

Online Appendix

Throwing Away the Umbrella

A Preclearance Definition

Our definition of “covered” counties (those previously subject to preclearance under Section 4 of the VRA) is drawn largely from a [list](#) provided by the Department of Justice.

We include all counties in fully-covered states as covered, as well as the individual counties included in the DOJ’s list. There are also several townships in Michigan and South Dakota that were covered as of 2013; we conservatively include the counties containing these townships as covered in our county-level analyses, though some jurisdictions in these counties were not covered.

In the case of jurisdictions in Virginia and New Hampshire that had “bailed out” of coverage by 2013, we continue to include them as covered here if they bailed out after the year 2003. Many of these bailouts occurred in the decade immediately preceding the *Shelby* decision, meaning that in many ways officials would still need to act as if they were covered (the decade-long “recapture period” would allow them to immediately be bailed back in if they did anything that would have prevented a bailout in the first place: see the Department of Justice’s public information about Section 4 [here](#)).

B Validating Catalist data against other datasets

We validated the Catalist data we use in this project by comparing it to several other datasets, in hopes of noticing any strange patterns or major errors. We began with a comparison to Current Population Survey estimates. The CPS is often used to produce estimates of turnout by race at the state level, so we aggregated the Catalist dataset to the state level for comparison. We used state-level estimates of citizen voting age population from the ACS (for 2010-2018) to turn the raw Catalist turnout counts into turnout rates comparable to the ones calculated from CPS data. When calculating CPS turnout rates, we rely on the “cpsvote” R package ([Lee and Gronke, 2020](#)), using its “Hur-Achen” approach to nonresponse and its provided weights to handle over-time changes in response rates.

We note that the CPS is not a perfect source of group-specific turnout estimates and should not be treated as the “ground truth,” but we nevertheless think it is useful to see how the Catalist-derived estimates we produce compare to the CPS ones. The left panel of Figure 6 shows that comparison for state-specific Hispanic turnout estimates from 2010-2018. The Catalist estimates are on the horizontal axis and CPS estimates are on the vertical axis, with the Black diagonal line showing the 45-degree line (along which estimates are exactly the same across the two datasets). Points are scaled by population size (states with larger Hispanic populations appear larger) and shaded by year. These datasets look similar, with points clustered along the 45-degree line. There are some points above and below it, where one source shows much higher turnout than the other, but for the most part these are states with small Hispanic populations (where we expect more measurement error, which is part of why we weight our main estimates by population size). The

years cluster somewhat, as expected (turnout in 2016 was higher than in 2014 almost everywhere), but there is not a clear pattern of one year straying farther from the 45-degree line than others.




Figure 6: Comparing Catalist Hispanic and Black turnout estimates to CPS-derived estimates

The right panel of Figure 6 compares Catalist and CPS estimates of Black turnout. The diagonal line again shows equivalence between the Catalist and CPS estimates, though in this case the axis is stretched out by the presence of a few extreme outliers in the Catalist data. As noted in the main paper, there are a few places where small Black populations combined with measurement error in either the Catalist turnout estimates or the ACS estimates yield impossible turnout estimates of over 100%. The two points on the extreme right side of the plot are estimates from North Dakota, a state with a very small number of estimated Black eligible voters and thus a lot of room for measurement error to influence estimated turnout in fairly extreme ways. Given our population-weighted approach to the main estimates, we do not think counties in ND are likely to exert a large influence over our analyses. The estimates are broadly similar across the two datasets, particularly for places with large Black populations (represented by larger points), though the CPS estimates are on average slightly higher than the Catalist ones (consistent with turnout over-reporting on the CPS, as in [Ansolabehere, Fraga and Schaffner \(2020\)](#)).

Next, we compared our county-level Catalist estimates to estimates from David Leip’s county-level elections data (obtained for 2008-2020 through the MIT library system). Leip’s data reports aggregate registration and turnout counts for each county in each year, not estimates for specific racial groups. Still, we thought it worth summing up our Catalist data to produce county-level estimates of the total number of registered voters and ballots cast for each county year and comparing those to the Leip estimate to diagnose problems.

We were able to merge over 99% of the counties in our main dataset to counties in Leip’s data




Figure 7: Comparing Catalist county turnout/registration estimates to Leip data

using FIPS codes; the main source of missed matches was a difference in how Alaskan counties/-election districts were handled across datasets. The left panel of Figure 7 compares our Catalist turnout estimates to Leip’s, with the diagonal line representing equivalence in the two datasets’ estimates. The datasets have very similar county turnout numbers; slight differences (points off the line) do not appear systematic across years. The right panel of Figure 7 performs the same exercise for county registration numbers. Again, the estimates line up tidily on the 45-degree line for most county-years.

C Additional Analyses of Registration/Turnout

C.1 Hispanic-White Gaps

In this section, we present analyses of the Hispanic-white registration and turnout gap analogous to the Black-white analyses presented in the main paper. We begin, as in the main paper, by simply plotting the trends over time in both the registration and turnout gap in previously-covered and non-covered places in Figure 8. The pre-2013 trends in the registration gap appear broadly similar in covered and non-covered places, but the turnout gap trends do not appear as similar, so we present and interpret difference-in-differences estimates for this measure with caution. (For more formal consideration of these pre-trend issues, see section C.3 below, where a pre-treatment placebo test finds significant effects of *Shelby* on this outcome prior to 2013, which is of course impossible).

Table 5 presents analogous DiD estimates to Table 3 in the main paper. Again, the first four




Figure 8: Time trends in Hispanic-White registration (left) and turnout (right) gaps. Dotted lines represent weighted means for formerly-preclearance counties and solid lines represent weighted means for non-covered counties. Means are weighted by county-level Black citizen voting-age population (CVAP).

columns present estimates of *Shelby* effects in affected places on the registration gap, while the last four columns examine the turnout gap. Comparably to the main paper, we do not see evidence of a widening Hispanic-white registration or turnout gap in affected jurisdictions: most of these point estimates are non-significant, and all of them are positively rather than negatively signed (a negative point estimate would mean that the Hispanic-white gap was widening, with white turnout further outpacing Hispanic participation). We see one statistically significant estimate on the Hispanic-white turnout gap (column 6, when including county-level covariates for precision), but we hesitate to conclude that the Hispanic-white turnout gap has actually *narrowed* in previously-preclearance places. In light of the pre-trend differences shown in Figure 8, we place less weight on these estimates for the Hispanic-white turnout gap than we do on the main-paper estimates about Black-white gaps.

Figure 9 presents dynamic estimates analogous to the main paper’s Figure 4, this time focusing on Hispanic-white registration and turnout gaps. The first two point estimates (prior to 2013) serve as an additional way of considering possible parallel-trends violations, and as mentioned above, we see concerning pre-treatment “effects,” particularly on the turnout gap.

Table 5: Estimated Effects of Removing Preclearance on Participation Gaps (Diff-in-Diff)

	<i>Dependent variable:</i>							
	Hispanic-White Registration Gap				Hispanic-White Turnout Gap			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Preclearance x Shelby	0.00001 (0.011)	0.004 (0.007)	0.018 (0.014)	0.030* (0.014)	0.015 (0.010)	0.017* (0.007)	0.019* (0.009)	0.027* (0.008)
County Fixed Effects	✓	✓	✓	✓	✓	✓	✓	✓
Year Fixed Effects	✓	✓			✓	✓		
County Demographic Controls		✓		✓		✓		✓
State x Year Fixed Effects			✓	✓			✓	✓
Observations	17,167	17,165	17,167	17,165	17,167	17,165	17,167	17,165
Adjusted R ²	0.708	0.735	0.766	0.785	0.718	0.737	0.811	0.818

Note:

* $p < 0.05$

Standard errors clustered by state. Time-varying county-level controls include: total population, population density, proportion male, proportion over age 65, proportion nonwhite, proportion Hispanic, proportion married, proportion foreign-born, proportion high school graduates, and unemployment rate.

C.2 Group-specific Turnout

In this section, rather than considering participation gaps, we simply examine registration and turnout rates by specific groups in the dataset across time and space.³⁰ These group-specific estimates were the main estimates presented in an earlier working paper version of this paper.

We again begin by simply plotting registration and turnout rates over time in covered and non-covered places. Figure 10 plots each group's registration and turnout, weighted by county group populations (such that counties with a large Black population, for example, count more heavily in the Black population trend than those with smaller Black populations).³¹ We also include as the fourth panel for each outcome a look at overall registration and turnout that does not incorporate any racial-classification data from Catalist (it simply looks at total registration and total turnout divided by total CVAP for the jurisdiction). For the most part, these outcomes appear to track closely in the pre-2013 period when comparing covered and non-covered places,

³⁰We focus on white, Black, and Hispanic voters as several large and geographically-dispersed groups of voters. Other groups could certainly be affected by the *Shelby* decision, but we are less sanguine about our ability to identify effects on their behavior using the county-level design in this paper.

³¹Alert readers will notice that registration rates in this dataset appear higher than many other sources would indicate, potentially due to outdated or "deadwood" registrations for people no longer living in the county. This overestimate should not pose a problem for the diff-in-diff setup unless there are specific time-varying geographic differences in registration purge patterns, which are unlikely to occur in a way that would produce positive (as opposed to negative) bias in the estimates. But we would not directly interpret the levels of registration shown here as true registration rates among current residents.




Figure 9: Dynamic estimates for Hispanic-White registration and turnout gaps

which lends plausibility to the parallel trends assumption needed for the difference-in-differences setup³²

Table 6 thus presents simple difference-in-differences estimates of the effect of the *Shelby* decision on group-specific registration and turnout rates in previously-covered places. As with estimates in the main paper, these are weighted by group population and drop counties with tiny (<100) populations, and standard errors are clustered by state. Consistent with the patterns shown in Figure 10 above, these estimates suggest that if anything, registration and turnout increased across the board in previously-covered places. Some of these point estimates are statistically distinguishable from zero and some are not, but all are positively signed and appear to rule out substantial *decreases* in minority voter participation in previously-covered places over the period studied.

³²We note some apparent trend differences in Hispanic registration rates, especially in 2012, which appear to trace back to variation in ACS population estimates during this period. We present estimates on Hispanic registration with caution.

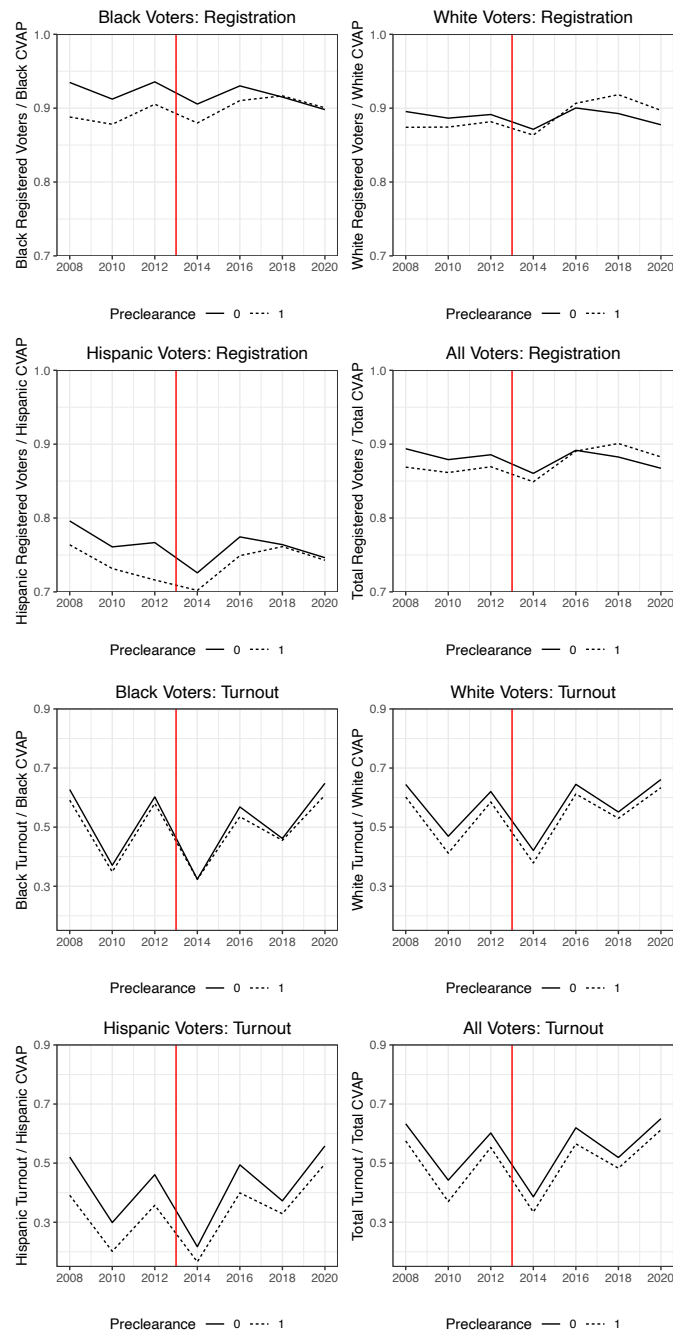


Figure 10: Registration and Turnout Trends (Group-specific)

Table 6: Group-Specific Registration and Turnout Estimates

	<i>Dependent variable:</i>					
	Black Reg.	Hispanic Reg.	White Reg.	Black Turnout	Hispanic Turnout	White Turnout
	(1)	(2)	(3)	(4)	(5)	(6)
Preclearance x Shelby	0.027 (0.014)	0.030* (0.007)	0.029* (0.010)	0.007 (0.007)	0.046* (0.011)	0.031* (0.015)
Observations	15,759	17,167	21,884	15,759	17,167	21,884
R ²	0.846	0.861	0.839	0.917	0.921	0.906
Adjusted R ²	0.816	0.834	0.812	0.901	0.906	0.890

Note:

* $p < 0.05$

C.3 Placebo Tests

One might wonder whether the kinds of estimates shown in the main paper could arise by chance, perhaps due to some other background “treatment” or some systematic issue with the data used. To assess this possibility, we run placebo tests where we implement the main analysis using “placebo” treatment years. We set false decision years for the *Shelby* case in 2009 and 2011 (rather than 2013, as in reality) and report the results of our estimation procedure under these assumptions. We rely on these years because they are the only pre-treatment years for which data is available; including post-treatment years would risk incorporating real effects from any real treatment period. Table 7 presents the resulting estimates: no choice of placebo year produces statistically significant effects on Black-white turnout or registration gaps (the main specification used in the paper), and the estimates vary in direction. Table 8 presents the same exercise for Hispanic-white registration gaps, and here we see some concerning patterns of apparent “pre-treatment effects” on Hispanic-white turnout gaps that should make us interpret the main DiD estimates for this measure with caution (as noted in Section C.1 above).

Table 7: Difference-in-Differences Results for Black-White Registration and Turnout Gaps in Placebo Treatment Years

	<i>Dependent variable:</i>			
	Black-White Registration Gap		Black-White Turnout Gap	
	(1)	(2)	(3)	(4)
Preclearance x Shelby	0.004 (0.009)	−0.001 (0.007)	0.009 (0.008)	0.0004 (0.007)
Placebo Year	2009	2011	2009	2011
Observations	6,528	6,528	6,528	6,528
Adjusted R ²	0.904	0.904	0.856	0.856

Note:

* $p < 0.05$; Standard errors clustered by state.

Table 8: Difference-in-Differences Results for Hispanic-White Registration and Turnout Gaps in Placebo Treatment Years

	<i>Dependent variable:</i>			
	Hispanic-White Registration Gap		Hispanic-White Turnout Gap	
	(1)	(2)	(3)	(4)
Preclearance x Shelby	0.007 (0.018)	−0.004 (0.010)	0.037* (0.016)	−0.016* (0.006)
Placebo Year	2009	2011	2009	2011
Observations	6,955	6,955	6,955	6,955
Adjusted R ²	0.879	0.879	0.831	0.826
<i>Note:</i>		* $p < 0.05$; Standard errors clustered by state.		

C.4 Robustness to Alternative Specifications

C.4.1 County-Specific Trends

Table 9 presents estimates of the *Shelby* decision's effect on Black-white registration and turnout gaps in a specification that allows for time-varying slopes for each county in addition to county and year fixed effects. Including variable slopes by county allows us to capture and account for county-specific time trends in election climates that might affect racial turnout gaps for reasons apart from *Shelby*.

Accounting for county-specific time trends in addition to fixed effects for county and year yields small positive coefficients, non-significant in the case of the registration gap and significant in the case of turnout. The direction of these estimates suggests that if anything, the Black-white registration and turnout gaps in previously-covered places *shrank* in the post-*Shelby* period. These estimates differ slightly from those in the main paper in that they are positively-signed (rather than the insignificant negative coefficients in the main paper) but they continue to bolster our central conclusion that the post-2013 period had not seen substantial reductions in minority participation or political power in previously-covered places.

Table 9: DiD Estimates for Black-White Registration/Turnout Gaps with Varying County Slopes

Dependent Variables: Model:	Black-White Registration Gap (1)	Black-White Turnout Gap (2)
<i>Variables</i>		
Preclearance \times Shelby	0.0059 (0.0044)	0.0153* (0.0077)
<i>Fixed-effects</i>		
County	Yes	Yes
Year	Yes	Yes
<i>Varying Slopes</i>		
County	Yes	Yes
<i>Fit statistics</i>		
Observations	15,759	15,759
R ²	0.94416	0.89681
Within R ²	0.00047	0.00429
<i>Clustered (state)) standard-errors in parentheses</i>		
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>		

C.4.2 State-Level Analyses

Following [Bertrand, Duflo and Mullainathan \(2004\)](#), we further validate our estimates by aggregating to the state level. Table 10 summarizes our main difference-in-differences specifications at the state level. Here, registration and turnout levels are summed over counties within each state and year and divided by corresponding group CVAP in order to generate registration and turnout rates by state. Following our previous analysis, we weight by group population in order to upweight states with large subgroup populations. States designated as preclearance include those states previously under statewide coverage (see Footnote 12 of the main paper); states that contain several covered jurisdictions, but are not covered statewide, are designated as untreated. However, these results are robust to the inclusion of North Carolina as a preclearance state. These estimates are consistent with those we report in the main paper: we do not find evidence of large or statistically-significant increases in Black-white registration or turnout gaps in previously-covered places after the *Shelby* decision.

Table 10: Difference-in-Differences Results for Registration/Turnout Gaps at the State Level

	<i>Dependent variable:</i>	
	Black Reg. Gap	Black Turnout Gap
	(1)	(2)
Preclearance x Shelby	−0.004 (0.006)	−0.004 (0.006)
Observations	349	349
R ²	0.847	0.780
Adjusted R ²	0.817	0.738
Residual Std. Error (df = 292)	19.640	18.752

Note:

* $p < 0.05$
Results weighted by state-group population in 2008.
See footnote 15 for preclearance criteria at the state level.

C.4.3 Different Time Periods

In addition to artificially re-setting treatment to years other than 2013 and finding the anticipated null effects, we run additional checks to ensure robustness over space and time. Figure 11 presents the estimates from an analysis in which we iteratively drop every year in our data, to show that our estimates are not dependent on events or data issues occurring in any one year. Estimates are broadly consistent with the main specification we present in the paper even when excluding any given year, though as noted in the paper the point estimates are even more precisely 0 when excluding 2020.

Further, in Table 11 we subset the main dataset to run separate analyses focused on midterm and presidential elections. These estimates are also consistent with our main specification, if noisier as a result of using fewer observations. As noted in the paper, 2020 appears somewhat different from other years, and accordingly the point estimates for the turnout gap in presidential years (the subset including 2020) appear more negative than those in midterm years, though still not statistically distinguishable from zero.




Figure 11: Difference in Differences Estimates for Dropped Years

C.4.4 Excluding Individual States

Another robustness concern is the possibility that our estimates are driven primarily by outcome changes in a single state, or perhaps by measurement error in one state's data. To investigate this possibility, we iteratively drop states from our analysis in order to examine whether differences in turnout and registration by group remain consistent. Figures 12 show that the results do not depend exclusively on the presence of specific states. Difference-in-differences estimates for registration and turnout gaps change very little when any individual state is excluded.

C.4.5 Weighting

As we discuss in Section 4, our main analysis weights counties by the size of the relevant minority group for which we analyze turnout and registration. Here, we verify that the conclusions we reach are not strictly an artifact of these population weights. We show this, in part, by using raw registration and turnout totals from Catalist in Section E below. In addition to this, we show the results of our main analysis of turnout and registration gaps without weighting in Table 12.³³ These estimates are consistent with our main analyses in the sense that they do not


³³These analyses continue to drop places with extremely small (<100) group population estimates and thus high chances of measurement error; the estimates presented in Section E remove this restriction as

Table 11: Difference-in-Differences Results for Registration/Turnout Gaps in Midterm/Presidential Election Years

	<i>Dependent variable:</i>			
	Black Reg. Gap Midterm Years	Black Turnout Gap	Black Reg. Gap Presidential Years	Black Turnout Gap
	(1)	(2)	(3)	(4)
Preclearance x Shelby	−0.002 (0.008)	0.001 (0.005)	−0.006 (0.011)	−0.017 (0.012)
Observations	6,752	6,752	9,007	9,007
R ²	0.904	0.887	0.846	0.824
Adjusted R ²	0.850	0.824	0.787	0.758
Residual Std. Error	6.011 (df = 4337)	3.373 (df = 4337)	7.068 (df = 6535)	5.099 (df = 6535)

Note:

* $p < 0.05$; Standard errors clustered by state.



(a) Black-white Registration Gaps

(b) Black-white Turnout Gaps

Figure 12: Difference-in-Differences Estimates for Turnout and Registration Gaps Excluding Individual States

show any indication of worsening Black-white registration or turnout gaps in previously-covered well.

places after *Shelby*; rather, they find significant and implausibly-large reductions in Black-white gaps over this period (a positive coefficient here denotes a shrinking Black-white gap). We think these estimates carry questionable assumptions about parallel trends; unweighted registration- and turnout-gap estimates in covered/uncovered places do not track nearly as closely in the pre-period as our weighted measures do, potentially due to strange population-estimate fluctuations in small places. As such, we present them with caution, but we hope to provide transparency into the impacts of our analytic decisions and to illustrate that removing weights from the analysis does not provide evidence that reverses the conclusions of the paper.

Table 12: Difference-in-Differences Results for Registration/Turnout Gaps with Unweighted Data

	Dependent variable:	
	Black Reg. Gap (1)	Black Turnout Gap (2)
Preclearance x Shelby	0.112* (0.021)	0.051* (0.012)
Observations	15,759	15,759
R ²	0.886	0.855
Adjusted R ²	0.865	0.827
Residual Std. Error (df = 13257)	0.244	0.140
Note: *p < 0.05; Standard errors clustered by state.		

C.4.6 The South

The VRA’s original target jurisdictions for preclearance were all states in the Deep South. While the preclearance formula expanded over time, the South’s large Black citizen population and robust history of minority vote suppression rendered it especially subject to federal scrutiny until *Shelby*. In Tables 13 and 14, we examine trends in turnout and registration gaps in the South alone. Following our pre-analysis plan, we use two different definitions. One approach narrowly defines “the South” as the 11 original Confederate states: AL, AR, FL, GA, LA, MS, NC, SC, TN, TX, and VA. Arkansas, Florida, North Carolina, and Tennessee were never preclearance in their entirety, so counties within these states form a comparison group within the southern region. We also use the U.S. Census Bureau’s broader definition of the South, which includes the following states: AL, AR, DE, Washington DC, FL, GA, KY, LA, MD, MS, NC, OK, SC, TN, TX, VA, and WV.

Table 13 relies on the Confederacy definition and Table 14 on the Census definition. These estimates are generally consistent with the paper’s results, showing no change in Black-white registration/turnout gaps in preclearance areas (relative to non-preclearance areas) after *Shelby*. The one exception is the small positive coefficient on the Black-white turnout gap when limiting to the former Confederacy. The positive direction of this estimate suggests that Black turnout growth may have slightly outstripped white increases in participation in previously-covered places over the post-2013 period, though we interpret these estimates with some caution given the smaller sample size (and fewer untreated places) when limiting to the former Confederacy. Nevertheless, this robustness test does not indicate that Black-white turnout gaps widened in previously-

covered places after *Shelby*, even when focusing our attention narrowly on the region of the country most often discussed as the target of the VRA.

Table 13: Difference-in-Differences Results for Registration/Turnout Gaps in Former Confederate States

	<i>Dependent variable:</i>	
	Black Reg. Gap (1)	Black Turnout Gap (2)
Preclearance x Shelby	0.010 (0.011)	0.020* (0.007)
Observations	7,237	7,237
R ²	0.857	0.820
Adjusted R ²	0.832	0.789
Residual Std. Error (df = 6155)	4.864	3.996

Note: * $p < 0.05$; Standard errors clustered by state.

Table 14: Difference-in-Differences Results for Registration/Turnout Gaps in Census South Region

	<i>Dependent variable:</i>	
	Black Reg. Gap (1)	Black Turnout Gap (2)
Preclearance x Shelby	-0.0003 (0.012)	0.009 (0.008)
Observations	8,770	8,770
R ²	0.856	0.809
Adjusted R ²	0.830	0.775
Residual Std. Error (df = 7447)	4.909	3.951

Note: * $p < 0.05$; Standard errors clustered by state.

D Looking for Evidence of Countermobilization

In this section, we look for evidence that efforts to register and mobilize Black and Hispanic voters increased after the *Shelby* decision, or that voters saw more reason to turn out. This question is important both because of a prominent hypothesis that there could be counterbalancing effects in the wake of *Shelby* (that is, effective vote suppression efforts being met with grassroots countermobilization efforts that fend off net decreases in voting), as well as simply providing a chance to consider the role of on-the-ground activism in shaping voter participation. In practice, we encounter substantial data limitations and find relatively inconclusive patterns. We encourage other researchers to continue to seek out new approaches to measuring grassroots voter activism.

There are prominent examples of GOTV efforts explicitly targeted to counter potential voter suppression in the wake of the decision: earlier in the paper, we noted the SPLC’s “Vote Your Voice” campaign and its references to *Shelby*. Similarly, major philanthropic donors gave to the Shelby Response Fund, set up to “support legal, organizing, and public education work focused

on protecting voting rights in states formerly covered under Section 5 of the Voting Rights Act.”³⁴ Though it is difficult to quantify all of the get-out-the-vote efforts of many disparate organizations, it is plausible that they ramped up in the wake of the *Shelby* decision. However, it is difficult to track such efforts systematically across space and time, and recent work has pointed out the limited size of any such expected effect given what is known about GOTV efforts ([Grimmer and Hersh, 2023](#)).

We first turn to survey data to look for evidence of such efforts. The Cooperative Election Study (CES, formerly CCES) is run every two years. In several recent election years (2006-2020, excluding 2008 and 2018), the survey asked whether people had been contacted during the election cycle by a campaign organization or candidate.³⁵ We use this question, combined with information about respondents’ county of residence and self-reported racial identity, to see whether campaigns’ GOTV efforts targeted at voters of color increased in previously-covered places after the *Shelby* decision. This question does not capture all possible mobilizing activity, since it is focused on campaigns and not other groups’ efforts, but it gives a consistent view of mobilization efforts across time and geography.

We present these results with caution, as the CES is designed to be a nationally-representative survey, not to yield precise estimates within small geographic areas or for segments of the population ([Grimmer et al., 2018](#)). It is also difficult to judge whether covered and non-covered places had similar pre-*Shelby* trends, since these questions were asked in only a handful of years before the decision. Still, we present these analyses as a preliminary look at the phenomenon of counter-mobilization. We approximately follow the specification of [Cantoni and Pons \(2019\)](#), though we focus on a “*Shelby v. Holder*” treatment rather than the voter ID laws they considered. Standard errors are clustered at the county level.

We begin by asking whether voters experience different mobilization trends in places that were and were not affected by the *Shelby* decision. Columns 1 and 2 of Table 15 indicates that voters in previously-covered places reported extra campaign mobilization after the *Shelby* decision. These voters saw about three percentage points’ higher rates of campaign contact after the decision than would otherwise have been expected (shown by the interaction between “preclearance” and “post-*Shelby*” in the table).

In the second half of Table 15, we ask whether that additional mobilization was focused on minority voters, as we would expect if it were driven by efforts from groups worried about voting rights. Here, the coefficient of interest is the interaction between “preclearance” (whether a jurisdiction was covered by preclearance before 2013), “Post-*Shelby*” (whether the observation is from before or after the 2013 *Shelby* decision), and “Non-white.” In both specifications, this coefficient is positive, suggesting more mobilization of nonwhite voters in previously-covered places after the *Shelby* decision. This pattern is consistent with a story about countermobilization, though the estimates are noisy and not statistically distinguishable from 0. We also acknowledge [Grim-](#)

³⁴See, for instance, the [MacArthur Foundation](#)

³⁵In most years, this question reads “Did a candidate or political campaign organization contact you during the [X] election”; in 2006, it read “During the November election campaign, did a candidate, party organization, or other organization contact you to get you to vote?”

Table 15: Self-Reported Mobilization (CES)

	<i>Dependent variable:</i>			
	Mobilization			
	(1)	(2)	(3)	(4)
Preclearance	−0.063* (0.022)	−0.070* (0.024)	−0.036* (0.011)	−0.040* (0.011)
Post-Shelby	−0.124* (0.013)	−0.110* (0.015)	−0.120* (0.012)	−0.096* (0.013)
Non-white			−0.006 (0.034)	0.010 (0.031)
Preclearance * Post-Shelby	0.034 (0.010)	0.031* (0.015)	0.027* (0.010)	0.022 (0.017)
Preclearance * Non-white			−0.018 (0.029)	−0.022 (0.025)
Post-Shelby * Non-white			−0.006 (0.021)	−0.052* (0.020)
Preclearance * Post-Shelby * Non-white			0.022 (0.013)	0.034 (0.022)
Constant	0.716* (0.023)	0.686* (0.025)	0.696* (0.012)	0.662* (0.012)
State FE's	X	X		
Year FE's	X	X		
Race-by-state FE's			X	X
Race-by-year FE's			X	X
Survey Weights		X		X
Observations	272,783	272,783	272,783	272,783
R ²	0.044	0.028	0.054	0.038
Adjusted R ²	0.044	0.028	0.054	0.037

Note:

*p<0.05

mer and Hersh (2023)’s point about the importance of considering the magnitude of these effects: taken at face value, they imply about a 5.6-percentage point increase in non-white voters reporting campaign mobilization efforts in previously-covered places after 2013. Making even generous assumptions about the effectiveness of campaign contacts and about how survey responses relate to actual mobilization efforts (imagining, for example, that increased campaign contacts might also be paired with non-campaign mobilization work that reached an even larger number of voters) still implies fairly limited effects of the kinds of mobilization reported here. Such increases in mobilizing contact could perhaps produce increases in participation of somewhere below one percentage point (among targeted population groups), but could not, for example, be offsetting what would otherwise be a large (multiple percentage point) reduction in voting driven by policy changes.

Next, we turn to another survey dataset to look at voters’ perceptions of the electoral system. As noted above, some research in political psychology finds that voters can react strongly to perceived attempts to disenfranchise them. Voters could potentially react to the *Shelby* decision or the electoral changes that followed with backlash, perhaps becoming more likely to vote in the wake of those changes.³⁶ We are not aware of any panel survey that asks voters directly whether they think voting rights are under threat. However, the Survey on the Performance of American Elections (SPAЕ) asks voters whether they believe their vote was counted as intended in the most recent election. Although this is not a question explicitly about voting rights, we expect it to capture respondents’ views about the integrity of the electoral system in their area, which should give an idea of whether they are concerned about voting access for people like them.

The SPAЕ surveys registered voters in a sample of states about their experiences in each federal election (Stewart, 2017). We rely on responses to the question “How confident are you that your vote in the General Election was counted as you intended?” from 2008, 2012, 2014, and 2016.³⁷ We focus on the share of voters reporting that they were “not too confident” or “not at all confident” that their votes were counted as intended; Figure 13 plots this measure over time for all respondents and for white and non-white voters separately. In these simple unweighted plots, it appears respondents in places previously covered by preclearance follow similar trends to those in other places before the *Shelby* decision, but then show much higher rates of skepticism about the electoral process after the decision; this pattern is particularly striking among non-white voters.

Table 16 again presents difference-in-differences estimates for all voters and then considers nonwhite voters specifically. Columns 1 and 2 show that after the 2013 *Shelby* decision, voters

³⁶Unlike the analysis of countermobilization above, which asked whether someone had been explicitly asked to vote, this mechanism could occur even without organizations or campaigns communicating with voters about the threat (if voters simply saw news stories, for example). However, we think it is possible that political organizations help spread the word about possible threats to the franchise.

³⁷The SPAЕ was not run in 2018, and though it resumed in 2020 we omit responses from that year out of concerns that the political environment in that year would lead respondents to interpret the question differently than they had before (that is, as an opportunity to embrace or reject Donald Trump’s claims of election fraud in the presidential election).




Figure 13: SPAE respondents’ lack of confidence in vote counting by race and by home county’s preclearance status.

in previously-covered places became several percentage points more likely to say they doubted their vote was counted as intended. In exactly the places where we might think voters would be turning out due to a sense of foreboding about voting rights, we see more voters expressing a lack of confidence in elections. Columns 3-4 ask whether this pattern is especially pronounced for non-white voters, as we would expect if minority voters were responding to perceived threats to voting rights after the *Shelby* decision. When we include an interaction between previous preclearance status, the post-*Shelby* period, and voter race, we see a pattern consistent with higher rates of concern among minority voters in affected places. The “preclearance * Post-*Shelby* * Non-white” coefficient suggests that nonwhite voters in affected places became several percentage points more likely to say they were not confident about vote-counting after the *Shelby* decision, though this coefficient is somewhat noisily-estimated and cannot be statistically distinguished from 0. Again, we also urge attention to the size of these estimates when thinking through how large a shift in voter participation could potentially be attributed to such voter concerns.

D.1 New Registrations from Catalist (Countermobilization)

In addition to the survey data presented above, we also look for evidence of countermobilization using a dataset constructed by Catalist of new voter registrations recorded in each county over each two-year election cycle from 2008 to 2018. For each election year (presidential and midterm), the dataset uses an aggregate snapshot of the voter file taken shortly after the election to tally up the number of new voter registrations added to the voter file in each county over the previous two years since the prior election.³⁸ For example, a person who moved to Cobb County,

³⁸This time window means that we have election years from 2010 through 2018 in this dataset: the 2010 observation captures new registrations taking place between the 2008 and 2010 elections. Aggregate data does not disaggregate by race.

Table 16: Lack of Confidence that Own Vote Counted Correctly (SPAE)

	<i>Dependent variable:</i>			
	Not Confident Own Vote Counted			
	(1)	(2)	(3)	(4)
Preclearance	0.002 (0.012)	0.004 (0.015)	0.002 (0.012)	0.011 (0.010)
Post-Shelby	0.002 (0.004)	0.004 (0.005)	−0.008 (0.005)	−0.008 (0.005)
Non-white			−0.023 (0.020)	0.008 (0.035)
Preclearance * Post-Shelby	0.019 (0.008)	0.025* (0.008)	0.006 (0.009)	0.007 (0.010)
Preclearance * Non-white			0.005 (0.019)	−0.028 (0.034)
Post-Shelby * Non-white			0.063* (0.013)	0.061* (0.014)
Preclearance * Post-Shelby * Non-white			0.028 (0.017)	0.048* (0.024)
Constant	0.039* (0.010)	0.040* (0.014)	0.045* (0.011)	0.041* (0.009)
State FE's	X	X	X	X
Year FE's	X	X	X	X
Survey Weights		X		X
Observations	37,860	37,860	37,806	37,806
R ²	0.012	0.012	0.018	0.019
Adjusted R ²	0.010	0.011	0.015	0.017

Note:

*p<0.05

Georgia and registered to vote in 2011 would be recorded in the 2012 “new registrations” data for that county, as would a person who had previously lived in the county unregistered but decided to register in summer 2012. These estimates are based not on comparing the total number of registrants in a county at different time points, but on the dates that each individual person’s registration appeared on the voter file.

This dataset should allow us to see whether new registrations increased in previously-preclearance counties after the *Shelby* decision. If voting-focused organizations worked to contact and register unregistered people or to help them update outdated registrations to reflect their current addresses, this dataset should capture the results of those efforts. We note that efforts to contact and turn out already-registered people would not be captured by this dataset.

We run a similar difference-in-differences analysis to the one presented in Section 4 above, but the outcome measure is now the share of voting-eligible residents of a county who appear as newly registered in a given year. As above, we cluster standard errors by state and weight by county population (in this case, total population rather than group-specific estimates, since we do not have new-registrations data by race). Table 17 displays the difference-in-differences estimate from this approach: the point estimate is positive, consistent with new registrations increasing very slightly in previously-covered places after 2013, but is noisily-estimated and cannot be distinguished from zero. It is possible that on-the-ground efforts to help voters register ramped up in counties affected by the *Shelby* decision, but this analysis does not allow us to say with certainty that those efforts occurred or succeeded.

Table 17: New Registrations (from Catalist) as a share of CVAP

	<i>Dependent variable:</i>
	New Registrations Rate
Preclearance x Shelby	0.002 (0.008)
County fixed effects	X
Year fixed effects	X
Observations	15,664
R ²	0.802
Adjusted R ²	0.753
<i>Note:</i>	* $p < 0.05$

E Other Outcomes: Total Registration/Turnout, Partisan Votes

Our main analysis focuses on registration and turnout among voters from specific groups that have faced disenfranchisement and political exclusion. This section looks at a broader measure: what happened to overall registration and turnout in previously-preclearance places after *Shelby*? We focus on raw counts of registrants and voters from two sources: the Catalist dataset used in the main analysis, and county-level data from Leip's election atlas. Using raw counts of registrants and voters makes these estimates slightly harder to interpret, but it also means we are not relying on any additional datasets (such as Census data) to calculate rates or to use as weights. And looking at overall registration and turnout means that we are no longer relying on Catalist's racial classifications of voters.

Table 18 presents difference-in-differences estimates calculated from the Leip data for 2008-2020, while Table 19 presents estimates from the Catalist dataset. The estimates vary slightly in size and are not statistically distinguishable from zero, but they generally point to increases in overall registration and turnout in previously-covered places after *Shelby*, consistent with our main findings and also with those of [Raze \(2021\)](#).

Table 18: Leip Difference-in-Differences Results for Preclearance After Shelby

	<i>Dependent variable:</i>	
	Leip Total Registration (1)	Leip Total Turnout (2)
Preclearance x Shelby	2,615.573 (1,678.148)	1,010.780 (1,108.047)
County fixed effects	X	X
Year fixed effects	X	X
Observations	20,744	21,728
R ²	0.992	0.936
Adjusted R ²	0.990	0.925

Note:

* $p < 0.05$

Table 19: Raw Registration and Vote Counts from Catalist data

	<i>Dependent variable:</i>	
	Total Registrations (count)	Total Votes Cast (count)
	(1)	(2)
Preclearance x Shelby	2,733.677 (1,520.524)	1,488.581 (1,051.388)
County fixed effects	X	X
Year fixed effects	X	X
Observations	21,909	21,909
R ²	0.994	0.951
Adjusted R ²	0.994	0.943
<i>Note:</i>		* $p < 0.05$




Figure 14: Time Trends in Democratic Voteshare

As part of considering whether the composition of the electorate shifts in some way that is not visible in our analysis of voter racial/ethnic demographics, we also look at partisan vote counts. Here, we use Leip data (again for 2000-2020) to use two-party Democratic voteshare as the outcome measure for an analogous difference-in-differences setup. Table 20 presents that analysis (again at the county level with standard errors clustered on state). There is no indication that Democratic voting declined in previously-covered places after the *Shelby* decision; if anything, the point estimate in the first column suggests a slight increase in Democratic voteshare in House races, though it is not statistically distinguishable from zero. The second column uses Democratic voteshare in gubernatorial races and finds small and non-significant negative shifts.³⁹

F More on Legislator Identity

In the main paper, we consider whether downstream representational outcomes shift after *Shelby*, including the identity of House members. We also look at state legislative representation, though data is less available here. Figure 15 plots our data on Latino and Democratic representation in state assemblies. Here, we could not find consistent over-time measures of Black legislative identity, but we relied on lists produced by the National Association of Latino Elected Officials to identify Latino legislators and used data from the MIT Election and Data Science Lab and Carl Klarner to identify legislators' terms, districts, and partisanship. We then mapped legislative districts to counties as above, using Geocorr crosswalks. As seen in Figure 15, the pre-2013 trends in these measures (of having a Latino or a Democratic state assembly representing all or part of a county) do not track perfectly across covered and uncovered places, so we approach the

³⁹Because state elections occur on various time frames, including all gubernatorial races would introduce strange compositional shifts in the panel from year to year. We focus here on the majority of states that hold their gubernatorial races in midterm years, so the data for this analysis includes 2010, 2014, and 2018.

Table 20: Two-Party Voteshare Difference-in-Differences Results for Preclearance After Shelby (using Leip elections data)

	<i>Dependent variable:</i>	
	2-Party D Voteshare, House Races	2-Party D Voteshare, Gubernatorial
	(1)	(2)
Preclearance x Shelby	0.019 (0.014)	−0.013 (0.021)
County fixed effects	X	X
Year fixed effects	X	X
Observations	21,611	6,469
R ²	0.846	0.890
Adjusted R ²	0.820	0.834

Note:

**p* < 0.05

difference-in-differences analysis with some caution. However, Table 21 presents DiD estimates for these two outcomes. Again, we see no clear pattern of representational shifts: point estimates suggest slightly more Latino representation and slightly less Democratic representation, but neither estimate is distinguishable from zero.

Here we also present analogous figures/tables on House representation to those shown in the main table for Black and Democratic representation, this time focusing on whether a county is represented (all or in part) by a Latino member of Congress. Again, we do not see perfect comparability in covered/non-covered trends in the pre-period (in Figure 16, though they look broadly similar. Similar to the estimates presented in the main paper, Table 22 estimates a small decrease in Latino House representation in previously-covered places, but this estimate is not statistically distinguishable from zero.

G More Detail on EAVS Analyses

Though the EAVS began in 2004, we use data from 2008 onward due both to low response rates and varying question formats in previous years (Feder and Miller, 2020). We focus on responses from counties, omitting observations provided at the state or township level, to keep our analyses comparable to other work on the EAVS as well as the rest of the paper. We clean the data to account for a variety of different numeric codes that have been used to indicate missing values, and also to remove some implausible values. The EAVS data often includes values of 0 when the information is in fact unknown, and where possible we replace those values with missingness. For example, if

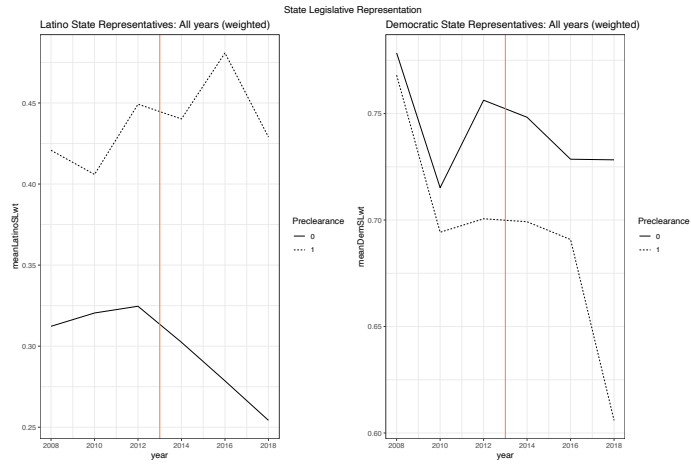


Figure 15: Time Trends in State Legislative Representation

Table 21: DiD Results for Preclearance After Shelby: State Legislative Representation

	<i>Dependent variable:</i>	
	Latino State Legislator	Dem State Legislator
	(1)	(2)
Preclearance x Shelby	0.083 (0.078)	−0.037 (0.043)
County fixed effects	X	X
Year fixed effects	X	X
Observations	14,861	14,861
R ²	0.782	0.804
Adjusted R ²	0.735	0.761

Note:

* $p < 0.05$




Figure 16: Time Trends in Congressional Representation in Covered and Non-Covered Jurisdictions

Table 22: DiD Results for Preclearance After Shelby: Congressional Representation

<i>Dependent variable:</i>	
Latino Congressperson	
Preclearance x Shelby	−0.026 (0.021)
County fixed effects	X
Year fixed effects	X
Observations	21,842
R ²	0.861
Adjusted R ²	0.837
<i>Note:</i> * $p < 0.05$	

all jurisdictions in a state report zero votes in a given year, we assume that those zeroes indicate a data issue rather than true vote counts.

We also adjust the data in several ways based on other work. Following concerns about data quality expressed in the EAVS codebook, we omit data from Iowa in 2018. And we use publicly-available code from the Pew Elections Performance Index (which relies on the EAVS dataset) to clean the code further.⁴⁰ In years where specific corrections are available for states with data issues (such as where the EPI team collected updated data directly from election officials and then manually corrected the EAVS dataset), we borrow those corrections from the EPI code. We also follow the EPI code in implementing a number of data quality checks, like making sure that subcategories (such as types of registrations) sum up to total categories (such as overall registration counts).

H Comparison to Other Recent Work

This is one in a growing series of papers seeking to estimate the effects of the *Shelby* decision on registration, turnout, and other political outcomes in areas formerly covered by preclearance. We welcome scholars investigating these questions using a wide array of data and approaches, and acknowledge that these investigations are also likely to lead to heterogeneous results and interpretations. One recent and widely publicized study by [Morris and Miller \(2024\)](#) reports, somewhat in contrast to the results presented here, that the turnout gap between white and nonwhite voters rose significantly in the wake of *Shelby v. Holder*. [Morris and Miller \(2024\)](#)'s is a thoughtful and serious analysis, but it differs from this one in several important ways worth examining in light of the importance of the research question.

First, our studies rely on different sources of raw data for turnout. While our study relies on county-level counts of registrants and voters by race provided exclusively by Catalist, [Morris and Miller \(2024\)](#) stitches together individual data from vendors Catalist for pre-*Shelby* years and L2 for post-*Shelby* years. Given that different vendors curate their records differently (e.g. prioritizing completeness of specific information fields or purging individuals who could not be contacted), this could be a consequential choice for a difference-in-differences analysis. Similarly, as [Morris and Miller \(2024\)](#) points out, we rely on Catalist's proprietary approach to estimating and aggregating race for registrants and voters while thier paper implements Bayesian Improved Surname Geocoding (BISG). Yet we successfully validate our estimates of registration and turnout for Black and Hispanic voters by comparing to the CPS and likewise compare total resulting registration and turnout counts to Leip's data (see Appendix B). Neither of those comparisons, nor [Morris and Miller \(2024\)](#), flag any systematic errors or suggest Catalist's estimates are ill-suited for this analysis.

[Morris and Miller \(2024\)](#)'s main specification also differs from ours in that it includes a term for the effect of *Shelby* on preclearance at the state level for formerly covered states and an additional term for the additional effect of *Shelby* on preclearance counties in covered states. The directions

⁴⁰See <https://doi.org/10.7910/DVN/WOV3HY>

of the coefficients on these terms are not always consistent in [Morris and Miller \(2024\)](#)’s results. Since it is the case that preclearance, if applied statewide, affected all smaller government jurisdictions within that state, and there are no cases in which counties within covered states were subjected to additional federal monitoring beyond preclearance under Section 5, the rationale behind this specification is unclear.

Perhaps the biggest difference [Morris and Miller \(2024\)](#) points out concentrates on how each study weights its results. We weight our results by raw population for the relevant group, while [Morris and Miller \(2024\)](#) employs entropy weighting to balance preclearance and non-preclearance counties. While there are certainly tradeoffs to consider for any approach to weighting the data, we choose the population weighting approach for several reasons. First, this approach helps minimize the impact of mis-classification errors on the results. For counties with relatively small (but present) minority groups, mis-classifying race for just a handful of voters under any set of modeling assumptions might produce large swings in the estimated proportion of each county’s voting-age minority population that turned out year over year. Since many counties both within and outside of the South have precisely those demographics—small but present Black minority populations—entropy balancing wouldn’t remove or downweight them, meaning that even entropy-balanced estimates might be driven by classification errors made across a large series of counties. Additionally, population-weighted estimates are substantively more interpretable if we want to know about outcomes for voters, not just counties.

[Morris and Miller \(2024\)](#) aptly points out that counties are a meaningful unit for *election administration*, and we agree with this point: when our paper considers election policy decisions made at the state (voter ID laws) or county (registration list management, pollworker hiring, etc.) level, we use states or counties as the unit of analysis without weighting by population. But when considering voter participation or legislative representation, counties are not the most relevant unit of aggregation: they rarely correspond to politically-meaningful boundaries such as legislative districts. Instead, counties are a convenient unit both for observing Section 5 coverage status and for obtaining the necessary Census data for calculating Black-white participation gaps. We use them for this calculation, but then weight by population to avoid distorting the view of overall voter experiences by disproportionate attention to small and less-well-measured places.

[Morris and Miller \(2024\)](#)’s results are fully consistent with the results reported in this paper when presenting analyses weighted by population (in SI Table A5). When doing this, the paper finds that the white-nonwhite gap in turnout decreased after *Shelby*, and observes no significant increase in the white-Black turnout gap. We take the simplest approach to addressing concerns about the application of weighting schemes to our dataset in Section C.4.5, where we show that Black-white registration and turnout gaps do not increase in formerly preclearance areas after *Shelby* even with no weights applied to the data.

Finally, in addition to these differences in specification and data source, [Morris and Miller \(2024\)](#) extends the analysis beyond 2020 into 2022. Dynamic estimates suggest that the Black-white turnout gap in previously-covered places grew further in 2022, meriting careful continuing observation by researchers who focus on voting rights.

I Pre-registration

Although this is an observational analysis and not an experiment, we pre-registered our design before purchasing the Catalist data used in this project. Here we included an anonymized copy of that preregistration document, along with notes on how the analyses presented here depart from it. The pre-analysis plan and other details of our preregistration can be found [here](#).

For the most part, we have adhered closely to the pre-registration, with several exceptions:

- The preregistration described our preferred approach as the simplest difference-in-difference specification with the outcome measure being group-specific registration and turnout rates. We have been convinced by journal reviewers that it makes sense substantively, given the history of the VRA, to center the analysis focused on racial gaps in registration and turnout rather than the simpler levels specification. But SI section C.2 continues to present those simpler group-specific estimates and they are discussed in the main paper.
- This paper focuses primarily on the main set of outcome measures described in the pre-registration document, those related to minority registration and voting. The pre-registration describes several additional outcomes that we hoped to collect about substantive or descriptive representation of minority voters. Data about the identity of legislators as well as the mapping of districts to counties over time is scarce, and the measures we have been able to construct do not map exactly to those described in the pre-registration. We have been unable to find systematic over-time data on Black state legislators, so our state legislative analysis focuses on Latino identity. However, we do examine both Black and Latino identity within Congress. Further, we realized on examination of the LCCR scores that they do not vary much within party, so we thought it more straightforward and equally informative to consider substantive representation via legislators' partisanship rather than their LCCR scores.
- The pre-registration did not discuss measurement error or whether the main Catalist analyses would be weighted or unweighted. As we discuss in the main paper, we think it makes sense to weight by group population size both because of the question we are interested in (we care about voters' experiences regardless of where they live, not about counties') and because places with very small minority populations are prone to measurement error. But in this SI (above), we present unweighted analyses and also estimates based on raw registration and turnout counts, not rates; we believe both these approaches indicate that our decision to weight the main analyses by group size does not drive the conclusions of the paper.
- Similarly, we described a robustness test that would use various ACS population estimate windows (1-year versus 5-year) to make sure that time lag in the population estimates was not driving the observed patterns. We do not present that test here because we think it is clearer and more apt to simply present the raw-counts analyses that fully drop the ACS data rather than using different variations of it.

Voting and Representation After *Shelby*: Did pre-clearance matter?

January 2019

The passage of the Voting Rights Act in 1965 had an immediate effect on voter registration, turnout, and representation, particularly in southern states that had been systematically disenfranchising African-Americans (Schuit and Rogowski, 2017). Five decades later, the Supreme Court, in the *Shelby County v. Holder* decision, dramatically changed the voting rights landscape by invalidating Section 4 of the VRA. This effectively ended the “pre-clearance” process, under which localities with a history of discrimination were required to get Justice Department approval of changes to their election procedures. The majority decision suggested that, while the VRA had once been useful, such strict monitoring was no longer necessary. Defenders of the VRA and of pre-clearance argued that the law continued to provide important protections to minority voters, and that removing this portion could have catastrophic effects (Wilson, 2015; Herron and Smith, 2016).

Since the 2013 decision, several federal elections have taken place. These do not provide enough data to test long-run theories about shifts in the electorate or slow-moving policy changes, but they do merit a simple examination of whether there have been clear shifts in minority voting registration and turnout, as well as in legislative representation, since the decision.

In this project, we examine data from recent elections, both before and after the *Shelby* decision, in counties covered and not covered by Section 5 of the VRA. We will focus on the south, as this was the region most noted for large-scale disenfranchisement before the VRA, and most covered jurisdictions are located there. We will run a simple difference-in-differences analysis comparing trends in minority voter turnout and representation before and after the *Shelby* decision across covered and non-covered places. If we find evidence of immediate shifts in registration, turnout, or representation after 2013, we will dig further into possible mechanisms, such as changes in local electoral practices that could deter minority voters.

If we do not find that any of our outcome variables have been substantially affected by the *Shelby* decision, we will conclude that some of the concerns about immediate effects of the decision have not been borne out. This could be due to concerted effort by advocates to prevent electoral changes through other legal means, or to activists who organized to mobilize minority voters in the wake of the decision, or it could be because our time frame

is too short to see longer-run effects that may materialize later. A null result here will not necessarily mean that pre-clearance was unimportant or that Justice Roberts was right that it was no longer needed, but it will rule out short-run changes in voting and representation. More data will be available when time has passed and more elections have taken place. Given the importance of this question in light of American histories of vote suppression and political exclusion, we nonetheless believe it is worth using the available data to make an early analysis of the effects of the *Shelby* decision.

Data

We will collect data on several outcome measures that capture the main goals of the VRA as we understand them: ensuring the opportunity to register and vote, especially for minority groups that have historically faced discrimination, and improving meaningful representation in government.

1 Registration and Turnout

Our main analysis will focus on voter registration and turnout in the wake of the *Shelby* decision. For this analysis, we will need local (in most cases, county-level) estimates of registration and turnout within racial or ethnic categories. We do not trust survey estimates of turnout by race at this level of aggregation, both because of misreporting and because political survey samples are generally not set up to provide valid population estimates at the level of the county, much less county-level estimates within-race. Instead, we turn to voter file data: we will use actual individual-level records of registration and turnout, combined with imputed race and ethnicity.¹ These county-level estimates will be purchased from Catalist, a firm that collects and cleans state voter files to maintain a national database. We are in the process of negotiating with Catalist to purchase this data for 2008, 2010, 2012, 2014, and 2016; we may also include 2018 data if it becomes available during the time we are working on this project.

Voter files can give us an estimate of the total number of people (by race) that were registered and/or voted in a given election, but they can't give us meaningful estimates of

¹In states where race is recorded on the voter file, we will have voters' self-reported race. In other places, we rely on Catalist's imputation of race, discussed further in Ansolabehere and Hersh (2012) and Hersh (2015), to construct estimates of the number of people who were registered to vote, and who voted, in several recent elections.

turnout rates: the proportion of *eligible* voters that actually voted.² To calculate turnout, we will use county-level estimates of the citizen voting-age population (CVAP) by race from the American Community Survey (drawn from tables B05003A-I) as the denominator. We will use the 5-year ACS population estimates (so for 2012 turnout, we will use the 2007-2012 estimates) because these provide the most complete data for the counties we are interested in.

We have not yet purchased the Catalist data for this analysis; we will do so after filing this pre-analysis plan.

2 Representation

We hope to use two measures of representation, one intended to capture the level of descriptive representation experienced by minority voters, and one intended to capture substantive representation.

We will measure descriptive representation using the proportion of Black or Latino state legislators representing a given county in a given election cycle. To the extent this increases, we will interpret that as more descriptive representation for Black/Latino voters. We will consider a county “represented” by a legislator if that legislator’s district includes any part of the county³

We will collect data on representatives’ race from a number of sources: interest groups such as NALEO often publish lists of elected officials, and we will supplement such available codings with our own codings (based on internet searches) to fill in any gaps.

We will also attempt to collect data on candidate emergence and primary elections, including the presence of minority candidates on primary ballots and their success in primary contests. However, we are uncertain about our ability to collect a complete and accurate dataset with these measures.

Then, we will measure substantive representation using at least one measure of pro-minority voting: House members’ voting scores from the Leadership Coalition on Civil Rights. These exist back to 1969, so they provide an over-time measure of pro-civil rights voting.⁴ We will use the approach described by Groseclose, Levitt and Snyder (1999) to make the scores comparable across time.

²We are less interested in the proportion of registrants that voted, since registration counts can vary for many reasons.

³We will also perform a robustness check where rather than looking at representation of “any part” of the county, we will allow for fractional representation (that is, measuring whether a Black legislator’s district covers 1/4 of the county, or 1/2, etc.).

⁴ex: <http://civilrightsdocs.info/pdf/voting-record/Voting-Record-October2016.pdf>

We acknowledge that these scores do not capture our ideal measure (a comparison of minority voters’ issue opinions and the votes cast by their representatives), but they do provide an accessible and useful measure of whether representatives appear to be voting in minority constituents’ interests. We will continue exploring other avenues for measuring substantive representation at the local level, including survey measures.

We have not yet collected any of the above measures or merged them to county-level data on Section 4 coverage; we will begin this process after filing this pre-analysis plan.

Design

We will set this up as a difference-in-differences analysis, using data from before and after the decision and from places that were and were not affected by the decision (pre-clearance and non-pre-clearance places). Our outcome measures will be:

1. Black voter turnout rates
2. Hispanic voter turnout rates
3. Black voter registration rates
4. Hispanic voter registration rates
5. Proportion of Black officials representing any part of the county in Congress/state legislature (Black officials divided by all officials)
6. Proportion of Hispanic officials representing any part of the county in Congress/state legislature
7. Average LCCR voting scores for Congresspeople representing the county.

Our main analysis will focus on counties within the South, as defined by the Census Bureau. We will present estimates separately for presidential and midterm elections.

The actual specification for the difference-in-differences setup will depend on how plausible we find the parallel trends assumption for the simplest possible specification. Once we receive/collect the county-level outcome data described above, we will examine pre-treatment trends to see whether, for example, Black voter registration in covered and non-covered counties (places that were and were not affected by the Shelby decision) followed similar trends prior to 2013. If they do appear to follow parallel trends (and placebo tests find no “effect” for covered places in periods before the Shelby decision happened), then we will use the

simplest possible difference in differences specification. We will simply predict each outcome measure (such as Black turnout) using a dummy variable for whether the county was covered by preclearance, another dummy for whether the observation was taken after the Shelby decision, and the interaction between the two variables (this is what we are interested in).

However, if we find that pre-treatment trends for covered and non-covered places look quite different, we will instead use a triple-differences approach to try to find a better comparison for the first six outcome measures.⁵ In this case, we will use white turnout (or registration/representation) in each county to try to capture time-varying forces that shape local participation and representation. We begin by presenting the example of Black voter turnout. In the triple-differences specification, we will use county-level turnout estimates, with each row of the dataset representing county turnout for a given racial group (black/white) for a given year. We will then predict turnout using, as before, a dummy variable for whether the county was covered by preclearance and a dummy variable for whether the observation is from after 2013. However, we will also include a dummy variable indicating whether the observation is for Black turnout or not, and then will include all two-way interactions between the three dummy variables, as well as the triple interaction (Covered * PostShelby * Black), which should yield the desired estimate of whether Black turnout dropped in affected counties in the wake of the Shelby decision (relative to white turnout in the same places). We will conduct analogous triple-differences analyses for the first six outcome measures laid out above.

We note that the preferred design here depends on pre-treatment trends in our observational data, which we do not yet have. If we can use the simplest diff-in-diff setup and think that the parallel trends assumption is reasonable, we would prefer to do that (it is simpler and should be better-powered). But if we find evidence that the parallel trends assumption is implausible, we will instead favor the triple-differences approach just described. Whichever approach we use, we will include the other’s estimates in an appendix or online appendix, with discussion of how we made the decision to privilege one over the other. If we find that neither approach is tenable given the data (if, for example, we find that pre-treatment placebo tests using either approach consistently yield impossible “effects” from the Shelby decision before it even happened), we will conclude that the data and design we have chosen are not well-suited to address this question, and we will give up on this entire research project.

⁵There is no meaningful white analogue to the LCCR scores, so if we don’t think the simple approach described above will work, we will drop this outcome measure.

3 Additional Tests

We will run several robustness checks, including:

- Including different states in our analysis: using a measure of “South” based on Confederacy membership, rather than Census designation, and just including all states.
- Using 1-year ACS population estimates to calculate registration and turnout, rather than 5-year estimates. This will necessarily shrink our sample of counties, as many small counties will not have population estimates reported. But the 1-year estimates are more current than the 5-year estimates in our main specification, so this specification should let us make sure any findings aren’t being driven by population shifts that throw off our population estimates.
- In some specifications, we will include controls for some potential time-varying confounders that could be driving turnout, such as election competitiveness, though we note that in some circumstances these measures could introduce post-treatment bias.
- On measures where it is possible, we will also try to run a within-North-Carolina design, taking advantage of the fact that a substantial number of NC counties were covered while others were not.

4 Extensions

If we find that minority turnout and registration (and possibly representation) decreased after *Shelby*, we will then try to discover the mechanisms by which this happened. We could examine changes to local election processes after the decision. We would also try to collect data on the racial composition of primary candidate fields in the wake of the case, to get a sense of whether candidate recruitment has changed.

If we find no effect, we will look into possible countervailing forces. For example, some political scientists have speculated that activism and mobilization would keep minority turnout relatively high for the few elections after *Shelby*. We could test this by looking at CCES self-reports of whether people were contacted during the campaign season and asked to vote, especially by non-campaign actors. Similarly, we could look for data on spending by national get-out-the-vote groups, as a measure of whether mobilization efforts increased in an attempt to counterbalance any effects of the *Shelby* decision.

References

- Ansolabehere, S. and E. Hersh. 2012. “Validation: What Big Data Reveal About Survey Misreporting and the Real Electorate.” *Political Analysis* 20(4):437–459.
URL: <http://pan.oxfordjournals.org/cgi/doi/10.1093/pan/mps023>
- Groseclose, Tim, Steven D Levitt and James M Snyder. 1999. “Comparing Interest Group Scores across Time and Chambers : Adjusted ADA Scores for the U.S. Congress.” *American Political Science Review* 93(1):33–50.
- Herron, Michael C and Daniel A Smith. 2016. “RACE, SHELBY COUNTY, AND THE VOTER INFORMATION VERIFICATION ACT IN NORTH CAROLINA.” *Florida State University Law Review* 43:465–506.
- Hersh, Eitan D. 2015. *Hacking the electorate: How campaigns perceive voters*. Cambridge University Press.
- Schuit, Sophie and Jon C. Rogowski. 2017. “Race, Representation, and the Voting Rights Act.” *American Journal of Political Science* 61(3):513–526.
- Wilson, McKenzie. 2015. “Piercing the Umbrella: The Dangerous Paradox of Shelby County v. Holder.” *Seton Hall Law Journal* 39(xii):181–205.