the possibility of distributor self-censorship, I scrape the webpages of all movies produced between 1960 and 1995 from the Movie Database (TMDb). I compare films to those with at least one keyword and produced two years before, since this was the median delay in movies reaching Chile. To ensure the validity of TMDb results, I first show that substantive effects hold when using this data (Figure E1).

I then report that there is little correlation between keywords associated with a given theme of interest and review during the dictatorship, providing suggestive evidence that films were not less likely to be distributed to Chile based on content (Table E1). There is a small negative correlation between gore and review, though this is also true prior to the dictatorship, suggesting something in the structure of the market rather than a deliberate effort by distributors to withhold certain content.

Figure E1: Substantive effects, using Movie Database keywords.

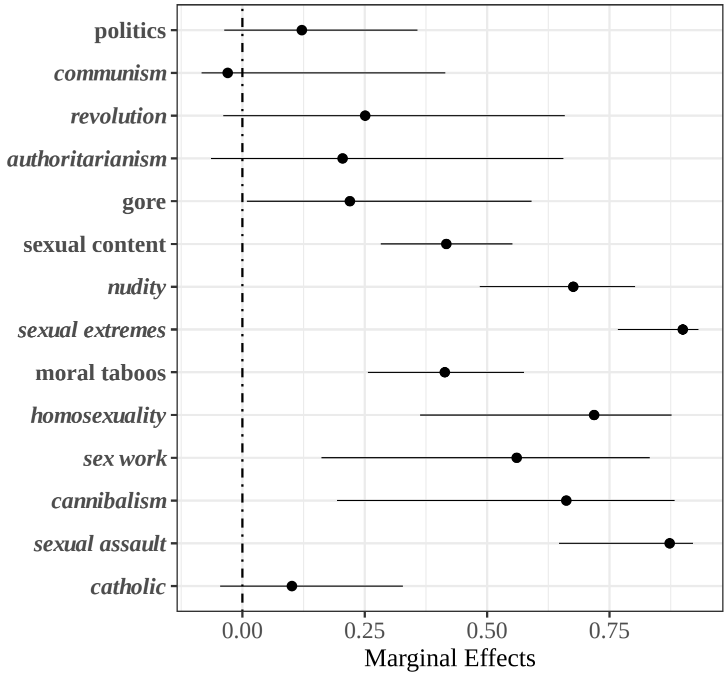

N=4,127, with 95% confidence intervals. Controls are genre, year, number of keywords, and video.

Table E1: Probability of Review

	(1)	(2)	(3)	(4)
Politics	0.002			
	(0.008)			
Sexual Content		0.003		
		(0.006)		
Gore			-0.033^{*}	
			(0.019)	
Moral Taboos				-0.005
				(0.008)
Constant	0.075^{***}	0.075^{***}	0.074^{***}	0.075^{***}
	(0.014)	(0.014)	(0.015)	(0.014)
Year FE	Y	Y	Y	Y
Controls	Y	Y	Y	Y

*p<0.1; **p<0.05; ***p<0.01
Logit with robust standard errors clustered at the day of review.

I next look for evidence of distributor self-censorship at two points: the creation of the CCC in 1974 and regime reorganization in 1979, when communism-related themes stop being significantly correlated with censorship, using a difference-in-differences framework. I define political films as those with at least one political keyword, though results hold when using alternate levels. Table E2 shows that, while the rate of review does increase considerably in 1974 – largely because of the virtual cessation of review in 1973 – this change is nearly identical for political and non-political films. In case self-censorship also affects sexual, gory, or morally taboo content, I also compare political films to films not associated with any of the four themes, with the same result.

Table E2: Probability of Review, Political and Non-Political Films

Before and After CCC			
	1969-1973	1974-1978	Difference
Political movies	10.62%	30.23%	19.61%**
Non-political movies	10.46%	21.63%	11.17%***
No-theme movies	10.96%	21.67%	10.7%***
Before and After 1979			
	1974-1978	1979-1983	Difference
Political movies	30.23%	27.51%	-2.72%
Non-political movies	21.63%	21%	-0.63%
No-theme movies	21.67%	17.7%	-3.97%***

Difference in means using t-tests. Results compare the rates of review before and after key years of change.

To formalize this I use logistic regression where the dependent variable is a binary indicator (*Reviewed*) for whether a film was reviewed by the CCC. The data is restricted to the three years directly before and directly after the change of interest, whether 1974 or 1979. I fit the following model:

$$Reviewed_i = \beta_0 + \beta_1 ChangeYear_i + \beta_2 Political_i + \beta_3 ChangeYear_i * Political_i + \chi_i + \gamma_i + \epsilon_i$$

Political is a binary variable capturing whether a film has at least one political keyword in the TMDb text data, though results hold using alternate thresholds. *Change Year* captures whether the year of release or review was before or after (inclusive) 1974 or 1979. The coefficient of interest is β_3, which reflects whether there is a significant decrease in the probability that a political film is reviewed compared to non-political films. χ_i represent controls, here genre fixed effects and number of keywords. β_0 is a constant, ϵ_i are error terms, and γ_i are year fixed effects. Results in Table E3 show political movies do not become any less likely to be reviewed in 1974 or 1979.

Table E3: Political and Non-Political Movies at Inflection Years

	1974	1979	1974	1979
	(1)	(2)	(3)	(4)
Politics	0.301	−0.207		
	(0.310)	(0.221)		
Change Year	1.043***	−0.201***		
	(0.078)	(0.060)		
Politics * Change Year	−0.482	−0.153	−0.201	−0.334
	(0.377)	(0.321)	(0.220)	(0.238)
Constant	−3.591***	−2.566***	−3.018***	−2.487***
	(0.103)	(0.077)	(0.131)	(0.109)
Year FE	N	N	Y	Y
Controls	Y	Y	Y	Y
Observations	7,335	8,502	7,335	8,502
Log Likelihood	−2,803.268	−3,751.534	−2,734.023	−3,727.115
Akaike Inf. Crit.	5,642.535	7,539.067	5,516.046	7,504.231

*p<0.1; **p<0.05; ***p<0.01

Difference in differences design, restricted to the three years before and after the inflection year of interest. Controls are genre fixed effects and number of keywords.

I additionally re-run results using the full list of films produced by the CCC, which covers review through 1994. This data is less complete than the *expedientes*, but it permits looking at any changes that occur during democratization in 1990. Since restrictions on political materials were removed from the CCC guidelines in the new constitution, an increase in the percent of political films reviewed after 1989 would provide evidence of distributor self-censorship. Results confirm the findings above (Tables E4, E5). There is no change in the percent of political movies before and after democratization.

To additionally confirm that there was no distributor self-censorship of immoral content, I show that there is no change in rates of review for movies classified by the U.S. MPAA as R, NC-17, or X-rated. Not all films received or reported an MPAA rating, and those that did were more likely to be rated high. Still, it shows no apparent change with the coup in the

Table E4: Probability of Review, Political and Non-Political Films, 1970-1993

Before and After Democratization			
	1986-1989	1990-1993	Difference
Political movies	36.36%	31.66%	-4.71%
Non-political movies	26.03%	26.44%	0.41%
No-theme movies	23.61%	23.88%	$-.27\%$

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Difference in means using t-tests. Results compare the rates of review before and after key years of change.

Table E5: Comparing Political and Non-Political Movies at Inflection Points, 1970-1993

	1974	1979	1990	1974	1979	1990
	(1)	(2)	(3)	(4)	(5)	(6)
Politics	0.508	−0.173	0.356*			
	(0.404)	(0.248)	(0.216)			
Change Year	2.702***	−0.642***	0.652***			
	(0.115)	(0.073)	(0.055)			
Politics * Change Year	−0.654	−0.182	−0.311	−0.197	−0.316	−0.043
	(0.471)	(0.377)	(0.262)	(0.251)	(0.295)	(0.157)
Constant	−5.602***	−2.888***	−2.585***	−6.830***	−2.721***	−3.660***
	(0.145)	(0.090)	(0.071)	(0.473)	(0.122)	(0.150)
Year FE	N	N	N	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y
Observations	11,099	8,443	10,628	11,099	8,443	10,628
Log Likelihood	−2,107.922	−2,786.648	−5,148.268	−2,034.854	−2,751.746	−4,963.057
Akaike Inf. Crit.	4,251.844	5,609.297	10,332.540	4,127.707	5,553.493	9,974.114

*p<0.1; **p<0.05; ***p<0.01

Difference in differences design, restricted to the three years before and after the inflection year of interest. Controls are genre fixed effects and number of keywords.

rate of R-rated films entering the country, either positive – suggesting that censorship was a mechanistic response to a change in the type of films entering the country – or negative, suggesting distributor self-censorship. Similarly, there is no increase or decrease in rates of review for R-rated films following democratization (Figure E2). Re-running the difference-in-difference analyses above with MPAA ratings similarly shows null results (not reported).

Figure E2: Percent of Films Rated as R or higher by MPAA

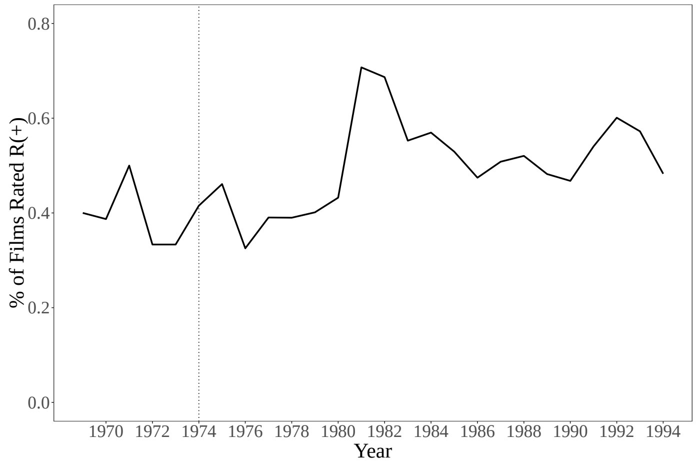

E.2 Censor Preferences

To demonstrate that the institutional environment plays a role in censorship independent of councilor preferences, Table E6 shows that democratization significantly decreases the probability of censorship even when controlling for year and genre fixed effects.

Table E6: Democratization and censorship

	(1)	(2)
Democratization	-1.923^{***}	-1.855^{***}
	(0.350)	(0.374)
Constant	-1.032^{***}	-0.957^{***}
	(0.323)	(0.350)
Year FE	Y	Y
Genre FE	Y	Y

*p<0.1; **p<0.05; ***p<0.01.
Logit with robust standard errors.

Appendix F: Political Content

F.1 Censorship Change

The results of this paper do not suggests that political themes played no role in censorship. Table F1 demonstrates that the *change* in censorship of political content is roughly equal to the change in rates of censorship for immoral content.

Table F1: Rates of Censorship Before and After CCC Reform

1970-1989			
	1970-1973	1974-1989	Difference
Political content	1.33%	7.16%	5.82%***
Sexual content	3.92%	13.24%	9.31%***
Gore	2.27%	12.12%	9.84%***
Moral taboos	2.76%	11.92%	9.17%***

F.2 Hand-Coded Political Content

To ensure that results are not a reflection of the data generating process, I additionally hand-code films for whether they contain negative depictions of right-wing authoritarianism or sympathetic portrayals of revolution or communism. As with the text analysis measure, political content related to collective action impacts censorship predominantly during the early years of the regime.

Figure F1: Censorship and political content (hand-coded), pre- and post-1978

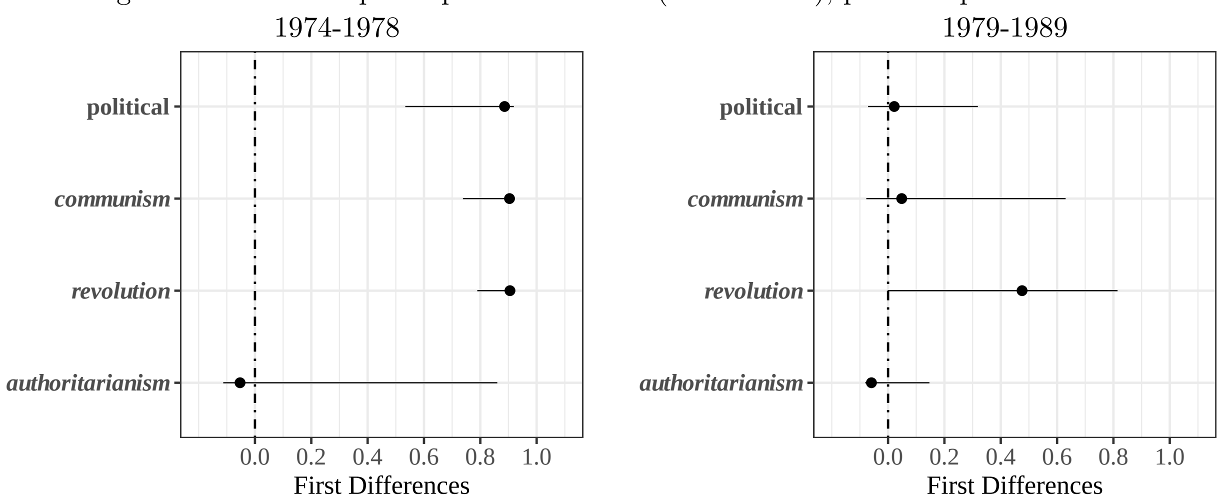

$N = 1,346$ (left) and $N = 5,015$ (right), with 95% confidence intervals (5,000 iterations) block bootstrapped at the day of review. Genre, year, number of keywords, and video controls.

References

Fong, Christian, and Justin Grimmer. 2016. "Discovery of Treatments from Text Corpora". *Proceedings of the Annual Meeting of the Association for Computational Linguistics.* `https://aclweb.org/anthology/papers/P/P16/P16-1151/`.

— . 2019. "Exploratory and Confirmatory Causal Inference for High-Dimensional Text Interventions". Working Paper. https://www.dropbox.com/s/b3f3soiqz3vuwob/dexp3.pdf?dl=0.

Hainmueller, Jens, Daniel Hopkins, and Teppei Yamamoto. 2014. "Causal Inference in Conjoint Analysis: Understanding Multi-Dimensional Policy Choices Via Stated Preference Experiments". *Political Analysis* 22 (1): 1–30.

Mimno, David, et al. 2011. "Optimizing Semantic Coherence in Topic Models". *Proceedings of the Conference on Empiricla mEthods in Natural Language Processing*: 262–272.

Roberts, Margaret E., Brandon M. Stewart, and Edoardo M. Airoldi. 2016. "A Model of Text for Experimentation in the Social Sciences". *Journal of the American Statistical Association* 111 (515): 988–1003.