

**KEEPING YOUR MOUTH SHUT:
SPIRALING SELF-CENSORSHIP IN THE UNITED STATES***

POLITICAL STUDIES QUARTERLY, FORTHCOMING.

James L. Gibson

Sidney W. Souers Professor of Government
Department of Political Science
Washington University in St. Louis
Campus Box 1063
St. Louis, MO 63130-4899
United States
jgibson@wustl.edu

Fellow, Centre for Comparative and International Politics
Professor Extraordinary in Political Science
Stellenbosch University (South Africa)

Joseph L. Sutherland

Visiting Assistant Professor
Department of Quantitative Theory & Methods
Emory University
36 Eagle Row, 5th Floor
Atlanta, GA 30322
United States
jls2316@columbia.edu

Fellow, Weidenbaum Center on the Economy, Government, and Public Policy
Washington University in St. Louis

Version 170, December 6, 2022

©

* The Freedom and Tolerance Surveys on which this article partially relies were funded by the Weidenbaum Center at Washington University in St. Louis. The authors appreciate the support made available for this research by Steven S. Smith, Director of the Center. Additional support for these surveys has been provided by the Law and Social Sciences Program of the National Science Foundation grants [SES-0943389 to J.L.G. and Chintan Turakhia and SES-0553156 to J.L.G.]. Any opinions, findings, and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the National Science Foundation. The 2020 AmeriSpeak Freedom and Tolerance Survey was also funded by the Weidenbaum Center at Washington University in St. Louis. The authors greatly appreciate the support provided for that research by Steven S. Smith, Director of the Center. We also appreciate the comments of Christopher Claassen and Taylor Carlson on an earlier version of this paper. We are particularly indebted to Professor Jacob M. Montgomery and The American Social Science Survey (TASS), funded by the Weidenbaum Center at Washington University in St. Louis, for support for our 2022 survey.

Abstract

Over the course of the period from the heyday of McCarthyism to the present, the percentage of the American people not feeling free to express their views has tripled. In 2020, more than four in ten people engaged in self-censorship. Our analyses of over-time and cross-sectional variability suggest that: First, self-censorship is connected to affective polarization among the mass public, with greater polarization associated with *more* self-censorship. Second, levels of mass opposition to full civil liberties bear no relationship to self-censorship. Third, those who perceive a more repressive government are slightly more likely to self-censor. Fourth, conservatives report engaging in *more* self-censorship than liberals (but this is not true when comparing Republicans to Democrats). Together, these findings suggest that one's larger macro-environment may have little to do with self-censorship. Instead, micro-environment sentiments—such as worrying that expressing unpopular views will isolate and alienate people from their friends, family, and neighbors—may be the driver of self-censorship.

There can be little doubt that Americans today are deeply divided on their ideological and partisan attachments, many issue preferences, and even their values.¹ Indeed, these divisions can go as far as to extend to the question of whom—or what kind of person—their children should marry!² A concomitant of these divisions is that political discourse has become coarse, abrasive, divisive, and intense. When it comes to politics today, it is increasingly likely that even an innocent but misspoken opinion will cause a kerfuffle to break out.

It therefore should not be unexpected to find that a large segment of the American people engages in self-censorship when it comes to expressing its views. Self-censorship we define as “intentionally and voluntarily withholding information from others in [the] absence of formal obstacles.”³ In an influential study, MacKuen refers to this more simply as deciding to “talk” or “clam.”⁴ Among social scientists of many stripes, self-censorship has long been a topic of considerable interest and concern. Earlier studies of self-censorship have taken many forms, ranging from philosophical inquiries⁵ to cases studies⁶ to studies of those deciding to withhold crucial evidence of human rights abuses⁷ to studies of self-censorship among racial minorities.⁸ Undoubtedly, the most theoretically rich and empirically comprehensive recent study of self-censorship at the micro-level is Carlson and Settle’s 2022 work.⁹

In a nationally representative survey we conducted in 2020 (see Online Appendix A in the online appendices that accompany this article), we asked a question about self-censorship that Stouffer first put to the American people in 1954. The question asks: “What about you personally? Do you or don't you feel as free to speak your mind as you used to?” While we readily acknowledge that there are a number of potential frailties with this item, its utility is that

the same question has been repeated over a number of surveys between 1954 and 2020 (Online Appendix C addresses several potential threats to the validity of the indicator, concluding generally that, like many, if not most, analyses of change in public opinion over time, the value of investigating how responses to the query have evolved exceeds the limitations of the question). The 2020 results are sobering: fully 46 percent of the American people today reported being less free to speak their minds than they used to.¹⁰

In comparison, a widely noticed contemporaneous 2020 poll by the Cato Institute produced considerably higher estimate of levels of self-censorship, using responses to the statement that “The political climate these days prevents me from saying things I believe because others might find them offensive.” The author concludes: “Nearly two-thirds—62%—of Americans say the political climate these days prevents them from saying things they believe because others might find them offensive. The share of Americans who self-censor has risen several points since 2017 when 58% of Americans agreed with this statement.”¹¹ While, it is difficult to know precisely how to compare the results using the Stouffer question with those from the Cato formulation, both surveys indicate that self-censorship is widespread in the contemporary United States. (See Online Appendix C for some empirical evidence on the comparison.) That so many Americans withhold their views is remarkable—and portentous.

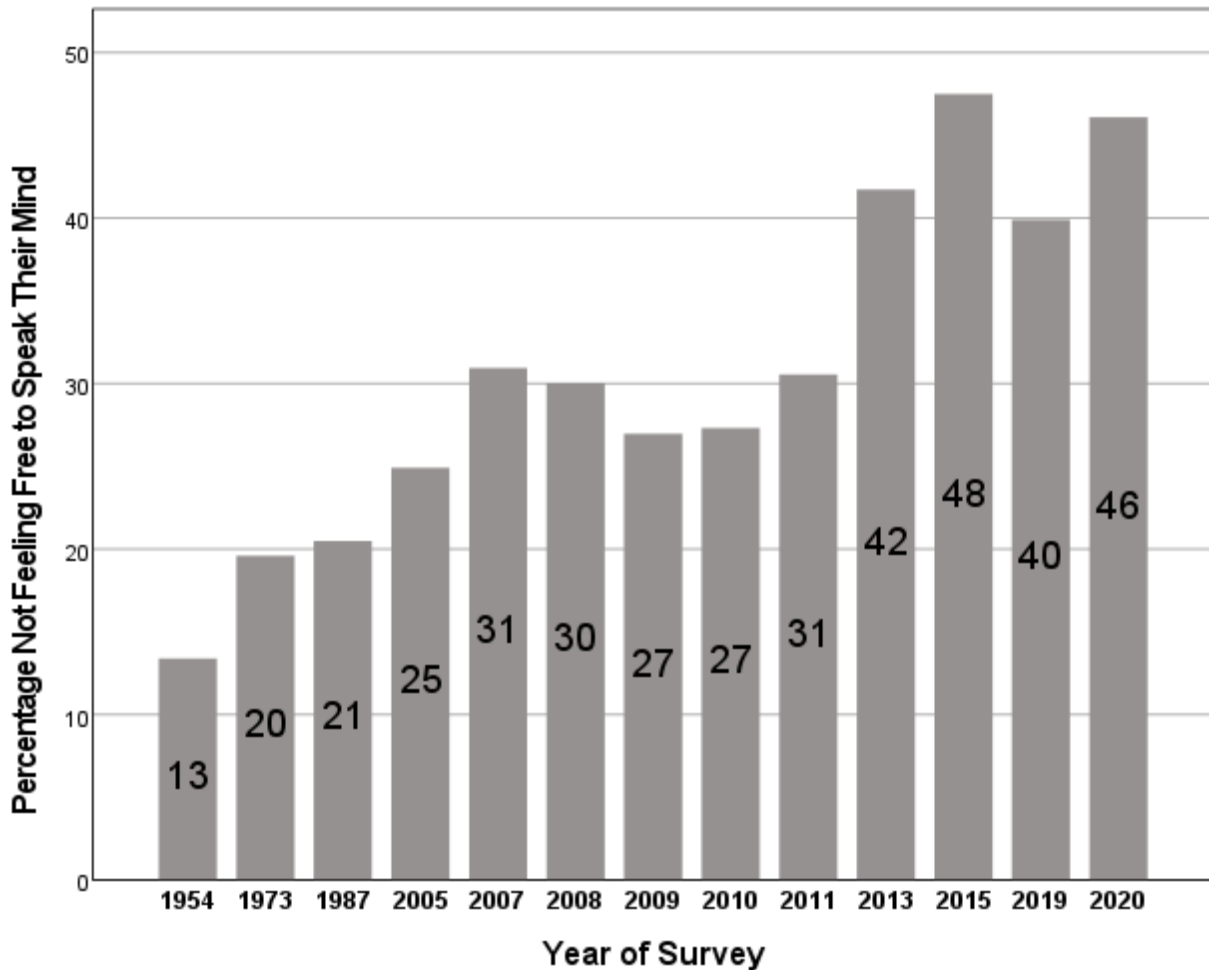
What may be even more unnerving, however, is the steady change in levels of self-censorship over the past 70 or so years. At the height of the 1950s Red Scare, when a circumstantial miscue could land a person in jail, Stouffer (1955) discovered that only 13.4 percent of the American people felt less free to “speak their mind” than they used to; fully 84.7

percent said they did feel free to speak their minds.¹² During an era in which many perhaps *should have been* fearful of speaking their minds, in fact very few Americans seemed to have personally felt the heavy hand of the Republican Senator from Wisconsin and his many allies and followers.

Because the same Stouffer question has been asked in a number of surveys, change in levels of perceived political freedom can be tracked over time. Figure 1 reports the percentages of people who do not feel free to speak their minds, from 1954 to 2020 (see Online Appendix B for the sources). Using a collection of surveys in which the Stouffer question was asked, these data reveal an overall steady erosion of levels of perceived freedom, although there seems to be two points of significant decline—2005 to 2007, and 2011 to 2013. Since about 2013, levels of unwillingness to speak one’s mind have remained fairly constant, with more than four in ten Americans feeling unfree. Public opinion on most matters is slow, perhaps even very slow, to change. But perceptions of the unavailability of individual-level perceived freedom rose rather dramatically from 1954 to 2020. While some might understand these data to indicate that those with “bad” views are no longer free to express themselves, which may be a good thing, we have no means of discerning whether the speech lost is “good” or “bad” speech. Owing to the benefits of deliberations among citizens for democratic politics, most democratic theorists would regard these results as too important to ignore.¹³

FIGURE 1.

Change in Levels of Unwillingness to Speak One's Mind, 1954—2020



Notes:

The Stouffer question asks: “What about you personally? Do you or don't you feel as free to speak your mind as you used to?”

See Online Appendix B for a discussion of the sources of these surveys. Note that this graph is based on available survey data; readers should carefully note the “Year of Survey” on the X-axis in the graph.

What accounts for this remarkable loss of perceived freedom in the United States? How is it that four in ten of the American people do not feel free to express themselves today? Is this loss of free speech a function of fear of being misunderstood by friends and colleagues, or are the causes more systemic, such as government surveillance of social media, telephone, and email discussions?¹⁴ Is the explanation associated with a culture of “political correctness” that many conservatives rail against, or is the source even more elementary, reflecting little more than growing political polarization and incivility as well as rising political intolerance in the country?

Unfortunately, little recent research has investigated the causes and consequences of changes in levels of political freedom reported by the American people, so we are largely in the dark as to why such a large proportion of Americans do not feel free to express their opinions. Earlier studies have examined various aspects of self-censorship but practically none has investigated change over time.¹⁵ As Wells et al. noted in 2017: “There is little evidence—and no record in the literature—of widespread breaking down of talk over political disagreement.”¹⁶ Similarly, Carlson and Settle call for more research on how tendencies toward self-censorship may have changed over time.¹⁷ Fortunately, our data allow a response to these calls of researchers to investigate change in levels of self-censorship.

Our purpose in this article is to explore several hypotheses about the correlates of self-censorship, both at the aggregate and individual levels. Our analysis here is assuredly not comprehensive or definitive, but, in light of the presumed importance of unbridled political discourse for the health of democracies, our findings raise many troubling issues for American democracy. Our most imperative objective in this article is to use these provocative results to

spur additional research on why people seem to have learned that keeping their mouths shut is the best thing to do.

To be clear at the onset, our analysis makes few claims to causal certitude in the relationships it investigates. We feel reasonably (but not perfectly) comfortable in drawing causal conclusions when two variables are unrelated (“X does not cause Y”). When the direction of the relationship is contrary to our stated expectations, we obviously reject the original causal process underlying the hypothesis. Our cross-sectional analysis is particularly vulnerable to causal doubt (although most demographic attributes are unlikely to be consequences of political attitudes, for most people). We contend that determining what goes with what—and what does *not* go with what—is a valuable first step in understanding how and why people engage in self-censorship.

H₁: Affective Polarization

Carlson and Settle conjecture that change in levels of self-censorship may be related to increasing levels of polarization.¹⁸ The trend exhibited in Figure 1 looks familiar and suggests the obvious hypothesis that self-censorship is related to growing affective political polarization. Because people so dislike each other and detest each other’s views and values, they may perceive a great cost associated with sharing their opinions publicly, for little or no reward (on a cost/benefit approach to the decision to express oneself see Menzner and Traunmüller).¹⁹ Therefore, they keep their mouths shut and refrain from expressing their true opinions.²⁰ Perhaps a “saving grace” is that people tend to live in politically homogenous silos or “echo chambers,”²¹

and therefore are most likely to encounter mainly like-minded folks. Nevertheless, unless one can completely isolate oneself from the toxic political environment of contemporary America, it could be prudent to withhold one's views, at least in certain contexts. Free speech has never been free; but the cost of such speech today seems to have skyrocketed—and, to some, the cost may have become exorbitant and out-of-reach.


Figure 2 adds to Figure 1 the trends in partisan and ideological affective polarization among the American people (as measured by Rogowski and Sutherland).²² (See Online Appendix D for a discussion of how we reconciled the varying year data points from different surveys.) Note that owing to the large gap between the data points in 1954 and 1973, we start this data series in 1973.

Clearly, there is no one-to-one relationship between levels of affective polarization and levels of self-censorship, but, also clearly, there are substantial inter-connections over time, whatever the causal link between these two concepts. As polarization has increased, so too has self-censorship. It seems unlikely to us that levels of self-censorship have caused levels of polarization because we cannot see how silence—“keeping one's mouth shut”—would contribute to increases in political differences and animosity. While grounded in a small number of data points, the relationship is actually quite strong: Based on filling in yearly scores using linear interpolation, the correlation of levels of self-censorship with ideological polarization is .91 (N = 48), with partisan polarization, $r = .86$ (N = 43). It seems that one concomitant of rising affective polarization may have been the loss of considerable quantities of freedom to speak.

It is also possible to consider the polarization hypothesis at the micro-level using 2020 data from the Freedom and Tolerance Surveys (FATS)—although the hypothesis makes a bit more sense to us as a measure of the degree of polarization of one’s environment rather than the degree of polarization of oneself. Following the literature, we created measures of affective ideological and partisan polarization as the absolute values of the difference in affect between “Conservatives” and “Liberals,” and between “Republicans” and “Democrats” (rescaled to vary from 0 to 1). The measures indicate that partisan polarization is somewhat stronger than ideological polarization (which is in part a function of the typical finding that there are about twice as many “moderates” as “independent independents”). The two types of polarization are not synonymous, but they are fairly strongly correlated ($r = .63$). We note, however, that the vast literature on ecological fallacies allows for the possibility that macro- and micro-level relationships can be and often are quite different. Indeed, the “fallacy” is specifically the drawing of micro-level conclusions from macro-level findings.

FIGURE 2.

The Relationship Between Self-Censorship and Affective Polarization



Notes:

“Self-Censorship” is the percentage of people not feeling free to speak their minds.

The polarization measures are taken from Rogowski and Sutherland 2016.²³ “Ideological Affective Polarization” is the absolute value of the average difference in thermometer scores assigned to “Liberals” and “Conservatives,” with higher scores indicating greater degrees of affective polarization. “Partisan Affective Polarization” is the absolute value of the average difference in thermometer scores assigned to “Democrats” and “Republicans,” with higher scores indicating greater degrees of affective polarization.

To facilitate interpretation, the graph interpolates the data points between each survey, for each variable (see Online Appendix D).

Note that owing to the availability of data, the time series begins in 1973.

Neither measure of affective polarization is significantly correlated with self-censorship at the micro-level (both correlations are on the order of $-.04$ or so). Those who are themselves more polarized do not withhold their views more. There may be several reasons for the lack of relationship, but the best explanations most likely have to do with how friends and associates of people *react to* polarization, rather than the degree of one's own polarization itself. That the relationship is insignificant also suggests that some who hold polarized views might seek out like-minded friends and therefore feel free to express themselves ("echo chambers"), but that others holding such views do not (or are not able to do so) and therefore do not feel free. This heterogeneity could well produce correlations near zero. Note, however, that this lack of relationship pertains to the 2020 survey, a period of relatively high affective polarization. The available data do not allow us to say how this relationship has changed over time at the micro-level.

In 2019, Ripley, Tenjarla, and He published an analysis of "partisan prejudice" in the United States.²⁴ Based on an application of the statistical technique "multi-level regression with post-stratification" of survey data, the measure they created for each county in the country is to some considerable degree an indicator of affective partisan polarization (although the authors sometimes refer to this as a measure of "political tolerance"—for the methodology and survey questions used in creating the measure, see Online Appendix E). While scholars of political tolerance would not at all accept their measure as indicating the willingness to extend civil liberties to objectionable groups and ideas, using their indicator, we can test the hypothesis that there is more self-censorship in counties in which partisan polarization is higher. For this

analysis, we again focus on the 2020 FATS micro-level data.

Once more, we find no relationship whatsoever between the degree of partisan polarization in the county and levels of individual self-censorship. This we suspect is a function of the county being a much too encompassing measure of the “environment” of the individual. For example, San Bernardino County, the largest county in the United States, has a population of about 2.2 million people and covers about 20,105 square miles.

Whatever causes people to withhold their views no doubt takes place at a much more intimate level than at the level of the county. Variability in the general climate of opinion over time seems to matter, but the across-time national level contextual variability seems not to be manifest in between-county differences in 2020 or so.

H₂: Support for Civil Liberties

Polarization is sometimes narratively associated with unwillingness to recognize the civil liberties of all, especially one’s despicable opponents, and therefore a reasonable hypothesis to account for the dramatic increase in self-censorship in the United States is that people feel less free because reluctance to grant civil liberties to everyone has increased in the country. As it turns out, that does not seem to be so.

Since 1972, the General Social Survey (GSS) has asked in its nationally representative surveys three civil liberties questions (that is, queries about three different activities) with regard to various groups presumed to be widely unpopular with the American people. Five groups were included in virtually all the surveys. One group (Muslims: “a Muslim clergyman who preaches

hatred of the United States”) was added in 2008 and has been included since then. In light of our interest in longitudinal analysis, we have focused on the five groups that were asked over roughly the entire time series.²⁵


We have created an index of group-based support for civil liberties for each of the surveys/years in the GSS time series. The index summarizes the responses to the five groups and three activities asked about: it is the average response to fifteen civil liberties items (after scoring the responses to each item as either (0) Supportive or (1) Not). The yearly averages are reported in Figure 3, which adds levels of civil liberties opposition to the self-censorship data shown in Figure 1. We also carry over from Figure 2 the measures of affective polarization.

Before analyzing the GSS data, we must acknowledge a significant debate about how political tolerance ought to be measured and that many if not most scholars reject the view that the GSS measures political intolerance.²⁶ At a minimum, however, the GSS measures willingness to extend full liberties to these five groups. We therefore will typically refer to this concept as opposition to civil liberties for all.

The GSS evidence indicates a fairly steady and reasonably significant *decrease* over the last several decades in unwillingness to recognize civil liberties for all in the United States. From the data in Figure 3, it is obvious that *decreasing* levels of opposition to civil liberties cannot easily account for *increasing* levels of unwillingness to speak one’s mind. The simple hypothesis that personal political freedom has been lost to a growth in mass unwillingness to grant civil liberties (intolerance) to these groups and groups like these can be clearly rejected.

FIGURE 3.

Self-Censorship, Opposition to Civil Liberties, and Ideological Affective Polarization, Over Time



Notes:

See the notes to Figure 2.

“Oppose Civil Liberties” is the average number of “rejection of civil liberties” responses to fifteen questions (three activities for five groups).

To facilitate interpretation, the graph interpolates the data points between each survey, for each variable (see Online Appendix D).

On the other hand, it may be that national trends are of little consequence because individuals are affected by the intolerance of their local communities, not the nation as a whole. Fortunately, Claassen and Gibson have provided *least-liked* measures of political intolerance in the metropolitan areas of the United States.²⁷ These data can be linked to the location of the 2020 survey respondents and used to test the hypothesis that individuals living in relatively more intolerant communities are more likely to engage in self-censorship. Although there are some important caveats (for example, partisan polarization is measured at the county-level while political tolerance is measured at the MSA-level), we note that the correlation of polarization and tolerance is $+0.12$ ($p < .001$, $N = 1096$), which indicates that more tolerant counties also tend to be more polarized by partisanship.

Using the 2020 FATS, little relationship exists between community intolerance and self-censorship ($r = .07$, $p = .056$, $N = 711$). From this analysis, it seems entirely clear that a community's political intolerance is not a factor driving people to engage in self-censorship, most likely because, for most, the metropolitan area is a much too large definition of the "environment" of the individual. In contrast to defining one's context as a large metro area, Gibson (1995) offers a more detailed analysis of the ways in which local community norms and values can impinge upon the political liberty of individual citizens.²⁸ The answers to the empirical puzzle of increasing and widespread self-censorship must be found elsewhere.

Finally, following Gibson²⁹, we consider whether people who are more politically intolerant are more likely to engage in self-censorship. Using a "least-liked" measure of intolerance, the data reveal only a weak correlation between intolerance and self-censorship ($r =$

.06). However, we defer further consideration of this relationship to the multivariate analysis reported below.

H₃: Perceived Political Repression

It appears that the intolerance of *ordinary citizens* as a whole is not an important factor influencing self-censorship. Perhaps, instead, it is perceptions that the *government* would attempt to prohibit certain types of speech. Since perceptions of government restrictions on free speech were measured in the 2020 survey, it is possible to test the hypothesis that those who view the government as more repressive are more likely to keep their mouths shut. In essence, this hypothesis posits that one source of self-censorship is fear of governmental efforts to quash unpopular speech. While it is beyond the scope of this paper to address all the ways in which governments might be perceived as a significant threat to free speech (and without taking a normative position on the issue), we note in passing the efforts in Georgia of the state to prosecute individuals for the lyrics in the rap songs they write.³⁰ We have little doubt that such artists might respond that they feel less free today than they used to.

Following Gibson, we asked the respondents whether they thought the government would allow them to engage in certain types of political activity (see Online Appendix E).³¹ We have calculated a simple index of the number of activities (of three) that the respondent believes the government would not allow. In 2020, approximately 68 percent of the American people said they felt the government would prohibit none of these activities; about 13 percent said the government would prohibit all these forms of political action.

The hypothesis that those perceiving more government repression engage in more self-censorship receives support, albeit only limited, from the 2020 data. Those (relatively few) asserting that the government would allow none of the activities engaged in significantly more self-censorship. We need not make any inference here about the direction of the causal flow. We only conclude that people who self-censor more tend, to some degree, to view their governments as unwilling to allow various form of political expression. It may be that perceptions of government repression cause self-censorship, but it may also be that unwillingness to express oneself produces a tendency to see the government as repressive. (Our own view is that the former relationship is more plausible than the latter.) Still, even with a dichotomous measure of perceived repression (no activities allowed versus some allowed), the relationship in the 2020 survey to keeping one's mouth shut is weak ($r = .13$), although it is statistically significant ($p < .001$). Thus, perceived government constraints on individual freedom provide a significant but nevertheless modest contribution toward understanding political self-censorship.

H4: Demographic Correlates of Self-Censorship and Multivariate Analysis


It is instructive to consider some plausible relationships between self-censorship and various demographic attributes of the respondents, especially inasmuch as some of the results are counterintuitive. In the FATS 2020 data, the strongest relationship with self-censorship is ideological self-identifications (see Table 1). At the bivariate level, the strongest correlation we observe is also between self-censorship and ideological self-identifications. That relationship is reported in Figure 4.

Figure 4 depicts a significant but not especially strong relationship between ideological self-identification and self-censorship, with conservatives being more likely to withhold their views. That the statistical relationship is not stronger is in part a function of the relatively small number of respondents claiming an “extreme” ideological view (as shown in Figure 4). Still, “Liberals” and “Conservatives” differ by more than 15 percentage points, with about one-third of liberals engaging in self-censorship while more than one-half of conservatives censor themselves.

This bivariate relationship is also reproduced in the multivariate analysis reported in Table 1, with ideological self-identification being the best predictor of self-censorship. (Because the intolerance of the respondent’s metropolitan statistical area bears no relationship to individual-level self-censorship and because including it in the multivariate analysis results in significant (list-wise) missing data, we have excluded that measure from further consideration.) Moreover, the strength of ideological identifications is also significantly related to self-censorship, as is ideological affective polarization, although neither relationship is particularly strong (and neither is as strong as ideology itself). Thus, the clearest conclusion from the analysis here is that, like many things in American politics, ideological divisions significantly structure self-censorship. While we have no data speaking to what sorts of views are withheld, we would not be surprised to learn that conservatives self-censor on a wider range of issues than liberals, perhaps being reluctant to express themselves on matters of race (for example, “woke-ism”), gender and sexuality (including “me-too-ism”), MAGA-ism, and perhaps many other issues as well.

FIGURE 4.

The Relationship Between Ideological Self-Identification and Self-Censorship, 2020



Notes:

Authors' analysis of FATS 2020 data.

At the same time, it is noteworthy that self-censorship is not much associated with partisanship. Whatever bivariate relationship exists is completely eliminated in the multivariate analysis, largely in deference to the association with ideological self-identification. Furthermore, neither the strength of partisan attachments nor partisan affective polarization is connected to self-censorship.

Two other variables are significantly related to self-censorship: Perceptions of repression and the respondent's own political intolerance ("least-liked"). As already noted, those who perceive more repression feel a little less free. More intolerant people also tend to perceive less freedom available to themselves (just as Gibson observed), perhaps under some theory of reciprocity ("just as I am denied freedom, I would deny freedom to others").³²

The analysis also reveals some relationships of no significance at all. Self-censorship is no more common in metro areas than in non-metro areas, which is perhaps a bit surprising since homogeneity, such as that found in many rural communities, is often associated with pressures for conformity. This finding is also unexpected because, as Gibson notes in an earlier analysis of perceived political freedom: "Rural and smaller communities are more likely to have norms since they are more likely to be homogeneous and consensual. Communities with deep divisions of opinion may be normless in the sense that advocates of many points of view can find adherents and defenders."³³ In addition, self-censorship is entirely unrelated to gender, which

TABLE 1.

Predictors of Feeling Free to Speak Your Mind, Logit Results, 2020

Predictor	b	AME	s.e. (b)	p	r
Community Partisan Polarization	-.51	-.12	.46	.274	.01
Political Intolerance	.98	.22	.30	.001	.06
Perceived Repression	-1.30	-.29	.34	.001	-.12
Level of Education	.38	.09	.33	.251	.08
Live in Metro Area	.03	.01	.22	.894	.03
Live in the South	-.02	-.01	.16	.884	-.01
Whether Not White	.17	.04	.18	.359	.04
Whether Owns Home	.20	.04	.20	.316	.04
Age	.18	.04	.39	.645	.01
Gender	.02	.01	.16	.880	-.00
Household Income	.50	.11	.33	.134	.10
Whether Born Again	.23	.05	.20	.261	-.02
Conservative Ideological Identification	-1.64	-.37	.42	.001	-.16
Republican Party Identification	-.04	-.01	.33	.904	-.12
Strong Ideology	.68	.15	.29	.019	.03
Strong Partisanship	-.03	-.01	.26	.900	.01
Ideological Affective Polarization	-1.01	-.23	.36	.005	-.06
Partisan Affective Polarization	-.20	-.04	.31	.530	-.05
Opinion Leader	.05	.01	.15	.757	.01
Political Knowledge	-.53	-.12	.35	.132	.02
Church Attendance	.09	.02	.28	.742	.00
Intercept	.80		.52	.122	

Authors' analysis of FATS 2020 data.

Notes:

N = 792

Equation:

Nagelkerke $R^2 = .11$; -2 Log Likelihood = 1024.9; Percent Correctly Classified = 61%

The AME column reports average marginal effects for each variable. Values are the change in the probability of feeling free to speak your mind in response to a one unit change in the variable. Since all variables are scaled from 0 to 1, an AME represents the change from the lower extreme of the variable to the upper extreme of the variable. AMEs use the distribution of realizations in the sample to set the values at which other variables are held.

perhaps runs counter to the findings of Karpowitz, Mendelberg, and Shaker.³⁴ Nor is social class connected to withholding one's opinions. And the better educated feel no freer than the poorly educated, just as the relationship with identifying as being "born again" is also trivial.

Table 1 also reports average marginal effects (AME) for each predictor. These coefficients are interpretable as the predicted percentage point change in the probability of self-censorship given an increment in the variable. No surprises here: the AMEs demonstrate consistency with the effects implied by the significant coefficients discussed above. The average marginal effect reported for conservative ideological identification suggests a 37-point difference in the probability of feeling free between extreme liberals and extreme conservatives. This seems to be a substantively large difference—which we examined in Figure 4—but the coefficient in part reflects the fact that the numbers of extreme liberals and extreme conservatives in the FATS sample are relatively small.

Perhaps the most important conclusion of this limited analysis of the correlates is that self-censorship is most grounded in ideological identifications, with conservatives perceiving greater constraints on their ability to express themselves. Liberals certainly perceive some restrictions on their freedom but not at the same level as conservatives.³⁵

Summary and Discussion

We readily accept—and perhaps even embrace—the charge that this article raises as many questions as it resolves. Frankly, we did not expect to find that self-censorship is as widespread as it is in the United States, and we were even more nonplused to discover how it has increased over time. That macro affective polarization is perhaps part of the explanation of our data seems quite reasonable, but our failure to draw a connection between unwillingness to recognize civil liberties for all and self-censorship is, while not unprecedented in the literature, a fecund and not-entirely-expected finding. While self-censorship has risen rather dramatically over the past few decades, anti-civil liberties sentiment has actually declined, suggesting that the general political culture may not be the primary source of encouragement to keep one’s mouth shut. Nor is a more fine-tuned measure of intolerance in one’s local community any more useful as a predictor of self-censorship. With the exception of the all-pervasive and much discussed climate of political polarization, one’s macro political environment seems not to be the main source of pressures to censor one’s views.

It is also noteworthy and perhaps unexpected that those who engage in self-censorship are not those with the most limited political resources: self-censorship is no more common among the poorly educated and less knowledgeable than it is among the better educated and more informed. And to the extent that level of education is a sign of greater political sophistication and more extensive socialization to dominant social norms, we might have expected the better educated to engage in more self-censorship. That seems not to be so.

This analysis has not been very successful in discovering which “deviant” political views

are most subject to social disapproval. We note that Carlson and Settle find that the topic under consideration actually has little effect on tendencies toward self-censorship.³⁶ In our data, those who describe themselves as “born again” do not engage in more self-censorship; nor do strong partisans (although we find a weak relationship among strong ideologues). It seems likely that what constitutes unacceptable political views is highly idiosyncratic to specific contexts. For instance, in certain circumstances, those who favor abortion keep their mouths shut; in other contexts, those who oppose abortion keep their mouths shut. Context-dependent patterns are generally difficult to discern and analyze with survey data such as these.

These findings have important potential consequences for the health of democracy in the United States. As Mutz has written: “Exposure to dissimilar views has been deemed a central element—if not the sine qua non—of the kind of political dialogue that is needed to maintain a democratic citizenry.”³⁷ But if those holding dissimilar views keep their mouths shut, then the unbridled discourse about which Mutz writes may become a thing of the past.

Next Steps: A Spiral of Silence?

In Noelle-Neumann’s terms, self-censorship is the driving mechanism in the “Spiral of Silence.”³⁸ In a nutshell, the theory hypothesizes that “opinion climate perceptions are related to the willingness to speak out.”³⁹ That is, the fear of being socially isolated owing to holding minority viewpoints encourages people to keep their mouths shut. Over time, people test whether their views are acceptable to others; when they find they are not, they shut up. Self-censorship makes it even less likely that positive reinforcement for minority viewpoints will be encountered

by others. Without reinforcement, more people remain quiet, with the consequence that orthodox views are established and ascend into a perhaps undeserved domination.

The Spiral of Silence theory is not a theory of government repression; rather, it is a social theory. It is not the officials who are the source of disapproval; rather, it is one's friends, family, and neighbors. From their impressive meta-analysis of research on the Spiral of Silence theory, Matthes, Knoll, and von Sikorski conclude that:

We also found that the type of the target of opinion expression matters for the silencing effect. Interestingly, expressing dissenting views to family, friends, or neighbors has a stronger effect as compared with strangers, politicians, or the media. ... The explanation for this novel finding can be found in the social accountability argument put forth by Mutz. The "need for social accountability creates anxiety because interpersonal disagreement threatens social relationships, and there is no way to please all members of one's network and thus assure social harmony." When voicing dissent to close others, social accountability is arguably higher compared with voicing disagreement to journalists, politicians, or strangers. Thus, citizens suppress their views to maintain social harmony. In terms of the spiral of silence, this means that disagreement with friends and family may lead to a higher state of fear of social isolation as compared with strangers.⁴⁰

That fear, we hypothesize, is what drives self-censorship. For a useful depiction of how the decision to self-censor is made see the scenario described by Carlson and Settle.⁴¹

Investigating the influence of social norms on people is by no means an easy task,

especially if placed within the dynamic context of the Spiral of Silence. Not everyone is aware of norms; not everyone acquiesces to norms in the same way; the sanctions for violating norms are experienced differently by different people; and so on. Designing a study capable of providing an understanding of these processes is difficult and demanding, to say the least.

Our analysis has only scratched the surface in terms of beginning to understand the causes and consequences of self-censorship. Refusing to express what one thinks are unpopular points-of-view may be good for smoothing social and familial relationships, but is, we submit, inimical to a vibrant democracy. Finding ways to lower the costs of expressing dissenting views should be a priority for those who favor the liberal democratic form of governance. After all, dissent is the essence of democracy. As Chief Justice Earl Warren wrote: “Mere unorthodoxy or dissent from the prevailing mores is not to be condemned. The absence of such voices would be a symptom of grave illness to our society.”

Endnotes

¹ James N. Druckman and Matthew S. Levendusky, “What Do We Measure When We Measure Affective Polarization?” *Public Opinion Quarterly* 83 (2019): 114–122.

² Shanto Iyengar, Yphtach Lelkes, Matthew Levendusky, Neil Malhotra, and Sean J. Westwood, “The Origins and Consequences of Affective Polarization in the United States,” *Annual Review of Political Science* 22 (2019): 129-146.

³ Keren Sharvit, Daniel Bar-Tal, Boaz Hameiri, Anat Zafran, Eldad Shahar, and Amiram Raviv, “Self-Censorship Orientation: Scale Development, Correlates and Outcomes,” *Journal of Social and Political Psychology* 6 (2018): 331–363, 331.

⁴ Michael MacKuen, “Speaking of Politics: Individual Conversational Choice, Public Opinion, and the Prospects for Deliberative Democracy.” In John A. Ferejohn and James H. Kuklinski, eds., *Information and Democratic Processes*. (Urbana, IL: University of Illinois Press, 1990), 59-99.

⁵ Matthew Festenstein, “Self-Censorship for Democrats.” *European Journal of Political Theory* 17 (2019): 324-342.

⁶ Chris Wells, Katherine J. Cramer, Michael W. Wagner, German Alvarez, Lewis A. Friedland, Dhavan V. Shah, Leticia Bode, Stephanie Edgerly, Itay Gabay, and Charles Franklin, “When We Stop Talking Politics: The Maintenance and Closing of Conversation in Contentious Times.” *Journal of Communication* 67 (2017): 131-157.

⁷ Daniel Bar-Tal, “Self-Censorship as a Socio-Political-Psychology Phenomenon: Conception and Research,” *Advances in Political Psychology* 38 (2017): 37-65.

⁸ James L. Gibson, “Being Free in Obama’s America: Racial Differences in Perceptions of Constraints on Political Action,” *Daedalus* 141 (2012): 114-129.

⁹ Taylor N. Carlson and Jaime E. Settle, *What Goes Without Saying: Navigating Political Discussion in America* (New York: Cambridge University Press, 2022).

¹⁰ Samuel C. Stouffer, *Communism, Conformity and Civil Liberties* (New York: Doubleday, 1955).

¹¹ Emily Ekins, “Poll: 62% of Americans Say They Have Political Views They’re Afraid to Share.” 22 July 2020, accessed at <https://www.cato.org/survey-reports/poll-62-americans-say-they-have-political-views-theyre-afraid-share#introduction>, 22 July 2020.

¹² Stouffer, *Communism, Conformity and Civil Liberties*. See also James L. Gibson, “Intolerance and Political Repression in the United States: A Half-Century After McCarthyism,” *American Journal of Political Science* 52 (2008): 96-108, 99, and James L. Gibson, “The Political Consequences of Intolerance: Cultural Conformity and Political Freedom,” *American Political Science Review* 86 (1992): 338-356.

¹³ Diana C. Mutz, “Cross-Cutting Social Networks: Testing Democratic Theory in Practice,” *American Political Science Review* 96 (2002): 111-126, and James N. Druckman, “Political

Preference Formation: Competition, Deliberation, and the (Ir) Relevance of Framing Effects,”
American Political Science Review 98 (2004): 671–686.

¹⁴ Rachel Levinson-Waldman, Harsha Panduranga, and Faiza Patel, “Social Media Surveillance by the U.S. Government: A Growing and Unregulated Trend of Online Surveillance Raises Concerns For Civil Rights And Liberties,” 7 January 2022, accessed at <https://www.brennancenter.org/our-work/research-reports/social-media-surveillance-us-government>, 7 January 2022.

¹⁵ See for examples Gibson, “The Political Consequences of Intolerance”; Gibson, “Intolerance and Political Repression;” Jaime E. Settle and Taylor N. Carlson, “Opting Out of Political Discussions,” *Political Communication* 36 (2019): 476-496; and Carlson and Settle, *What Goes Without Saying*.

¹⁶ Wells et al., “When We Stop Talking Politics,” 135.

¹⁷ Carlson and Settle, *What Goes Without Saying*, 254-255.

¹⁸ Carlson and Settle, *What Goes Without Saying*, 255.

¹⁹ Jan Menzner and Richard Traunmüller, “Subjective Freedom of Speech: Why Do Citizens Think They Cannot Speak Freely?” *Politische Vierteljahresschrift* (2022): 1-27.

²⁰ Pew Research Center, “More Now Say It’s ‘Stressful’ to Discuss Politics With People They Disagree With,” 5 November 2018, accessed at <https://www.people-press.org/2018/11/05/more-now-say-its-stressful-to-discuss-politics-with-people-they-disagree-with/> 29 February 2020.

²¹ Kathleen Hall Jamieson and Joseph N. Cappella, *Echo Chamber: Rush Limbaugh and the Conservative Media Establishment* (Oxford: Oxford University Press, 2008).

²² Jon C. Rogowski and Joseph L. Sutherland, “How Ideology Fuels Affective Polarization,” *Political Behavior* 38 (2016): 485-508. See also Iyengar et al., “The Origins and Consequences of Affective Polarization.”

²³ Rogowski and Sutherland, “Affective Polarization.”

²⁴ Amanda Ripley, Rekha Tenjarla, and Angela He, “The Geography of Partisan Prejudice,” 4 March 2019, accessed at <https://www.theatlantic.com/politics/archive/2019/03/us-counties-vary-their-degree-partisan-prejudice/583072/>, 14 March 2022.

²⁵ For examples of over-time analysis of the GSS data see Dennis Chong and Morris Levy, “Competing Norms of Free Expression and Political Tolerance,” *Social Research* 85, no. 1 (2018): 197-227; Joshua T. Davis and Samuel L. Perry, “White Christian Nationalism and Relative Political Tolerance for Racists,” *Social Problems* 68 (2021): 513-534; and Dennis Chong, Jack Citrin, and Morris Levy, “The Realignment of Political Tolerance in the United States,” *Perspectives on Politics* (2022), forthcoming.

²⁶ James L. Gibson, “Measuring Political Tolerance and General Support for Pro-Civil Liberties Policies: Notes, Evidence, and Cautions,” *Public Opinion Quarterly* 77 (2013): 45-68.

²⁷ Christopher Claassen and James L. Gibson, “Does Intolerance Dampen Dissent? Macro-Tolerance and Protest in American Metropolitan Areas,” *Political Behavior* 41 (2019): 165-185.

²⁸ James L. Gibson, “The Political Freedom of African Americans: A Contextual Analysis of Racial Attitudes, Political Tolerance, and Individual Liberty,” *Political Geography* 14 (1995): 571-599.

²⁹ Gibson, “Intolerance and Political Repression.”

³⁰ Briana Carter, “Lyrics for Lockups: Using Rap Lyrics to Prosecute in America,” *Mercer Law Review* 69 (2017): 917-946. See also Austin Vining, “Rap as a Proxy for Blackness: How the Prosecution of Rap Lyrics Continues to Unconstitutionally Restrict Free Speech Rights,” *University of Florida Journal of Law and Public Policy* 31 (2020): 77-95.

³¹ Gibson, “Intolerance and Political Repression.”

³² Gibson, “Intolerance and Political Repression.”

³³ Gibson, “The Political Freedom of African Americans,” 574.

³⁴ Christopher F. Karpowitz, Tali Mendelberg, and Lee Shaker, “Gender Inequality in Deliberative Participation,” *American Political Science Review* 106 (2012): 533-554.

³⁵ We tested for a number of interactions among the variables reported in Table 1, focusing especially on interactions with gender and race, but found none of any magnitude or substantive interest.

³⁶ Carlson and Settle, *What Goes Without Saying*, 125, see also Table 5.3, page 127.

³⁷ Mutz, “Cross-Cutting Social Networks,” 122.

³⁸ Elisabeth Noelle-Neumann, *The Spiral of Silence: Public Opinion, Our Social Skin* (Chicago, IL: The University of Chicago Press, 1993).

³⁹ Jörg Matthes, Johannes Knoll, and Christian von Sikorski, “The ‘Spiral of Silence’ Revisited: A Meta-Analysis on the Relationship Between Perceptions of Opinion Support and Political Opinion Expression,” *Communication Research* 45 (2018): 3–33, 22.

⁴⁰ Matthes, Knoll, and von Sikorski, “The ‘Spiral of Silence’ Revisited,” 23.

⁴¹ Carlson and Settle, *What Goes Without Saying*, 109-110.

ONLINE APPENDICES:

**KEEPING YOUR MOUTH SHUT:
SPIRALING SELF-CENSORSHIP IN THE UNITED STATES**
Version 4

Online Appendix A: The 2020 AmeriSpeak Freedom and Tolerance Survey (FATS)

Funded and operated by NORC at the University of Chicago, **AmeriSpeak®** is a probability-based panel designed to be representative of the U.S. household population. Randomly selected U.S. households are sampled using area probability and address-based sampling, with a known, non-zero probability of selection from the NORC National Sample Frame. These sampled households are then contacted by U.S. mail, telephone, and field interviewers (face-to-face). The panel provides sample coverage of approximately 97 percent of the U.S. household population. Those excluded from the sample include people with P.O. Box only addresses, some addresses not listed in the U.S.P.S. Delivery Sequence File, and some newly constructed dwellings. While most AmeriSpeak households participate in surveys by web, non-internet households can participate in AmeriSpeak surveys by telephone. Households without conventional internet access but having web access via smartphones are allowed to participate in AmeriSpeak surveys by web. AmeriSpeak panelists participate in NORC studies or studies conducted by NORC on behalf of governmental agencies, academic researchers, and media and commercial organizations.

A general population sample of U.S. adults aged 18 and older was selected from NORC's AmeriSpeak Panel for this study. This survey was offered only in English and was administered on the web and over the phone. Invitations to participate in the survey were initiated on July 1, 2020, and the last interviews were completed on July 24, 2020. In total, NORC collected 1,006 interviews, 950 by web mode and 56 by phone mode.

To encourage study cooperation, NORC sent five email reminders to sampled web-mode respondents. Panelists were offered the cash equivalent of \$5 for completing the study. Interviewed respondents took 29 minutes (median) to complete the survey. The interview was divided into two modules, with the FATS questions asked first. The 29 minutes median is the total length of interview. NORC applied standard cleaning rules to the survey data for quality control by removing responses in the main study interview to questions from non-eligible respondents. These respondents provided responses indicative of speeding through the survey and skipping survey questions. These respondents were not included in the final dataset.

The data are weighted, with various factors going into the construction of the final study weight. These include panel base sampling weights, final panel weights, study-specific base sampling weights, and nonresponse adjusted survey weights.

The weighted AAPOR Response Rate #3 recruitment rate was 23.6 percent, with a weighted household retention rate of 84.8 percent and a survey completion rate of 28.4 percent. A weighted AAPOR Response Rate #3 cumulative response rate of 5.7 percent was achieved. The survey has a margin of error of 4.17 percent, and a design effect of 1.82.

For additional technical information about the sample or the study, email AmeriSpeak-BD@norc.org or visit www.AmeriSpeak.norc.org.

This research was approved by the Washington University in St. Louis Institutional Review Board (IRB). That IRB judged this project to be in the “exempt” category owing to the

fact that participation in the survey was voluntary, no harm was afflicted on the respondents, and no identifiers were connected to the database generated, among other factors.

Online Appendix B: The Surveys Used in the Time Series

The data included in the time series reported in Figure 1 are derived from: Stouffer's 1954 nationally representative, face-to-face survey¹; the Nunn, Crockett, and Williams 1973 replication of the Stouffer survey (1978)²; the 1987 General Social Survey, as supplemented by the Gibson re-interview³; the 2005—2011 Freedom and Tolerance Surveys (FATS), with the 2005 survey being the U.S. part of the 2005 European Social Survey⁴; (5) the 2013—2015 American Panel Surveys (TAPS: <https://wc.wustl.edu/american-panel-survey>); and (6) the 2019 and 2020 AmeriSpeak Freedom and Tolerance Surveys (documented in Online Appendix A, above).

As with any compilation of independent studies, these data are not, strictly speaking, directly comparable owing to different modes of survey administration (and other factors). (For a similar approach to assembling various surveys within a single database, see Enns.)⁵ The Stouffer survey, the 1987 GSS/Gibson survey, and the 2005 ESS/US survey were conducted face-to-face. The FATS surveys from 2007 through 2009 were based on RDD samples. The FATS surveys in 2010 and 2011 used RDD samples supplemented with cell phone samples. The TAPS surveys were internet surveys. The 2019 and 2020 surveys were mixed mode surveys (telephone and internet) with internet administration of the instrument very strongly dominant. Care must therefore be taken not to make strict and precise comparisons across the surveys.

Nevertheless, the trend in these data is completely obvious and is strong enough to defeat any explanation of the trend in the data based on survey mode.

¹ Samuel C. Stouffer, *Communism, Conformity and Civil Liberties* (New York: Doubleday, 1955).

² Clyde Z. Nunn, Harry J. Crockett, and J. Allen Williams, *Tolerance for Non-Conformity* (San Francisco: Jossey-Bass, 1978).

³ James L. Gibson, "The Structure of Attitudinal Tolerance in the United States," *British Journal of Political Science* 19 (1989): 562-570.

⁴ James L. Gibson, "Measuring Political Tolerance and General Support for Pro-Civil Liberties Policies: Notes, Evidence, and Cautions," *Public Opinion Quarterly* 77 (2013): 45-68.

⁵ Peter K. Enns, *Incarceration Nation: How the United States Became the Most Punitive Democracy in the World* (New York, NY: Cambridge University Press, 2016).

Appendix C: The Measurement of Self-Censorship

Several potential threats to the validity of the Stouffer measure of self-censorship can be imagined. Here we address several of those.

The Use of a Single-Item Indicator: Why the Stouffer Measure?

Single-item indicators are, of course, used in a wide variety of research, especially when policy preferences or institutional attitudes are of interest (for example, death penalty, confidence in institutions such as the U.S. Supreme Court, presidential popularity, satisfaction with democracy). When it comes to measures of self-censorship, single-item indicators are much more common than multi-item indicators.

Still, we readily acknowledge that using a single-item indicator to measure self-censorship introduces some amount of measurement error (for a comprehensive set of measures of perceived political freedom see Gibson 1992).⁶ Since a central purpose of this research is to measure change in self-censorship—and since the Stouffer survey, conducted during the McCarthy Red Scare, is both one of the oldest and most highly regarded analyses of political intolerance in the United States—we accept the limitations of the measure in order to analyze how responses to the item have evolved over time. This is a tradeoff, to be sure, but it is the only feasible way of getting purchase on the central question of change in public opinion.

Defects in the Question Wording?

Three complaints might be imagined: the question does not explicitly refer to “political speech,” the comparison to “as you used to” is vague, and no response options for “more free” was offered by the question.

Political Speech

The first issue we address is whether the Stouffer item measures perceived freedom to speak *about politics*. Obviously, the question itself does not mention politics so some might consider that the responses are overly broad in the sense that they conflate and combine political and non-political speech.

We investigated that possibility with a question-wording experiment fielded in a nationally representative survey in September 2022. The survey utilized NORC’s AmeriSpeak Panel, one of the highest quality probability-based surveys available. (The 2019 and 2020 FATS surveys were also conducted via NORC’s AmeriSpeak panel.) Respondents were randomly assigned to hear one of three versions of the “freedom to speak” question. The first version is the original item asked by Stouffer in 1954 (the item reported in the time series in this paper). The second version, simply adds to the Stouffer question the preface “When it comes to politics . . .” The third version was constructed by the Cato Institute and was used in widely noticed 2017/2020 national surveys. The CATO item has traditionally revealed more self-censorship

⁶ James L. Gibson, “The Political Consequences of Intolerance: Cultural Conformity and Political Freedom,” *American Political Science Review* 86 (1992): 338-356.

than the Stouffer item. We begin by considering the null hypothesis that the version without a reference to politics (“All Speech”) produces the same results as the version with an explicit reference to politics (“Political Speech”). The results are shown in Table A.1.

Clearly, the data indicate that the Stouffer version is understood as asking, *de facto*, about political speech: the differences in the responses between the question delimited to politics and the question silent on what sort of speech is being asked about are absolutely trivial (see the first two columns in the table). We conclude therefore that the Stouffer item is not over-inclusive and is in fact measuring speech relevant to politics.

“as you used to” Comparison

As for the “as you used to” comparison, we first note that a recent survey in Germany also found it useful to use the past as a referent for assessing contemporary freedom. Menzner and Traunmüller measured self-censorship with the following item: “People like me are no longer allowed to express their opinions freely in public.”⁷

It is also commonplace for survey researchers to ask “as you used to” questions. For example, Pew has asked:

Since Donald Trump's (2016) election, would you say you are paying more, less, or about the same amount of attention to politics as you used to?

And:

Since you started reading newspapers online, are you reading the paper version of the newspaper more often, less often, or about as much as you used to?

CBS News has asked:

In general, because of the economic recession, when you go shopping, are you spending more money on things other than the basic necessities--I don't mean essential food and clothes, but the extras--or less money on these things, or about the same amount of money as you used to?

Fox News has asked:

And what about you--are you personally more grateful for what you have these days or angry you don't have as much as you used to have?

Despite imprecision in knowing exactly what period the respondents are using as the baseline, as a measure of contemporary trajectory, the question should be understood as asking whether “things are getting better or getting worse.” That seems to us to be a quite valid approach.

⁷ Jan Menzner and Richard Traunmüller, “Subjective Freedom of Speech: Why Do Citizens Think They Cannot Speak Freely?” *Politische Vierteljahresschrift* (2022): 1-27.

TABLE A.1.

Alternative Measures of Perceived Freedom to Speak, 2022

Response	<i>All Speech</i>	<i>Political Speech</i>	<i>Prevents Me from Saying</i>
Free to Speak	42.9	41.1	32.5
Not Free to Speak	57.1	58.9	67.5
Total	100.0	100.0	100.0
N	630	345	346

Source: NORC AmeriSpeak 2022

Notes:

Table entries are percentages (except for N).

All Speech: What about you personally? Do you or don't you feel as free to speak your mind as you used to?

Yes, do feel as free

No, feel less free

Political Speech: What about you personally? When it comes to politics, do you or don't you feel as free to speak your mind as you used to?

Yes, do feel as free

No, feel less free

Prevents Me from Saying: The political climate these days prevents me from saying things I believe because others might find them offensive.

Strongly agree

Somewhat agree

Somewhat disagree

Strongly disagree

Source: NORC AmeriSpeak 2022

No “more free” Option Offered

We are not much concerned about this issue. First, we expect that not many would say that they are “more free” (especially in light of the Cato findings that we cite). Second, the question in essence asks, “are you less free, yes or no?” We readily concede that those who say they are not less free include two types of people: those for whom there has been no change in their perceived level of freedom and those who see themselves as having more freedom. Our analysis focuses on describing and understanding how the replies of “less free,” which we contend are *not* ambivalent, have changed over time and how they vary across individuals and communities. So, the heterogeneity in the “not less free” category is not of much concern to us.

Convergent Validity

How do the responses to the Stouffer item stack up against other measures of self-censorship?

The Cato Measure

The Cato measure of self-censorship—“The political climate these days prevents me from saying things I believe because others might find them offensive”—seems not to suffer from any obvious question-wording infirmities. It refers to speech in a political context, identifies the source of constraints on speech as being external to the respondent, and is a contemporaneous measure (not requiring any comparison to the past). As noted, a portion of the AmeriSpeak sample was assigned (randomly) to respond to this item.

The data in Table A.1 (above) reveal that in 2022 the Stouffer item provides more conservative estimates of self-censorship than the Cato measure (which was also true in the 2020 survey), with the Stouffer measure registering about 10 percentage points less self-censorship than the Cato measure. In its 2017 survey, Cato found that 58 percent of Americans engaged in self-censorship; in its 2020 survey, the number had risen to 62 percent. Still, the alternative measures of self-censorship all support that conclusion that a large proportion of the American people do not believe they enjoy unrestricted freedom of speech.

Finally, it should be noted that the 2022 data indicate for both the Stouffer version and the Cato question considerably more self-censorship in 2022 than in 2020.

The 1987 General Social Survey

We also endeavored to validate the Stouffer measure with some items that were first used in a 1987 GSS reinterview survey (reported by Gibson 1993) pertaining to reluctance to speak one’s mind. Figure A.1 reports the relationship between the responses in the 2022 survey to the Stouffer item and these measures of hesitancy about speaking. The statistic reported in the graph is the difference in the percentage of those who self-censor according to the Stouffer measure who express reluctance to speak out for various reasons minus the percentage for those who do not self-censor. Note that our expectations are that the Stouffer responses will be related to perceived external constraints on freedom speech but not to perceived internal constraints.

The figure reports interesting and expected relationships and non-relationships. Regarding the former, those who currently feel less free to speak their minds are also much more likely than those who do feel free to express their views to assert that talking about politics

creates enemies, that they worry about what people might think of them, that they do not like arguments, and that they expect that others would judge their views to be strange. All of these seems to be beliefs that reinforce or explain their perceptions of a lack of freedom by references to external constraints.

At the time, however, there is little or no difference between those feeling unfree and those not feeling unfree on beliefs that they are insufficiently knowledgeable to speak, that their associates do not care about politics, or that they hold similar views to their associates, and for other reasons as well. In general, those feeling less free identify the expected external constraints on their freedom to speak and fail to identify the expected internal constraints on their freedom, providing at least some confidence in the validity of the Stouffer question.

Our 2022 validation survey also repeated a question from the 1987 GSS on the degree to which people are worried about expressing their political views to others. The question asked about being worried to express their views in a variety of contexts, ranging from publicly in their community to their immediate family. The 1987/2022 question read:

How worried are you about expressing your political views to. . . ?

Members of your immediate family

Your close friends

Your co-workers

Members of organizations to which you belong

Publicly, in your community

Your representatives in the government

The response set was:

You worry quite a bit about what they will think of your views

You worry only some about what they will think of your views


You don't worry very much about what they will think of your views

You don't worry at all about what they will think of your views

We created an index of the extent of being worried across the six contexts by scoring those worried "quite a bit" and "only some" as worried and those who "don't worry very much" and "don't worry at all" as not worried. Regarding the Stouffer item, those feeling free to express themselves had an average of 1.1 contexts in which they were worried, whereas those feeling less free worried about expressing themselves, on average, in 1.9 contexts. The difference between the two Stouffer groups is statistically significant at $p < .001$ (and the same results are produced using different summary measures of worries).

FIGURE A.1.

Correlates of Self-Censorship, 2022



Notes:

The table entries are the differences in the percentages of people who self-censor who regard the statement about reluctance to speak as true as pertains to them minus the percentage of those who do not self-censor.

The question read:

Some people have told us that they are occasionally reluctant to talk about politics with their families and friends. I would like to read you several statements and ask if they are true or false as they apply to you.

Creates enemies^{***}: I am sometimes reluctant to talk about politics because it creates enemies.
I lack knowledge: I am sometimes reluctant to talk about politics because I lack the information and knowledge to do so.
People would think my views strange^{***}: because people would think my political views were strange.
Few friends care about politics: because few of my friends and family care about politics.
I hold the same views as others: because I usually hold the same political views as those I am around.
I don't like arguments^{***}: I am sometimes reluctant to talk about politics because I don't like arguments.
I worry about what people think of me^{***}: because I worry about what people would think of me.
I worry that the government would find out: because I worry that the government might find out about me.

^{***} The difference in proportions between those feeling free and those not feeling free is statistically significant at $p \leq .001$
95 percent confidence intervals are shown for each difference of percentages.

Source: NORC AmeriSpeak 2022

Reliability

Finally, the reliability of the responses to single-item indicators is more difficult to assess. In such circumstances, test/re-test reliability is the most useful approach. In the 2005 FATS survey, a subsample of respondents was reinterviewed twice. The original interview was face-to-face; the follow-up interviews were conducted by telephone. Of the 1,001 respondents in the t_1 survey, 257 were reinterviewed at t_3 . The t_3 interview took place approximately one year after the t_1 interview. The Stouffer question was asked at t_1 and t_3 (but not at t_2). For 64 percent of the respondents, the same answer was given at t_1 and t_3 . A large majority of those saying they felt free at t_1 also said they felt free at t_3 . For those reporting not feeling free at t_1 , about one-half said they did not feel free at t_3 . These results are somewhat challenging to assess because such a small percentage of the t_1 respondents completed the t_3 interview (in part because completing the t_2 interview was a requirement for being reinterviewed at t_3), because of the shift in interview mode, and because a year elapsed between the t_1 and t_3 interviews. Still, that approaching two-thirds of the respondents gave the same response at both interviews suggests a reasonable degree of reliability in the measure.

Returning to the time-series data, we also observe that the pattern of change over time in levels of self-censorship seems relatively smooth, with considerable similarity in the results from surveys conducted in close temporal proximity to each other. Such a pattern does not seem

compatible with the hypothesis that the Stouffer measure is contaminated with a great deal of random error (a reliability matter).

Are Responses Driven by Contexts?

Every analyst much face the possibility that the context of questions and answers changes over time (even over fairly short periods of time). Most long-term repeated cross-sectional surveys (for example, the GSS, ANES) are biased in favor of not updating question wording even when the wording seems to have become less apposite to current conditions (for example, GSS's fixed-group tolerance questions). For instance, the question "should a communist be allowed to give a speech" no doubt is embedded in different contexts when asked in 1954 and when asked in 2022 (as is true of all the groups about which the GSS asks its tolerance questions.⁸ The same could be said about longitudinal analyses of racial attitudes, or, for that matter, virtually any issue (for example, health care, abortion, affirmative action, the death penalty). Even the composition of "liberals" and "conservatives" has undoubtedly changed over time.

Our view is that the description of public opinion and the analysis of the etiology of public opinion must be treated as different questions. We have no doubt that one possible explanation of the rise of self-censorship in the United States revolves around the growth of social media usage (although one should be careful not to over-estimate the importance of social media, especially for middle-aged and older Americans). Widely used available time-series measures—such as Stimson's measures of "mood" or Enns' measures of punitiveness—face the same issues.⁹ Indeed, the shifting definitions of concepts are one reason why it is important to use the same question wording over time.

Conclusions

Generally, these various empirical findings from the 2022 survey tend to support the view that self-censorship is being validity and reliably measured by the Stouffer item. We certainly do not maintain that the Stouffer item is a perfect summary measure of self-censorship, but we do contend that there is sufficient evidence of the psychometric properties of the item to warrant using the measure in this paper and to support the substantive conclusions we draw.

Our purposes in this paper are to describe change in levels of self-censorship in the United States, to provide some highly exploratory correlates (and non-correlates) of that change, to point to unanswered questions requiring further research, and to suggest a theoretical framework for such additional investigations. Focusing on the Stouffer item therefore seems to be quite reasonable. In the final analysis, we contend that our evidence of change over time is sufficiently dramatic that it most likely cannot be attributed to temporal frailties in the survey question.

⁸ Dennis Chong and Morris Levy, "Competing Norms of Free Expression and Political Tolerance," *Social Research* 85, no. 1 (2018): 197-227.

⁹ James A. Stimson, *Public Opinion in America: Moods, Cycles, and Swings*. 2nd Edition (New York: Routledge, 2019); Enns, *Incarceration Nation*.

Online Appendix D: The Time-Series Database

As we have noted in the text, our collection of various yearly time series does not overlap very well when we match the data by the exact year. For example, ideological affective polarization and partisan affective polarization are simultaneously available in only 15 years (for which the correlation is .70). Because our objective in this analysis is nothing more than to offer a rough comparison of trends in different time series, we have created a full data series in which scores between years are linearly interpolated.¹⁰ This database was used to create Figures 2 and 3 in the text. Across the time period from 1954 to 2020, the correlations of the variables shown in the graphs are (number of cases available in parentheses):

Self-Censorship	1.00				
Civil Liberties Support	-.73 (45)	1.00			
Ideological Affective Polarization	.91 (48)	-.65 (48)	1.00		
Partisan Affective Polarization	.86 (43)	-.82 (43)	.77 (43)	1.00	
Year	.85 (48)	-.95 (45)	.74 (48)	.94 (43)	1.00

Some might object to this database because the yearly coverage of surveys in the period prior to 2005 is sparse indeed. We therefore report the correlation matrix for these variables for the period from 2005 to 2020 (N = 16).

Self-Censorship	1.00				
Civil Liberties Support	-.92	1.00			
Ideological Affective Polarization	.89	-.97	1.00		
Partisan Affective Polarization	.70	-.77	.87	1.00	
Year	.87	-.93	.92	.73	1.00

Because the conclusions that we seek to draw from this analysis are minimalist, simply related to trends over time, and because these correlations are so strong and unambiguous and so little affected by the time period to which they are constrained, we have confidence that our figures do not mis-represent our data.

¹⁰ Avram Sidi, *Practical Extrapolation Methods: Theory and Applications* (New York: Cambridge University Press, 2003).

Online Appendix E: Measurement and Question Wording

Perceived Political Repression

The question read:

Suppose you felt very strongly that something the government was doing was very wrong and you wanted to do something about it. Do you think the government would definitely allow, probably allow, probably not allow, or definitely not allow you to . . .

Make a speech in public criticizing the actions of the government?

Organize public meetings to oppose the actions of the government?

Organize protest marches and demonstrations to oppose the actions of the government?

Community-Level Political Intolerance

We employ a measure of intolerance (“least-liked”) in the respondent’s metropolitan statistical area (MSA) developed by Claassen and Gibson.¹¹ The index of intolerance at the MSA level was computed via multi-level regression with post-stratification. For details, see the Claassen and Gibson article. Note that respondents not living in an MSA do not receive a community-level intolerance score.

Individual-Level Political Intolerance

At the individual-level, the measure of intolerance is a “least-liked” measure.¹² After identifying the most disliked groups, the respondents were asked:

[THE DISLIKED GROUP] should be allowed to make a speech in our community] ... Please indicate whether you agree or disagree with each of the following statements.

Strongly agree

Agree

Neither agree nor disagree

Disagree

Strongly disagree

[THE DISLIKED GROUP] should be banned from running for public office] ... Please indicate whether you agree or disagree with each of the following statements.

Strongly agree

Agree

¹¹ Christopher Claassen and James L. Gibson, “Does Intolerance Dampen Dissent? Macro-Tolerance and Protest in American Metropolitan Areas,” *Political Behavior* 41 (2019): 165-185.

¹² James L. Gibson, “Measuring Political Tolerance and General Support for Pro-Civil Liberties Policies: Notes, Evidence, and Cautions,” *Public Opinion Quarterly* 77 (2013): 45-68.

Neither agree nor disagree
Disagree
Strongly disagree

[THE DISLIKED GROUP] should be allowed to hold public rallies and demonstrations in our community] ... Please indicate whether you agree or disagree with each of the following statements.

Strongly agree
Agree
Neither agree nor disagree
Disagree
Strongly disagree

Community Partisan Polarization

According to the company that did the analysis, they used Random Device Engagement as their sampling technique to get 2,000 survey responses from around the country as well as the respondents' geo-coordinates. The survey consisted of 14 questions (see below) that were designed to measure partisan prejudice along with demographic information and party identification. The respondents were matched to the full voter file by their geo-coordinate history. They employed “highly evolved variants” of multi-level regression and post-stratification which enabled them to use their small sample to get estimates for wider areas. They “model the outcome of interest (here: political tolerance) based on urbanicity based on home address, age, gender, education, household composition, race, party affiliation, and two variables we use to describe the neighborhood: age variation and variation in partisan identification at the census block where the individual resides.” The model can give “estimates of political tolerance for millions of demographic combinations” and is “powerful because every single response is used to train all parameters.” Lastly, they weighted the estimates for the demographics “by the fraction of the demographic of interest in the target population” to get each county’s score. The full details of their methodology are posted on their website:

<https://www.predictwise.com/blog/the-atlantics-the-geography-of-partisan-prejudice-method-addendum>.

The items used in the MRP to estimate county-level partisan polarization are:

1. How would you react if a member of your immediate family married a Democrat?
2. How would you react if a member of your immediate family married a Republican?
3. How well does the term “Patriotic” describe Democrats?
4. How well does the term “Selfish” describe Democrats?
5. How well does the term “Willing to compromise” describe Democrats?
6. How well does the term “Compassionate” describe Democrats?

7. How well does the term “Patriotic” describe Republicans?
8. How well does the term “Selfish” describe Republicans?
9. How well does the term “Willing to compromise” describe Republicans?
10. How well does the term “Compassionate” describe Republicans?
11. How do you feel about the Republican Party today?
12. How do you feel about the Democratic Party today?
13. How do you feel about Democratic voters today?
14. How do you feel about Republican voters today?

Demographic Variables

For the following variables, standard measures were used (with the number of categories in the response sets shown in parentheses):

- Level of Education (5)
- Live in Metro Area (2)
- Live in the South (2)
- Whether Not White (2)
- Whether Owns Home (2)
- Age (numerous)
- Gender (2)
- Household Income (9)
- Whether Born Again (2)
- Conservative Ideological Identification (7)
- Republican Party Identification (7)
- Strong Ideology (4, folded Conservative Ideological Identification)
- Strong Partisanship (4, folded Republican Party Identification)
- Ideological Affective Polarization (numerous, (ABS(Conservative Affect – Liberal Affect)))
- Feeling Thermometer toward Conservatives (101)
- Feeling Thermometer toward Liberals (101)
- Partisan Affective Polarization (numerous, (ABS(Republicans Affect – Democrats Affect)))
- Feeling Thermometer toward Republicans (101)
- Feeling Thermometer toward Democrats (101)
- Opinion Leader (2)
- Political Knowledge (10)
- Church Attendance (9)

Table E.1. The Distributions of the Variables Reported in Table 1

Variable	Mean	Standard Deviation
Self-Censorship	.54	.50
Community Partisan Polarization	.44	.18
Political Intolerance	.51	.28
Perceived Repression	.24	.28
Level of Education	.53	.29
Live in Metro Area	.82	.39
Live in the South	.37	.48
Whether Not White	.35	.48
Whether Owns Home	.66	.48
Age	.39	.23
Gender	.49	.50
Household Income	.57	.29
Whether Born Again	.27	.45
Conservative Ideological Identification	.51	.26
Republican Party Identification	.48	.32
Strong Ideology	.39	.34
Strong Partisanship	.55	.34
Ideological Affective Polarization	.46	.31
Partisan Affective Polarization	.39	.31
Opinion Leader	.51	.50
Political Knowledge	.73	.29
Church Attendance	.34	.33

Notes:

All variables are scored to range between 0 and 1.

N = 792