Federal Reserve Bank of Cleveland Working Paper Series

From Pink Slips to Polling Stations: The Electoral Consequences of Mass-Layoff Announcements

Victor Hernandez Martinez, Pawel M. Krolikowski, and
Andrew H. McCallum

Working Paper No. 26-21

August 2026

Suggested citation:

Hernandez Martinez, Victor, Pawel M. Krolikowski, and Andrew H. McCallum. 2026. "From Pink Slips to Polling Stations: The Electoral Consequences of Mass-Layoff Announcements." Working Paper No. 26-21. Federal Reserve Bank of Cleveland. https://doi.org/10.26509/frbc-wp-202621.

Federal Reserve Bank of Cleveland Working Paper Series
ISSN: 2573-7953

From Pink Slips to Polling Stations: The Electoral Consequences of Mass-Layoff Announcements*

Victor Hernandez Martinez † Pawel M. Krolikowski ‡
Andrew H. McCallum §

August 24, 2026

Abstract

We study how mass-layoff announcements affect electoral outcomes using a newly assembled dataset of US establishment-level administrative notices from the Worker Adjustment and Retraining Notification (WARN) Act. A matched event study compares same-state counties with and without sizable WARN-announced layoffs. Controlling for local labor-market conditions, layoff announcements increase the Democratic-Republican vote margin by reducing Republican votes and increasing Democratic votes without clearly altering election winners or incumbent support. Effects appear across political offices, are not driven by manufacturing decline, and are geographically localized. The responses operate through turnout composition rather than vote switching and result in no policy-induced changes in transfer payments.

JEL codes: D72, J63, J65, R23.

Keywords: Mass layoffs, WARN Act, voter turnout, local labor markets, turnout composition, electoral accountability.

---

*We thank Daniel Carroll, Bruce Fallick, and Craig Sylvera for helpful advice. Dylan Jacobs provided excellent research assistance and early insights that helped inspire the project. The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors, the Federal Reserve Bank of Cleveland, or the Federal Reserve System.

†victor.hernandezmartinez@clev.frb.org, Federal Reserve Bank of Cleveland, 1455 E 6th St, Cleveland, OH 44114

‡Corresponding author, pawel.krolikowski@clev.frb.org, Federal Reserve Bank of Cleveland, 1455 E 6th St, Cleveland, OH 44114

§andrew.h.mccallum@frb.gov, Board of Governors of the Federal Reserve System, Washington, DC 20551

Mass layoffs can have severe consequences for displaced workers, sizable employment spillovers within local labor markets, and broader social costs for households and communities.[1] Prior work also suggests a mechanism that could link mass layoffs and electoral behavior: economic conditions and individual economic adversity shape political responses (Fair, 1978; Rosenstone, 1982; Margalit, 2019). Mass-layoff *announcements*, however, may have broader political consequences because they publicly signal local economic disruption to individuals not directly affected by the layoffs. And the stakes can extend beyond elections because changes in political participation can affect policy outcomes, including public spending (Husted and Kenny, 1997; Fujiwara, 2015). Yet, despite these potential consequences, we know little about how mass-layoff announcements affect electoral outcomes.

We find that mass-layoff announcements increase the Democratic-Republican (D-R) vote-rate margin by reducing Republican votes and increasing Democratic votes, which together imply lower turnout. We assemble a new dataset of about 85,000 US establishment-level administrative notices filed under the Worker Adjustment and Retraining Notification (WARN) Act. The act requires larger employers to provide advance notice before covered plant closings and mass layoffs. We link these data to county-level voting outcomes from 1996 to 2024. Our preferred empirical strategy compares counties with sizable WARN-announced layoffs to similar counties in the same state, month, and year, but without such events. Controlling for local labor-market conditions, WARN events increase the D-R vote margin per citizen voting-age population (CVAP) by 0.50 percentage points, driven by a 0.32-percentage-point reduction in Republican votes per CVAP and a 0.18-percentage-point increase in Democratic votes per CVAP in subsequent elections. Turnout falls by 0.14 percentage points. The estimated D-R vote-rate margin effect is comparable to that associated with a 1-percentage-point increase in the county unemployment rate. These vote-rate changes, however, do not detectably change which party wins the next election or yield clear evidence of electoral accountability. The estimates remain similar after residualizing WARN exposure on standard labor-market variables and after conditioning on realized job destruction, suggesting that WARN notices convey distinct, electorally relevant information.

Additional results clarify the scope, timing, and robustness of these responses. The effects appear across presidential, gubernatorial, US Senate, and US House elections, suggesting that the pattern is not specific to any political office. They are similar for events occurring within one year and between one to two years before the next election, suggesting that proximity to the election date is not the main driver. Our findings are not

[1] See, for example, Jacobson et al. (1993), Charles and Stephens (2004), Brand (2015), Foote et al. (2019), and Gathmann et al. (2020).

driven by manufacturing decline; restricting our analysis to nonmanufacturing announcements delivers similar results. The effects also appear to be geographically localized, with little evidence of spillovers to neighboring counties. Finally, the results are stable when we add richer local controls, address slower-moving local decline, and vary key specification, sample, and event-time choices.

We next examine evidence on the sources of these vote-rate responses and on subsequent changes in economic support received by county residents. The estimates are robust to controls for migration flows and demographic composition, suggesting that changes in the size and mix of the voting-age population do not explain the electoral responses. Individual-level survey evidence instead points to turnout composition rather than vote switching, consistent with slow-moving partisan identification documented by Green et al. (2004). We also find little evidence that differential candidate visits explain the electoral responses. One possible interpretation is that WARN announcements heighten concerns about economic protection, reinforcing participation incentives among Democratic-leaning voters (Petrocik, 1996). Among Republican-leaning voters, however, these concerns may conflict with pre-existing partisan attachments, leading some voters to abstain rather than switch parties (Endres and Panagopoulos, 2019). Turning to economic support, transfer payments to county residents increase after WARN events, but appear to reflect automatic stabilizers rather than local policy changes associated with electoral responses.

Our paper provides the first US evidence on the electoral consequences of mass-layoff announcements. Outside the United States, Braakmann and Vermeulen (2023) use media reports of layoff announcements to study electoral accountability in the United Kingdom, while Dehdari (2022) uses layoff notices to examine how economic distress shifts electoral support toward the political right in Sweden. Using administrative data spanning a longer period and outcomes from multiple offices, we find that the US response operates through partisan turnout composition favoring Democrats. Our findings also contrast with evidence that offshoring and individual job losses can reduce incumbent support (Rickard, 2022; Upward and Wright, 2024). Moreover, because our main conclusions persist when we restrict our analysis to nonmanufacturing announcements, they extend beyond prior work on trade-related and deindustrialization shocks (Margalit, 2011; Autor et al., 2020; Baccini and Weymouth, 2021; Che et al., 2022; Autor et al., 2024; Choi et al., 2024). WARN effects also differ from those associated with broader labor-market deterioration. Prior work finds that weaker labor-market conditions increase Democratic votes and leave Republican votes little changed, therefore raising turnout (Charles and Stephens, 2013; Burden and Wichowsky, 2014). These contrasts suggest that WARN events capture a distinct electoral margin associated with salient local economic disruption.

1 Data

1.1 WARN notices

We measure local layoff announcements using administrative notices filed under the Worker Adjustment and Retraining Notification (WARN) Act. Federal law requires employers with 100 or more employees to provide at least 60 days' advance written notice before qualifying plant closings and mass layoffs. We assemble and harmonize about 85,000 establishment-level WARN notices from state employment offices and related public sources for 1990–2025. Historical notice availability varies across states, so the resulting panel is unbalanced and coverage expands over time: 5 states in 1996, 26 states in 2006, and 37 states in 2016. Appendix A.1 describes the act, data construction, and coverage, and Krolikowski and Lunsford (2024) provide further details on data collection.

WARN notices typically report the employer, notice date, affected worksite, and number of affected workers ("WARN layoffs"). The notice date identifies when the announcement is made, and the worksite address allows us to assign announcements to counties. Announced layoffs may be revised, delayed, canceled, or partly realized. Specifically, about half of establishment-level employment declines within six months of a WARN notice are close to or exceed the announced number of affected workers (Fallick et al., 2024). Therefore, we interpret our estimates as effects of mass-layoff announcements rather than exclusively as effects of realized job loss.

1.2 Election returns

We use county-level election returns by party from the US Elections Atlas for presidential, gubernatorial, US House, and US Senate elections from 1996–2024 (Leip, 2026). We include general elections held from November 1 through November 8, and exclude a small number (35) of special and recall elections. The sample contains county-level returns from 8 presidential elections, 490 Senate elections, 6,510 House elections, and 361 gubernatorial elections. We exclude state legislative, municipal, county, school-board, and ballot-measure contests because available data lack consistent nationwide coverage of returns, lack coverage of candidate characteristics, or do not map uniformly to the county units used in our design.

Our main outcomes are the Democratic-Republican (D-R) vote-rate margin, party-specific vote rates, and turnout. Party-specific vote rates and turnout are given by

$$T^{p}_{cyo} = 100 \times \frac{V^{p}_{cyo}}{CVAP_{cy}}, \qquad T_{cyo} = 100 \times \frac{V_{cyo}}{CVAP_{cy}}, \tag{1}$$

in which V^{p}_{cyo} denotes votes cast in county c, in election year y, for office o, and party $p \in \{D, R, O\}$. The categories D, R, and O denote Democratic, Republican, and other candidates, respectively, with the last including independent and third-party candidates. $CVAP_{cy}$ denotes the citizen voting-age population from the US Census Bureau (CB, 2002, 2024, 2025, 2026; IPUMS, 2026). Total votes and turnout satisfy $V_{cyo} = \sum_{p\in\{D,R,O\}} V^p_{cyo}$ and $T_{cyo} = \sum_{p\in\{D,R,O\}} T^p_{cyo}$. The D-R vote-rate margin is given by $T^D_{cyo} - T^R_{cyo} = 100 \times (V^D_{cyo} - V^R_{cyo})/CVAP_{cy}$.

When possible, we link returns to candidate and office-holder information from the Federal Election Commission (FEC, 2026) for presidential and congressional contests, and Wikipedia for gubernatorial contests. These links yield two accountability outcomes: incumbent-candidate vote share and the vote share of the party holding the office. Appendix A.2 describes the election returns data, election sample, candidate-party harmonization, CVAP sources, and incumbent-candidate links, including assignment in House elections.

1.3 Auxiliary data and controls

We supplement the WARN and election data with county-level labor-force, employment, and unemployment measures (BLS, 2026a), per capita income (BEA, 2026b), migration flows (IRS, 2026), and COVID controls (Killeen et al., 2020). We use these data to construct treatment intensity and treated-control matches, add regression controls, and interpret the results.

Some robustness specifications add richer county employment and wage histories (BLS, 2026b), income components (BEA, 2026b), education and rural-urban classifications (USDA, 2026), age and race/ethnicity populations (SEER, 2026), and land area (CB, 2025). For interpretation, we link WARN events to county-level firm job-destruction rates from the Quarterly Workforce Indicators (QWI, 2026), individual-level Cooperative Election Study (hereafter CES) voting data (Kuriwaki, 2026), presidential and vice-presidential candidate visits (Devine, 2020), and county transfer payments (BEA, 2026a). We classify WARN events as manufacturing or nonmanufacturing using state-provided industry codes when available and manual or AI-assisted classification otherwise. Appendix A.3 provides details on these data and their construction.

2 Empirical strategy

2.1 Defining WARN events, treatment, and election-year windows

We define WARN events based on three-month rolling exposure,

$$R_{cm} = E_{cm} + E_{c,m-1} + E_{c,m-2}, \tag{2}$$

in which $E_{cm} = 100 \times W_{cm}/LF_{cm}$, W_{cm} denotes WARN layoffs in county c and month m, and LF_{cm} denotes the county labor force. A county-month observation has a sizable

WARN event if R_{cm} is at or above the 85th percentile among county-months with positive rolling WARN exposure and at least 50 workers are affected. County c is treated in month m if it has a sizable WARN event and has had no such event in the previous four years. Treatments may recur if they are at least four years apart.

To align monthly WARN events with elections, we define 12-month election-year windows around regular November elections. For an election in year y, the nearest window runs from November of year $y - 1$ through October of year y. Earlier and later election-year windows are defined analogously in 12-month intervals. Appendix B.1 provides details on WARN-event treatment and election-year windows.

2.2 Constructing treated-control comparisons

Counties experiencing sizable WARN events may differ from other counties before the event. Our matched design, closely related to Abadie (2005) and Schmieder et al. (2023), compares each treated event with potential controls that have similar pre-event political and economic histories.

We use the following algorithm to choose control observations for a treated county-month observation, (c, m), in year y. We restrict potential controls in the matched design to observations from the same state as county c, in month m and year y. Controls must have no sizable WARN event in the treatment election-year window, the four preceding windows, and up to four subsequent windows. Near the end of the sample, the post-treatment restriction applies only to subsequent election-year windows observed through 2024. Smaller WARN events do not disqualify control observations.

Among eligible controls, we use nearest-neighbor matching on recent turnout, the two most recent Republican vote shares, labor-force size, the unemployment rate and its change, and WARN exposure in each of the previous two years. When multiple offices are contested on the same date, we prioritize presidential, gubernatorial, Senate, and House elections, in that order. To avoid poor matches, we impose calipers—maximum treated-control differences on each matching variable (Caliendo and Kopeinig, 2008; Dauth et al., 2021). Among counties within the calipers, we rank controls by standardized Euclidean distance, retain up to four of the closest controls, and apply inverse-distance weights.

If a treated observation has no controls, we repeat the algorithm using a fallback that preserves state and treatment-year alignment but relaxes exact treatment-month matching. Appendix B.2 provides details on control eligibility, matching, calipers, and weights.

2.3 Treatment characteristics, matched support, and balance

Our matched analysis sample spans 1996–2024 and draws on an unbalanced panel of states with WARN and election-return data.

Panels A and B of Table 1 show summary statistics about treatment characteristics and electoral outcomes. Panel A shows that sizable WARN events in the matched sample are economically meaningful local shocks. The 85th-percentile threshold corresponds to 0.6 percent of a county's labor force (not shown). At the median, treated events involve 206 affected workers, or 1.1 percent of the county labor force, comparable to mass-layoff shocks studied in other work (Foote et al., 2019; Gathmann et al., 2020; Braakmann and Vermeulen, 2023). The median event occurs 10 months before the next election and consists of one WARN notice. Appendix B.3.1 provides further details on treated-event characteristics and matched support and shows their distribution over time. Panel B shows that, for example, mean turnout is 50.2 percent and the mean Democratic (Republican) vote rate is 19.4 (29.4) percent.

Panels C and D summarize matched-sample support and balance. Panel C shows that of 1,703 treated events across 1,300 counties, 1,508 (89 percent) are matched to at least one control. Three percent of matched events use the state-year fallback, and matched events receive 3.5 controls on average. Relative to matched events, unmatched events are similarly sized as a share of the local labor force, but occur in labor markets with a larger labor force. Panel D shows that matching substantially improves balance on the matching variables. Appendix B.3.2 provides additional support and balance diagnostics.

2.4 Event-study design

We organize event time by the chronological order of elections. For each matched comparison group g, anchored by a treated county's WARN-event date, let $h_g = 0$ denote the last pre-event election, $h_g = 1$ the first post-event election, and so on; we omit the g subscript hereafter. Elections held on the same date share the same chronological index but contribute separate county-office observations. Because this design is office-agnostic, outcomes may be compared with earlier elections for different offices within the same matched comparison. To account for office-specific differences in outcomes, the estimating equation relies on fixed effects.

To examine heterogeneity across election types, we also organize event time separately within each office. For each matched comparison group g and office o, $h = 0$ denotes the last pre-event election for that office, $h = 1$ the first post-event election, and so on. Although the office-specific and chronological designs use the same treatment events and matching algorithm, they may select different controls because the former matches on office-specific pretreatment political variables, whereas the latter matches on these variables without regard to office.

In both the chronological and office-specific designs, we combine observations from all treated-control comparison groups into a single estimation sample. Schmieder et al. (2023) discuss how this approach addresses problems arising in difference-in-differences designs with variation in treatment timing (Goodman-Bacon, 2021; Callaway and Sant'Anna, 2021).

We estimate the effects of WARN events on electoral outcomes using

$$Y_{cgoh} = \sum_{i \neq 0} \beta_i \left(D_{cg} \times \mathbf{1}\{h = i\} \right) + \mathbf{X}'_{cgh}\delta + \alpha_{cg} + \lambda_h + \kappa_{syo} + \varepsilon_{cgoh}, \tag{3}$$

in which s denotes county c's state, Y_{cgoh} denotes an outcome, and D_{cg} equals one for the treated county in comparison group g and zero for its matched controls. With $h = 0$ omitted, β_i measures the treated-control difference at election index i relative to the difference at the last pre-event election. The chronological design retains post-event indexes through $h = 4$, a horizon of up to eight years. Within each comparison group, g and h uniquely determine election year y, so the outcome subscripts are equivalent to those in Equation (1).

The vector $\mathbf{X}_{cgh}$ contains county-level controls. The baseline specification controls for the unemployment rate, labor-force size, per capita income, and cumulative-to-date COVID cases per 1,000 residents. The county-by-comparison fixed effect, α_{cg}, absorbs persistent within-comparison county differences, and λ_h absorbs average differences across relative election periods. The state-by-office-by-year fixed effect, κ_{syo}, absorbs shocks common to counties voting for the same office in the same state and election year, including race-specific political environments, election administration, and ballot context. We two-way cluster standard errors by county and matched comparison group to account for dependence induced by stacking and correlated sub-state shocks within matched comparison groups.

The identifying assumption is that, absent a sizable WARN event, electoral outcomes in treated and matched-control counties would have followed parallel paths. The main identification threat is latent local deterioration or other coincident shocks that would have changed electoral outcomes even without the announcement. We address this threat with controls for contemporaneous county labor-market conditions in our baseline specification, a residualized-exposure exercise in Section 3.1, and several robustness exercises in Section 3.4.

The chronological design is not the only reasonable way to pool outcomes across offices. We also report results from an alternative target-office design in which every pre- and post-WARN comparison is made within the same office. We stack the office-specific comparisons, so that the $h = 1$ coefficient pools observations from four office-specific designs, one for each office's first post-event election. Therefore, the $h = 1$ coefficient estimates a weighted average across offices at each office's first post-event election, regardless of calendar timing. We retain post-event indexes through $h = 3$, a horizon of up to 12 years. Appendix B.4 provides further details on event-time ordering and the office-specific design.

3 Main results

3.1 Baseline electoral effects of WARN events

This section presents the baseline electoral effects of WARN events, assesses their magnitude and implications for election winners and incumbent support, and asks whether the effects are captured by standard measures of local labor-market conditions and whether they persist after accounting for subsequent job destruction.

A sizable WARN event increases the Democratic-Republican (D-R) vote-rate margin by reducing Republican votes and increasing Democratic votes, which together imply lower turnout. Figure 1 reports estimates from Equation (3) for the D-R vote-rate margin, party-specific vote rates, and turnout. It also reports a *post* estimate from a specification with one post-event coefficient for $h \in \{1, 2, 3, 4\}$ and one pre-event coefficient for $h \in \{-3, -2, -1\}$. WARN events increase D-R votes per CVAP by 0.50 percentage points in subsequent elections (Panel A), reduce Republican votes per CVAP by 0.32 points (Panel B), and increase Democratic votes per CVAP by 0.18 points (Panel C). Other-candidate votes per CVAP are essentially unchanged (Table 2, Row 1). Together, these shifts imply a 0.14-percentage-point decline in turnout (Panel D), although the estimate is marginally insignificant (p-value = 0.14). We find no evidence of pre-trends in any outcome. Appendix C.1.1 shows that the target-office approach yields similar results, with a weaker Democratic increase and a larger, statistically significant turnout decline of 0.19 percentage points (p-value = 0.02).

The estimated effects are meaningful at the county level in two ways. First, in the baseline specification, the increase in the D-R vote-rate margin is comparable to that associated with a 1-percentage-point increase in the county unemployment rate. The Republican and Democratic vote-rate effects are comparable to those associated with 1.5- and 0.6-percentage-point changes, respectively, in the unemployment rate. The turnout decline is similarly comparable to that associated with a 1.4-percentage-point unemployment-rate decline. Second, about 2 percent of treated county-office outcomes at the next election have D-R vote-rate gaps smaller than the estimated 0.50-percentage-point WARN effect. For a treated county with the sample-mean baseline CVAP (about 60,000), this effect corresponds to roughly 300 votes.

WARN events shift partisan vote rates without measurably changing which party wins the next election. We examine winner outcomes in three ways. First, we reestimate Equation (3) using as the outcome an indicator for whether Democratic votes exceed Republican votes in the county. We find no significant increase in the unrestricted county sample, when the preceding election was close, or when the preceding election and the preceding county-level margin were both close. Second, after mapping WARN events to US House districts, we find no evidence that treated districts are more likely to elect a Democrat. For this exercise, we retain only treated events in county-election windows that can be unambiguously assigned to a district and control counties that can be mapped to districts that can be unambiguously determined to lack qualifying WARN exposure. These restrictions reduce the number of treated events from 1,703 to 1,112 and result in only about 300 district-election pairs. Third, for presidential, gubernatorial, and Senate elections, we find no evidence that statewide WARN exposure between consecutive elections is associated with a higher probability of a Democratic victory in the state. Appendix C.1.2 provides details.

The vote-rate responses are also not accompanied by clear evidence of electoral accountability. Although prior work finds that incumbent parties lose support after production offshoring and individual job loss (Rickard, 2022; Upward and Wright, 2024), Appendix C.1.3 shows no clear decline in incumbent-party vote shares after WARN events. Incumbent-candidate vote shares fall in some specifications, but the estimates are noisy.

We next examine whether WARN events convey distinct information affecting electoral outcomes in three ways. First, WARN events have electoral consequences beyond those that can be predicted by other labor-market variables. We regress WARN exposure on a battery of county-level labor-market variables and state, year, and month fixed effects.[2] We use the residual from this regression to reconstruct the treatment indicator: treated events are those above the 85th percentile among observations with positive residual WARN exposure, subject to a residualized baseline-equivalent 50-worker floor. We then reconstruct the treated-control comparisons and reestimate Equation (3). In this approach, treatment represents unusually high WARN exposure given existing local labor-market conditions. The residualized and baseline definitions identify related but distinct treated events. Half of residualized events coincide with baseline treatments, 5 percent qualify in the same county and election-year window but in a different month, 25 percent occur in baseline-treated counties but in different election windows, and the remaining 20 percent occur in counties not treated under the baseline definition. The resulting electoral effects of these residual WARN events are quantitatively similar to our baseline results, as shown in Table 2, Row 2. Appendix C.1.4 provides details about this exercise.

Second, post-WARN job destruction does not explain the partisan electoral response,

[2] Labor-market variables include log labor force; employment-to-population and unemployment rates; per capita income; manufacturing, goods-producing, and private-employment shares; average weekly wages; average establishment size; log population; and their respective 12-month and annualized 5-year changes.

leaving room for an announcement component. Because WARN notices generally precede covered layoffs, the electoral response may reflect the announcement itself, the job destruction that follows, or both. The residualized-WARN exercise cannot distinguish these channels. We therefore use county-level firm job-destruction rates (JDR) from the QWI, which measure jobs lost at contracting private-sector firms as a share of average county employment. Treated counties experience a 0.84-percentage-point larger increase in JDR than matched controls, confirming that WARN events are followed by higher realized job destruction. We then reestimate our baseline specification by interacting the h indicators with each county's change in JDR from before to after the WARN event. Table 2, Row 3 shows that conditioning on JDR only slightly attenuates the baseline electoral responses, indicating that measured destruction accounts for little of the partisan response. Although this exercise does not identify an announcement effect, the remaining response is consistent with an announcement component. Appendix C.1.5 provides details about this exercise.

Third, the electoral response to WARN events also differs from that associated with other measures of local labor-market deterioration. Prior work finds that weaker local labor-market conditions increase Democratic votes, leave Republican votes little changed, and raise turnout (Charles and Stephens, 2013; Burden and Wichowsky, 2014). WARN events instead increase the D-R vote-rate margin through both higher Democratic votes and lower Republican votes, while reducing turnout. Together with the residualized-exposure and job-destruction results, this contrast suggests that WARN notices capture a distinct margin of local economic disruption rather than broader labor-market weakness.

3.2 Office-specific effects

The baseline estimates average across offices that differ in salience, campaign intensity, and electorate composition. Figure 2 reports the D-R vote-rate margin, Republican and Democratic vote rates, and turnout effects separately for presidential, gubernatorial, US House, and US Senate elections, allowing us to assess whether the response is concentrated in a particular office.

The D-R margin shifts in the same direction across all four office types, and the underlying estimates show similar qualitative patterns, particularly for turnout and Republican votes. The responses for Democratic votes are more heterogeneous, largest in House races and smallest in Senate elections. Nevertheless, the pooled estimates appear to reflect a broader change in party-specific vote rates and turnout rather than a single-office artifact.

3.3 Heterogeneity and geographic scope

This subsection examines heterogeneity in the main responses documented in Section 3.1. We focus on three margins that bear on interpretation, including the timing of the announcement relative to the next election, whether the event occurs outside manufacturing, and whether the response spills over to nearby counties. The estimates point to a localized response that is neither narrowly timed nor manufacturing specific.

First, the response is not concentrated among WARN events closest to the next election. A WARN event shortly before an election may be especially salient to voters and campaigns, whereas earlier events allow more time for political mobilization. We examine this timing margin by dividing treated events into those occurring within one year of the next election ($\tau = 0$) and those occurring one to two years before it ($\tau = 1$). The results are shown in Table 2, Panel B. Although the estimates are noisier because we are splitting the treated sample, both groups show higher D-R vote-rate margins, driven by lower Republican votes per CVAP and higher Democratic votes per CVAP, and lower turnout. Overall, the results do not support a purely short-run election-salience story. Appendix C.2.1 provides details and estimates by election index.

Second, the response is not confined to manufacturing announcements. Prior work examines the political consequences of manufacturing decline, including trade exposure and deindustrialization (Autor et al., 2020; Baccini and Weymouth, 2021). We show that the responses to sizable WARN events capture a broader form of economic disruption than simply manufacturing decline. Table 2, Panel C reconstructs treatment and the matched comparisons using only nonmanufacturing notices and reestimates Equation (3). The results are similar to the baseline results, with a somewhat stronger Democratic response and a slightly weaker Republican decline, so that the turnout effect is closer to zero. Appendix C.2.2 provides details and shows that sizable nonmanufacturing WARN events are rarely accompanied by sizable manufacturing WARN events, so these estimates do not mask a manufacturing-driven response.

Third, the response appears geographically localized. We assess spillovers to nearby counties by performing a donut exercise that keeps the county-level design but excludes potential control counties in the treated county's commuting zone. The resulting donut estimates remain close to the baseline, as shown in Table 2, Panel D, indicating limited spillovers to neighboring counties. Relatedly, these results also reduce the concern that the baseline estimates are attenuated by using nearby controls that are partially exposed to the same WARN disruption. Appendix C.2.3 reports estimates by election index.

3.4 Endogeneity checks and robustness of the main vote-rate and turnout effects

3.4.1 Pre-existing local deterioration and estimate validity

The main estimates are not driven by observed pre-existing local deterioration or other observable economic variables. Table 2, Panel E varies how we account for local economic conditions. Rows 8 and 9 omit controls or add richer local controls to the baseline specification, including the manufacturing employment share, average weekly wages, and population density. Row 10 addresses slower-moving local decline by rematching on variables that include longer pre-event trajectories and variables from the Quarterly Census of Employment and Wages. Row 11 adds forecast-score-by-event-index controls, allowing counties with different predicted-deterioration scores to follow different event-time paths even absent treatment. Across these exercises, post-event estimates remain close to the baseline estimates and pre-event coefficients remain small. Appendixes C.3.1 and C.3.2 provide details.

3.4.2 Robustness of the estimates

The main vote-rate and turnout estimates are stable across key specification, sample, and design choices in the matched event study. Table 2, Panel F shows similar estimates under alternative WARN-event thresholds (Rows 12 and 13), minimum-worker floors (Rows 14 and 15), rolling exposure windows (Rows 16 and 17), and up to 10 matched controls (Row 18). The main patterns also remain similar when we use comparison-by-election-index-by-office fixed effects, which restrict identification to within matched comparison groups (Row 19); apply log-CVAP weights (Row 20); and exclude WARN events during the 2020 treatment-election-year (Row 21). Appendix C.4 provides details about these exercises. Appendix C.4.5 also shows that the results are similar when we use a log-odds transformation to address bounded support under the linear specification in Equation (3).

4 Evidence on electoral mechanisms and policy outcomes

4.1 Population mix, turnout composition, and vote switching

The county-level results in Section 3.1 show that WARN events have electoral consequences, but they do not by themselves identify the underlying margin. These responses could reflect changes in the size or composition of the county population, changes in who turns out, or changes in how voters choose among parties. We distinguish these possibilities using population-composition checks and individual-level CES voting data. The evidence points to turnout composition rather than population change or vote switching, with little role for differential candidate visits. We conclude by discussing an interpretation consistent with these findings.

Changes in migration or demographic composition do not explain the county-level vote-rate patterns. Table 3, Panel A reports the estimates from a specification that adds controls for county-level migration flows and demographic structure to our baseline specification. These checks test whether WARN events affect measured vote rates by changing the county population rather than voting behavior among existing residents. Migration is especially relevant because residential mobility can reduce turnout through registration frictions (Squire et al., 1987). So the turnout decline could reflect post-announcement mobility rather than a direct participation response to the announcement. Nevertheless, the estimates remain similar with these controls. Appendix D.1 presents details.

We next use individual-level CES data to distinguish turnout composition from vote switching. Our estimated declines in turnout rule out a pure vote-switching story, but aggregate county rates can hide individual entry, exit, and switches. We use the matched comparison groups described in Section 2.2, and attach each individual in the CES to their county, election year, and relevant office. Then we estimate Equation (3) at the individual level including basic demographic controls.

The CES data provide little evidence of major-party switching. The first (second) column of Table 3, Panel B, restricts the sample to respondents who intended to vote Republican (Democratic) for the relevant office and later report voting in the election. The outcome equals one if they report voting Democratic (Republican). The estimated switching responses are insignificantly different from zero, suggesting no clear major-party switching among voters after a WARN event.

In contrast, differential turnout among major-party leaners points toward a compositional response.[3] The third column of Table 3, Panel B, restricts the sample to major-party leaners and uses as the outcome a dummy variable for officially validated respondent turnout; column four further restricts the sample to respondents who intended to vote before the election. In both columns, we estimate an augmented version of Equation (3) in which the pre- and post-treatment dummies and county controls and fixed effects are also interacted with whether the respondent leans Republican. The reported point estimates correspond to the pre- and post-interaction coefficients. The estimates indicate that the post-WARN turnout response among Republican leaners is 4.2 percentage points lower than among Democratic leaners overall and 5.1 points lower in the sample of respondents who intended to vote.

These changes in electoral mobilization do not appear to be explained by visits of presidential-ticket candidates to affected counties. Campaigns may target economically

[3]Leaners are respondents who identify as Democratic or Republican or are independents who lean toward one of the major parties.

distressed areas with candidate travel, which could directly generate differential local mobilization and contribute to the turnout-composition response. Table 3, Panel C shows how vice-presidential and presidential candidate visits respond to WARN events using our baseline specification. We find no relative increase in Democratic versus Republican campaign visits to WARN-treated counties, measured by either any visit (first column) or visit counts (second column). These results weaken a simple candidate-travel explanation, although the visit data capture only one narrow form of campaign attention. Appendix D.3 provides details.

Taken together, the evidence points to a turnout-composition response that is not explained by population change, vote switching, or candidate visits. One interpretation is that mass-layoff announcements heighten concern about future economic risk and thereby affect participation incentives differently across partisan groups. More generally, political messages can mobilize voters facing economic hardship (Aytaç et al., 2020). Economic insecurity can increase demand for social insurance (Rehm et al., 2012; Margalit, 2013), a policy domain traditionally dominated by the Democratic Party (Petrocik, 1996; Faricy, 2015). Therefore, if a WARN announcement heightens concerns about economic protection, it may reinforce existing partisan preferences among Democratic-leaning voters and increase their incentive to participate.[4] Among Republican-leaning voters, the same concerns may conflict with partisan attachments, leading some to abstain rather than switch parties (Endres and Panagopoulos, 2019). Such a mechanism would be consistent with the turnout-composition response documented above. Admittedly, we do not directly observe the salience of WARN events or the degree of resulting partisan tension, so we view this as an interpretation consistent with the evidence rather than a mechanism identified by our design.

4.2 Transfer payments after WARN events

We examine whether WARN events are followed by changes in economic support through transfer payments to residents, and whether any changes appear to reflect the electoral response. Prior work shows that participation and the composition of turnout can affect public spending and redistribution (Husted and Kenny, 1997; Fujiwara, 2015). In our setting, transfer payments may rise for at least two reasons. First, WARN events may directly increase economic need, raising payments from programs such as unemployment insurance and food assistance. Second, if the electoral response changes

[4] Braakmann and Vermeulen (2023) find limited effects of mass-layoff announcements on individuals' expectations about their own economic situations. This finding does not necessarily rule out political responses to mass-layoff announcements, however, because voting decisions need not be based only on individual economic expectations (Riker and Ordeshook, 1968; Kinder and Kiewiet, 1981; Rosenstone and Hansen, 1993).

local political incentives or representation, it could affect the allocation of more discretionary public resources. Therefore, we ask whether transfer payments increase after WARN events and whether these increases reflect automatic stabilizers or electoral responses.

Table 3, Panel D estimates the effects of WARN events on log real county-level transfer payments per capita. Payments rise after WARN events in the shock-responsive categories (first column), including unemployment insurance and the Supplemental Nutrition Assistance Program, but not in more discretionary categories (second column), mainly education and training assistance. Appendix D.4 examines whether the transfer-payment response is associated with post-WARN electoral changes by comparing estimates with and without controls for those changes. This difference is negligible, suggesting that electoral outcomes play little role in transfer changes. Taken together, these findings suggest that transfer payments respond to local distress, mostly through automatic stabilizers rather than through local policy changes associated with electoral responses.

5 Conclusion

This paper shows that mass-layoff announcements have electoral consequences. We use a newly assembled dataset of US establishment-level WARN notices. A matched event-study design shows that sizable WARN events increase the D-R vote-rate margin by reducing Republican votes and increasing Democratic votes. Our results suggest that sizable WARN events can alter electoral participation in ways not captured by other labor-market variables and realized job destruction. Effects appear across political offices, are not driven by manufacturing decline, and are geographically localized.

Additional evidence helps interpret these observed electoral responses. Survey evidence suggests that the party-specific vote-rate effects are more consistent with changes in turnout composition among voters with different pre-existing political preferences than with vote switching across parties. A natural next step for future work is to distinguish more sharply why layoff announcements change turnout composition, by investigating the roles of private economic insecurity, partisan conflict, media attention, and campaign responses.

References

Abadie, Alberto (2005). "Semiparametric Difference-in-Differences Estimators." *The Review of Economic Studies*, 72(1), pp. 1--19. doi:10.1111/0034-6527.00321.

Autor, David, Anne Beck, David Dorn, and Gordon H. Hanson (2024). "Help for the Heartland? The Employment and Electoral Effects of the Trump Tariffs in the United States." URL https://www.nber.org/papers/w32082, NBER Working Paper.

Autor, David, David Dorn, Gordon Hanson, and Kaveh Majlesi (2020). "Importing Political Polarization? The Electoral Consequences of Rising Trade Exposure." *American Economic Review*, 110(10), pp. 3139--3183. doi:10.1257/aer.20170011.

Aytaç, S. Erdem, Eli Gavin Rau, and Susan Stokes (2020). "Beyond Opportunity Costs: Campaign Messages, Anger and Turnout among the Unemployed." *British Journal of Political Science*, 50(4), pp. 1325--1339. doi:10.1017/S0007123418000248.

Baccini, Leonardo and Stephen Weymouth (2021). "Gone For Good: Deindustrialization, White Voter Backlash, and US Presidential Voting." *American Political Science Review*, 115(2), pp. 550--567. doi:10.1017/S0003055421000022.

BEA (2026a). "CAINC35: Personal Current Transfer Receipts." https://www.bea.gov/itable/discontinued-data-tables. Accessed July 2026.

BEA (2026b). "Regional Economic Accounts." https://www.bea.gov/data/economic-accounts/regional. Accessed July 2026.

BLS (2026a). "Local Area Unemployment Statistics." https://www.bls.gov/lau/. Accessed July 2026.

BLS (2026b). "Quarterly Census of Employment and Wages." https://www.bls.gov/cew/. Accessed July 2026.

Braakmann, Nils and Wessel N. Vermeulen (2023). "Do Mass Layoffs Affect Voting Behaviour? Evidence from the UK." *British Journal of Industrial Relations*, 61(4), pp. 922--950. doi:10.1111/bjir.12749.

Brand, Jennie E. (2015). "The Far-Reaching Impact of Job Loss and Unemployment." *Annual Review of Sociology*, 41(1), pp. 359--375.

Burden, Barry C. and Amber Wichowsky (2014). "Economic Discontent as a Mobilizer: Unemployment and Voter Turnout." *The Journal of Politics*, 76(4), pp. 887--898. doi:10.1017/S0022381614000437.

Caliendo, Marco and Sabine Kopeinig (2008). "Some practical guidance for the implementation of propensity score matching." *Journal of Economic Surveys*, 22(1), pp. 31--72. doi:10.1111/j.1467-6419.2007.00527.x.

Callaway, Brantly and Pedro H. C. Sant'Anna (2021). "Difference-in-Differences with Multiple Time Periods." *Journal of Econometrics*, 225(2), pp. 200--230. doi:10.1016/j.jeconom.2020.12.001.

CB (2002). "Citizen Voting Age by Race and Ethnicity - Census 2000." https://www.census.gov/programs-surveys/decennial-census/about/voting-rights/cvap/2000-cvap.html. Accessed July 2026.

CB (2024). "Intercensal Estimates." https://www.census.gov/programs-surveys/popest/technical-documentation/research/intercensal-estimates.html. Accessed July 2026.

CB (2025). "Gazetteer files." https://www.census.gov/geographies/reference-files/time-series/geo/gazetteer-files.html. Accessed July 2026.

CB (2026). "American Community Survey 5-Year Data (2009--2024)." https://www.census.gov/data/developers/data-sets/acs-5year.html. Accessed July 2026.

Charles, Kerwin Kofi and Melvin Stephens, Jr. (2004). "Job Displacement, Disability, and

Divorce." *Journal of Labor Economics*, 22(2), pp. 489--522. doi:10.1086/381258.
Charles, Kerwin Kofi and Melvin Stephens, Jr. (2013). "Employment, Wages, and Voter Turnout." *American Economic Journal: Applied Economics*, 5(4), pp. 111--143. doi:10.1257/app.5.4.111.
Che, Yi, Yi Lu, Justin R. Pierce, Peter K. Schott, and Zhigang Tao (2022). "Did Trade Liberalization with China Influence US Elections?" *Journal of International Economics*, 139, p. 103,652. doi:10.1016/j.jinteco.2022.103652.
Choi, Jiwon, Ilyana Kuziemko, Ebonya Washington, and Gavin Wright (2024). "Local Economic and Political Effects of Trade Deals: Evidence from NAFTA." *American Economic Review*, 114(6), pp. 1540--1575. doi:10.1257/aer.20220425.
Dauth, Wolfgang, Sebastian Findeisen, and Jens Suedekum (2021). "Adjusting to Globalization in Germany." *Journal of Labor Economics*, 39(1), pp. 263--302. doi:10.1086/707356.
Dehdari, Sirus H. (2022). "Economic Distress and Support for Radical Right Parties---Evidence from Sweden." *Comparative Political Studies*, 55(2), pp. 191--221. doi:10.1177/00104140211024301.
Devine, Christopher J. (2020). "Campaign Visits Database." `http://www.christopherjdevine.com/data.html`. Accessed July 2026.
Endres, Kyle and Costas Panagopoulos (2019). "Cross-Pressure and Voting Behavior: Evidence from Randomized Experiments." *The Journal of Politics*, 81(3), pp. 1090--1095. doi:10.1086/703210.
Fair, Ray C. (1978). "The Effect of Economic Events on Votes for President." *The Review of Economics and Statistics*, 60(2), pp. 159--173.
Fallick, Bruce, Andrew Foote, Pawel M. Krolikowski, and Lee Tucker (2024). "Advance Layoff Notices and Worker Outcomes." Presentation at the 2024 AEA Annual Meeting. URL `https://www.aeaweb.org/conference/2024/program/powerpoint/BG4HetKi`, presented January 6, 2024.
Faricy, Christopher G. (2015). *Welfare for the Wealthy: Parties, Social Spending, and Inequality in the United States*. Cambridge University Press.
FEC (2026). "Candidate master file description." `https://www.fec.gov/campaign-finance-data/candidate-master-file-description/`. Accessed July 2026.
Foote, Andrew, Michel Grosz, and Ann Stevens (2019). "Locate Your Nearest Exit: Mass Layoffs and Local Labor Market Response." *ILR Review*, 72(1), pp. 101--126. doi:10.1177/0019793917753095.
Fujiwara, Thomas (2015). "Voting Technology, Political Responsiveness, and Infant Health: Evidence from Brazil." *Econometrica*, 83(2), pp. 423--464. doi:10.3982/ECTA11520.
Gathmann, Christina, Ines Helm, and Uta Schönberg (2020). "Spillover Effects of Mass Layoffs." *Journal of the European Economic Association*, 18(1), pp. 427--468. doi:10.1093/jeea/jvy045.
Goodman-Bacon, Andrew (2021). "Difference-in-Differences with Variation in Treatment Timing." *Journal of Econometrics*, 225(2), pp. 254--277. doi:10.1016/j.jeconom.2021.03.014.
Green, Donald, Bradley Palmquist, and Eric Schickler (2004). *Partisan Hearts and Minds: Political Parties and the Social Identities of Voters*. Yale University Press.
Husted, Thomas A. and Lawrence W. Kenny (1997). "The Effect of the Expansion of the Voting Franchise on the Size of Government." *Journal of Political Economy*, 105(1), pp. 54--82. doi:10.1086/262065.
IPUMS (2026). "IPUMS National Historical Geographic Information System: Version 21.0." Dataset. doi:10.18128/D050.V21.0. URL `https://data2.nhgis.org/main`.
IRS (2026). "SOI Tax Stats: Migration Data." `https://www.irs.gov/statistics/soi-tax-stats-migration-data`. Accessed July 2026.
Jacobson, Louis S., Robert J. LaLonde, and Daniel G. Sullivan (1993). "Earnings Losses of

Displaced Workers." *The American Economic Review*, 83(4), pp. 685--709. URL http://www.jstor.org/stable/2117574.

Killeen, Benjamin D., Jie Ying Wu, Kinjal Shah, Anna Zapaishchykova, Philipp Nikutta, Aniruddha Tamhane, Shreya Chakraborty, Jinchi Wei, Tiger Gao, Mareike Thies, and Mathias Unberath (2020). "A county-level dataset for informing the United States' response to COVID-19." *arXiv preprint arXiv:200400756*. doi:10.48550/arXiv.2004.00756.

Kinder, Donald R. and D. Roderick Kiewiet (1981). "Sociotropic Politics: The American Case." *British Journal of Political Science*, 11(2), pp. 129--161. doi:10.1017/S0007123400002544.

Krolikowski, Pawel M. and Kurt G. Lunsford (2024). "Advance Layoff Notices and Aggregate Job Loss." *Journal of Applied Econometrics*, 39(3), pp. 462--480. doi:10.1002/jae.3032.

Kuriwaki, Shiro (2026). "Cumulative CES Common Content." doi:10.7910/DVN/II2DB6.

Leip, David (2026). "Dave Leip's Election Data." https://uselectionatlas.org/BOTTOM/store_data.php. Accessed July 2026.

Margalit, Yotam (2011). "Costly Jobs: Trade-related Layoffs, Government Compensation, and Voting in US Elections." *American Political Science Review*, 105(1), pp. 166--188. doi:10.1017/S000305541000050X.

Margalit, Yotam (2013). "Explaining Social Policy Preferences: Evidence from the Great Recession." *American Political Science Review*, 107(1), p. 80–103. doi:10.1017/S0003055412000603.

Margalit, Yotam (2019). "Political Responses to Economic Shocks." *Annual Review of Political Science*, 22(Volume 22, 2019), pp. 277--295. doi:https://doi.org/10.1146/annurev-polisci-050517-110713.

Petrocik, John R. (1996). "Issue Ownership in Presidential Elections, with a 1980 Case Study." *American Journal of Political Science*, 40(3), pp. 825--850.

QWI (2026). "Quarterly Workforce Indicators." Longitudinal Employer-Household Dynamics Program. URL https://lehd.ces.census.gov/data/, accessed August 2026.

Rehm, Philipp, Jacob S. Hacker, and Mark Schlesinger (2012). "Insecure Alliances: Risk, Inequality, and Support for the Welfare State." *American Political Science Review*, 106(2), pp. 386--406. doi:10.1017/S0003055412000147.

Rickard, Stephanie J. (2022). "Incumbents Beware: The Impact of Offshoring on Elections." *British Journal of Political Science*, 52(2), pp. 758--780. doi:10.1017/S0007123420000757.

Riker, William H. and Peter C. Ordeshook (1968). "A Theory of the Calculus of Voting." *American Political Science Review*, 62(1), pp. 25--42. doi:10.2307/1953324.

Rosenstone, Steven J. (1982). "Economic Adversity and Voter Turnout." *American Journal of Political Science*, 26(1), pp. 25--46. doi:10.2307/2110837.

Rosenstone, Steven J. and John Mark Hansen (1993). *Mobilization, Participation, and Democracy in America*. Macmillan, New York.

Schmieder, Johannes F., Till von Wachter, and Jorg Heining (2023). "The Costs of Job Displacement over the Business Cycle and Its Sources: Evidence from Germany." *American Economic Review*, 113(5), pp. 1208--1254. doi:10.1257/aer.20200252.

SEER (2026). "US County Population Data." https://seer.cancer.gov/popdata/download.html. Accessed July 2026.

Squire, Peverill, Raymond E. Wolfinger, and David P. Glass (1987). "Residential Mobility and Voter Turnout." *American Political Science Review*, 81(1), pp. 45--65. doi:10.2307/1960778.

Upward, Richard and Peter W. Wright (2024). "Income Shocks, Political Support and Voting Behaviour." *Journal of Public Economics*, 239, p. 105,253. doi:10.1016/j.jpubeco.2024.105253.

USDA (2026). "County-Level Data Sets." https://www.ers.usda.gov/data-products/county-level-data-sets. Accessed July 2026.

Table 1: Treatment characteristics, electoral outcomes, matched support, and balance

Panel A. Sizable WARN event characteristics (matched events; all treated events in parentheses)

	Mean	Median	P25	P75
WARN layoffs in treated event	431 (721)	206 (222)	123 (128)	418 (483)
WARN layoffs as percent of labor force	1.811 (1.799)	1.145 (1.124)	0.811 (0.804)	1.906 (1.883)
Number of WARN notices in treated event	2.21 (4.34)	1.00 (1.00)	1.00 (1.00)	2.00 (2.00)
Months from WARN event to next election	11.5 (11.4)	10.0 (10.0)	7.0 (7.0)	18.0 (17.0)

Panel B. Electoral outcomes (for matched events)

	Mean	Median	P25	P75
Turnout	50.2	50.8	42.7	58.3
Democratic votes/CVAP	19.4	18.7	12.7	25.2
Republican votes/CVAP	29.4	29.2	22.7	35.7
Other votes/CVAP	1.4	0.9	0.1	1.7
Incumbent-candidate share	56.0	56.8	44.3	67.5
Pre-election office-holder-party share	54.1	54.6	42.3	66.1

Panel C. Matched analysis sample

	Value
Treated events	1,703
Treated counties	1,300
Treated states	36
Matched comparisons	1,508
Mean matched controls per treated event	3.46
Share matched in exact state-year-month cell	96.9
Share matched in fallback state-year cell	3.1
County-office-election obs. in baseline estimation	122,698

Panel D. Balance on matching variables

	Treated mean	Eligible control mean	Matched control mean	Difference	Absolute std. diff.
Turnout, last pre-event election	49.254	50.829	49.636	-0.382	0.032
Republican vote share, last pre-event election	59.197	59.739	59.723	-0.526	0.041
Republican vote share, second pre-event election	57.146	57.710	57.664	-0.518	0.040
Labor force	38,421	51,579	37,220	1,200	0.013
Unemployment rate	5.666	5.510	5.578	0.087	0.039
Change in unemployment rate	0.508	0.387	0.439	0.070	0.029
Prior WARN exposure 12-months before treatment	0.005	0.005	0.004	0.001	0.051
Prior WARN exposure 13 to 24 months before treatment	0.004	0.005	0.004	0.000	0.019

Note: This table summarizes the treated events, electoral outcomes, matched comparisons, and variable balance in the baseline event-study sample. Panel A reports characteristics of sizable WARN events, defined using the three-month rolling exposure measure in Equation (2), discussed in Section 2.1. "P25" and "P75" denote the 25th and 75th percentiles, respectively. Statistics for matched events are shown first; statistics for all treated events appear in parentheses. Panel B reports unweighted summary statistics for successfully matched treated observations at the last election across all offices before the WARN event; the two accountability rows use their own available samples. Panel C reports matched-support statistics, with details about how we construct treated-control comparisons in Section 2.2. We select up to four controls for each treated event. Exact-cell matches align treated and control counties by state, treatment election year, and treatment month; fallback-cell matches preserve state and treatment-year alignment while relaxing the exact treatment-month requirement. The 1,508 matched comparisons contain one treated county and 3.46 controls on average, or about 6,730 comparison-county units. Each comparison-county contributes roughly 18 office-election observations. Thus, $6{,}730 \times 18.1 \approx 122{,}000$ observations. Panel D reports balance on the variables used in the baseline nearest-neighbor matching procedure. We use the following variables: recent turnout, the two most recent Republican vote-share observations, log labor force (displayed here in levels), the unemployment rate, the change in the unemployment rate, and WARN exposure over the previous year and the year before that, as discussed in Section 2.2. Republican vote share is defined as one hundred times the number of Republican votes divided by the number of Republican, Democratic, and other votes. Eligible controls are counties satisfying the control-sample restrictions before nearest-neighbor selection. Matched controls are the up to four nearest-neighbor eligible controls that survive maximum-difference limits on each matching variable between treated and potential controls. Matched-control means use the inverse-distance weights defined in Appendix B.2. Differences are treated means minus matched-control means. Absolute standardized differences divide the absolute treated-minus-matched difference by the pooled standard deviation of the treated and matched-control observations. The turnout measures in Panels B and D differ slightly because Panel B uses all contests and Panel D prioritizes presidential, gubernatorial, Senate, and House elections when several offices are contested. Section 2.3 discusses the results in the table in more detail.

Table 2: Main results, heterogeneity, geographic scope, endogeneity, and robustness

Dependent variable ⟶	Democratic-Republican Votes/CVAP		Republican Votes/CVAP		Democratic Votes/CVAP		Other votes/CVAP		Turnout	
Specification ↓ / Independent variable ⟶	Pre	Post	Pre	Post	Pre	Post	Pre	Post	Pre	Post
Panel A: Main results										
(1) Baseline	0.041	0.501***	-0.032	-0.318***	0.010	0.182**	0.001	0.000	-0.021	-0.136
	(0.133)	(0.169)	(0.079)	(0.105)	(0.080)	(0.088)	(0.028)	(0.025)	(0.083)	(0.093)
(2) Residualized WARN exposure	-0.165	0.728***	0.136*	-0.451***	-0.029	0.277***	-0.027	0.006	0.080	-0.168*
	(0.140)	(0.178)	(0.080)	(0.106)	(0.084)	(0.097)	(0.027)	(0.025)	(0.083)	(0.095)
(3) Net of job destruction	0.103	0.464**	-0.082	-0.282**	0.020	0.182*	0.011	0.006	-0.051	-0.094
	(0.143)	(0.184)	(0.087)	(0.115)	(0.086)	(0.094)	(0.031)	(0.027)	(0.090)	(0.100)
Panel B: Election timing heterogeneity										
(4) $\tau = 0$	-0.068	0.530***	0.056	-0.336**	-0.012	0.194*	-0.018	-0.003	0.026	-0.145
	(0.161)	(0.201)	(0.101)	(0.131)	(0.098)	(0.106)	(0.034)	(0.031)	(0.109)	(0.123)
(5) $\tau = 1$	0.172	0.446*	-0.161	-0.277*	0.011	0.169	0.033	0.004	-0.117	-0.104
	(0.208)	(0.251)	(0.119)	(0.147)	(0.125)	(0.134)	(0.046)	(0.041)	(0.121)	(0.123)
Panel C: Industry heterogeneity										
(6) Nonmanufacturing	0.013	0.475***	0.037	-0.247**	0.050	0.228**	-0.053*	-0.001	0.033	-0.020
	(0.133)	(0.156)	(0.078)	(0.097)	(0.081)	(0.089)	(0.031)	(0.030)	(0.084)	(0.096)
Panel D: Geographic scope and spillovers										
(7) Donut	0.025	0.489***	-0.016	-0.302***	0.009	0.187**	-0.013	-0.003	-0.020	-0.119
	(0.138)	(0.172)	(0.082)	(0.108)	(0.083)	(0.090)	(0.030)	(0.026)	(0.085)	(0.096)
Panel E: Local deterioration & estimate validity										
(8) No controls	-0.066	0.499***	0.028	-0.344***	-0.039	0.155*	0.011	0.009	0.000	-0.180**
	(0.133)	(0.163)	(0.078)	(0.100)	(0.080)	(0.087)	(0.028)	(0.024)	(0.083)	(0.091)
(9) Richer controls	-0.016	0.596***	-0.010	-0.349***	-0.026	0.247***	0.002	0.007	-0.034	-0.096
	(0.132)	(0.170)	(0.079)	(0.107)	(0.080)	(0.088)	(0.028)	(0.025)	(0.083)	(0.092)
(10) Rematching trajectory	-0.089	0.317*	0.059	-0.238**	-0.030	0.079	0.002	0.004	0.030	-0.155*
	(0.149)	(0.182)	(0.084)	(0.109)	(0.085)	(0.095)	(0.023)	(0.022)	(0.076)	(0.093)
(11) Forecasted deterior.	0.011	0.524***	-0.020	-0.329***	-0.009	0.194**	0.001	0.001	-0.028	-0.134
	(0.133)	(0.169)	(0.079)	(0.106)	(0.080)	(0.088)	(0.028)	(0.026)	(0.082)	(0.092)
Panel F: Robustness										
(12) 80th percentile	-0.101	0.587***	0.047	-0.367***	-0.053	0.220***	0.000	0.035	-0.006	-0.112
	(0.125)	(0.151)	(0.071)	(0.093)	(0.075)	(0.081)	(0.028)	(0.026)	(0.073)	(0.085)
(13) 90th percentile	0.042	0.405**	-0.073	-0.257**	-0.031	0.147	0.039	0.032	-0.066	-0.078
	(0.145)	(0.166)	(0.088)	(0.105)	(0.085)	(0.091)	(0.032)	(0.027)	(0.091)	(0.102)
(14) 25-worker floor	0.051	0.489***	-0.046	-0.306***	0.005	0.183**	-0.007	-0.004	-0.048	-0.127
	(0.131)	(0.166)	(0.079)	(0.103)	(0.079)	(0.087)	(0.028)	(0.025)	(0.083)	(0.092)
(15) 75-worker floor	0.030	0.413**	-0.014	-0.256**	0.016	0.157*	0.001	0.001	0.003	-0.098
	(0.136)	(0.167)	(0.080)	(0.104)	(0.082)	(0.089)	(0.028)	(0.025)	(0.083)	(0.093)
(16) 1-Month window	-0.040	0.472***	0.009	-0.310***	-0.031	0.162**	-0.021	0.014	-0.042	-0.134
	(0.123)	(0.148)	(0.073)	(0.092)	(0.073)	(0.080)	(0.027)	(0.024)	(0.076)	(0.084)
(17) 6-Month window	0.037	0.604***	-0.008	-0.353***	0.030	0.251***	0.007	0.021	0.029	-0.081
	(0.144)	(0.168)	(0.084)	(0.105)	(0.086)	(0.089)	(0.028)	(0.026)	(0.086)	(0.096)
(18) Up to 10 controls	-0.052	0.455***	0.021	-0.277***	-0.031	0.177**	-0.001	0.004	-0.011	-0.096
	(0.131)	(0.163)	(0.078)	(0.102)	(0.079)	(0.085)	(0.028)	(0.025)	(0.081)	(0.091)
(19) Comparison-index FE	-0.027	0.484***	0.007	-0.312***	-0.020	0.172**	-0.005	0.003	-0.019	-0.137
	(0.120)	(0.149)	(0.072)	(0.091)	(0.074)	(0.082)	(0.027)	(0.024)	(0.080)	(0.085)
(20) $\ln(\text{CVAP})$ weights	0.026	0.505***	-0.021	-0.327***	0.004	0.178**	-0.007	-0.003	-0.024	-0.152
	(0.132)	(0.173)	(0.079)	(0.109)	(0.080)	(0.089)	(0.028)	(0.025)	(0.083)	(0.094)
(21) Exclude 2020 treatments	-0.053	0.499***	0.023	-0.315***	-0.030	0.183*	0.002	0.018	-0.006	-0.114
	(0.150)	(0.182)	(0.087)	(0.107)	(0.090)	(0.100)	(0.032)	(0.028)	(0.093)	(0.095)

Note: This table reports results from separate specification, sample, and diagnostic exercises. Each row of the table reports effects, in percentage points, for five dependent variables: Democratic-Republican votes per CVAP, Republican votes per CVAP, Democratic votes per CVAP, other votes per CVAP, and turnout, as defined in Equation (1). For each of these dependent variables, columns labeled "Pre" report effects on $D_{cg} \times \mathbf{1}\{h \in \{-3, -2, -1\}\}$, and columns labeled "Post" report effects on $D_{cg} \times \mathbf{1}\{h \in \{1, 2, 3, 4\}\}$, in which h is the election index, as discussed in Section 2.4. The omitted election is $h = 0$, the last election before the sizable WARN event. Standard errors, reported in parentheses, are two-way clustered by county and matched comparison group. Panel A is discussed in Section 2.4. Row 1 reports coefficients from the baseline specification, Equation (3). Row 2 reconstructs the treatment indicator and treated-control comparisons using residual WARN exposure. Row 3 allows the electoral path to vary with realized county job destruction in the sample with available job-destruction data. Panels B, C, and D are discussed in Section 3.3. Rows 4 and 5 show the results when we divide treated events into those that occur within one year of the next election ($\tau = 0$) or one to two years before it ($\tau = 1$). Row 6 restricts attention to nonmanufacturing WARN notices and reconstructs the treatment indicator and treated-control comparisons using those notices. Row 7 shows the results from a donut exercise that excludes potential control counties in the treated county's commuting zone. Panel E reports a set of checks that vary how we account for local economic conditions, as discussed in Section 3.4.1. Rows 8 and 9 show the results when we omit controls or add richer local controls to the baseline specification. Rows 10 and 11 adjust for longer pre-event economic, demographic, and sectoral histories. Panel F reports results across various specification, sample, and design choices, as discussed in Section 3.4.2. Rows 12 and 13 vary the thresholds for defining a sizable WARN event. Rows 14 and 15 vary the minimum-worker floor. Rows 16 and 17 vary the rolling exposure window. Row 18 allows up to 10 controls. Row 19 replaces the baseline event-index and state-by-office-by-year fixed effects with comparison-by-election-index-by-office fixed effects. Row 20 reweights by the natural logarithm of CVAP. Row 21 excludes WARN events assigned to treatment election-year 2020. Except for Rows 4, 5, and 18, the regressions use between 97,889 and 143,898 observations across specifications. Rows 4 and 5 use 71,508 and 51,190 observations, respectively, while Row 18 uses 228,052. $*$, $**$, and $***$ denote significance at the 10, 5, and 1 percent levels.

Table 3: Evidence on electoral mechanisms and policy outcomes

Panel A: Demographics and migration

Dependent variable ⟶	Democratic-Republican Votes/CVAP		Republican Votes/CVAP		Democratic Votes/CVAP		Other votes/CVAP		Turnout	
Independent variable ⟶	Pre	Post	Pre	Post	Pre	Post	Pre	Post	Pre	Post
	0.041	0.461***	-0.028	-0.274***	0.012	0.187**	0.004	0.011	-0.012	-0.076
	(0.131)	(0.171)	(0.079)	(0.102)	(0.079)	(0.093)	(0.028)	(0.025)	(0.083)	(0.089)

Panel B: CES

Dependent variable ⟶	R intent and D vote		D intent and R vote		R--D validated turnout gap		R--D val. turnout gap among those intending to vote	
Independent variable ⟶	Pre	Post	Pre	Post	Pre	Post	Pre	Post
	0.469	-0.220	-0.317	-0.342	-0.301	-4.150*	1.484	-5.110*
	(0.423)	(0.410)	(0.589)	(0.550)	(2.357)	(2.324)	(2.881)	(2.942)

Panel C: Campaign visits

Dependent variable ⟶	Democratic minus Republican any visit		Democratic minus Republican visit count	
Independent variable ⟶	Pre	Post	Pre	Post
	0.104	0.344	-0.017	0.007
	(0.961)	(0.864)	(0.017)	(0.012)

Panel D: Real local transfer payments per capita (natural logarithm)

Dependent variable ⟶	Shock-responsive safety net		Active and administratively mediated support		Slow-moving categorical entitlements		Private and community transfers	
Independent variable ⟶	Pre	Post	Pre	Post	Pre	Post	Pre	Post
	-0.001	0.005***	-0.003	0.000	0.001	0.001	-0.001	0.003*
	(0.001)	(0.002)	(0.004)	(0.005)	(0.001)	(0.001)	(0.001)	(0.002)

Note: This table reports coefficients for four different sets of exercises. Panel A reports coefficients from estimating Equation (3) with pre- and post-coefficients and controls for county-level migration flows and demographic structure for five outcome variables: Democratic-Republican votes per CVAP, Republican votes per CVAP, Democratic votes per CVAP, other votes per CVAP, and turnout. Panel B uses CES individual-level voting data and estimates Equation (3) at the individual level. In the first (second) column, we restrict the sample to respondents who intended to vote Republican (Democratic) for the relevant office and later reported voting in the election. The outcome equals one if the respondent reports voting Democratic (Republican). In the third column, we restrict the sample to major-party leaners and use as the outcome a dummy variable for officially validated respondent turnout. We estimate an augmented version of Equation (3) in which the pre- and post-treatment dummies are also interacted with whether the respondent leans Republican. The reported point estimates correspond to the coefficients on these interactions. The fourth column uses the same outcome variable and specification as the third column but restricts the sample to respondents intending to vote. All four Panel B specifications include controls for respondent age, age squared, and sex. Panel C shows how vice-presidential and presidential candidate visits respond to WARN events using Equation (3). The first outcome is an indicator for any Democratic visit minus an indicator for any Republican visit. The second is the number of Democratic visits minus the number of Republican visits. Section 4.1 discusses the results in Panels A, B, and C. Panel D reports the effects of WARN events on the natural logarithm of one plus real local transfer payments per capita in 2022 dollars. The first column uses shock-responsive categories as the outcome, including unemployment insurance, SNAP, the Earned Income Tax Credit, and other income-maintenance and public-assistance benefits. The second column uses active and administratively mediated support as the outcome, including education and training assistance and other government transfers to individuals. The third column uses slow-moving categorical entitlements as the outcome, including Social Security, retirement and disability, Medicare, Supplemental Security Income, and veterans' benefits. The fourth column uses private and community transfers as the outcome, including transfers from businesses to individuals and transfers to nonprofit institutions. Panels A and B use the chronological index: Pre combines $\{h \in \{-3,-2,-1\}\}$, Post combines $\{h \in \{1,2,3,4\}\}$, and $h=0$ is omitted. Panel C uses the presidential target-office index: Pre combines $\{h \in \{-2,-1\}\}$, Post combines $\{h \in \{1,2,3\}\}$, and $h=0$ is omitted. Panel D uses annual event time: Pre combines $\{k \in \{-3,-2\}\}$, Post combines $\{k \in \{0,1,2,3\}\}$, and $k=-1$ is omitted. Standard errors, reported in parentheses, are two-way clustered by county and matched comparison group. Panel A uses 114,969 observations. Panel B uses 368,146, 314,945, 1,049,470, and 721,123 observations across its four outcomes. Panel C uses 27,121 observations, and Panel D uses 41,651 observations. Section 4.2 discusses the results in Panel D. $*$, $**$, and $***$ denote significance at the 10, 5, and 1 percent levels.

Fig. 1: Effects of WARN events on the Democratic-Republican vote-rate margin, party-specific vote rates, and turnout

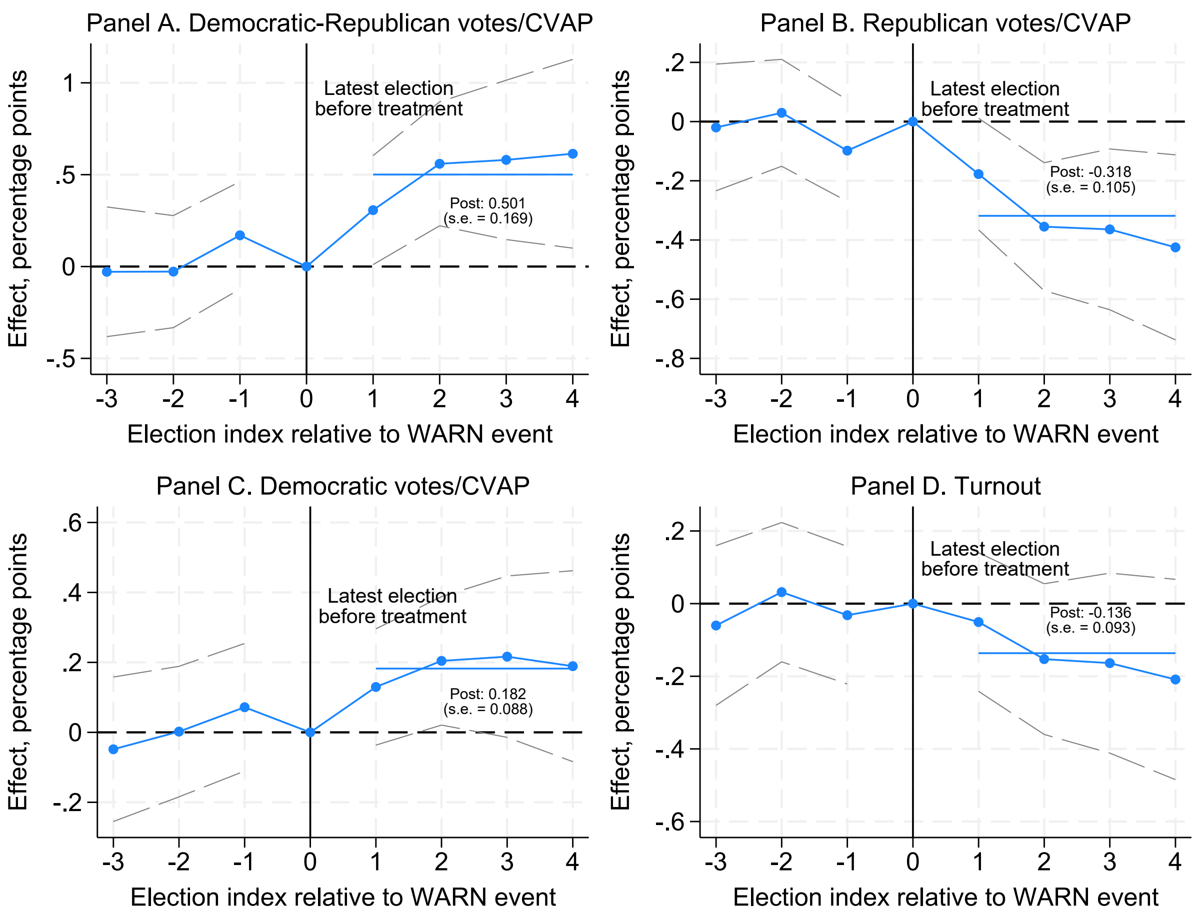

Note: This figure reports event-study estimates for four electoral outcomes using Equation (3) in Section 2.4: Democratic-Republican votes per CVAP (Panel A), Republican votes per CVAP (Panel B), Democratic votes per CVAP (Panel C), and turnout (Panel D). All outcomes are measured in percentage points. The election index on the x-axis follows the chronological design discussed in Section 2.4. That is, $h = 0$ denotes the last election before the WARN event, $h = 1$ denotes the first election after the WARN event, and so on. The omitted election index is $h = 0$. This design is office-agnostic, and so election outcomes are compared to earlier election outcomes without requiring the office to be the same. We retain post-event indexes through $h = 4$, corresponding to a horizon of up to eight years. The results use data from four election types: presidential, gubernatorial, US House and US Senate elections. The text in each panel reports the post-coefficient estimate from a specification that includes only one post-event coefficient on $D_{cg} \times \mathbf{1}\{h \in \{1, 2, 3, 4\}\}$ (standard errors in parentheses), in which D_{cg} equals one for the treated county in comparison group g and zero for its matched controls. That specification also includes only one pre-event coefficient on $D_{cg} \times \mathbf{1}\{h \in \{-3, -2, -1\}\}$. The regressions use 122,698 observations. Dashed gray lines represent 95 percent confidence intervals based on standard errors two-way clustered by county and matched comparison group. Section 3.1 describes the results.

Fig. 2: Effects of WARN events by election type

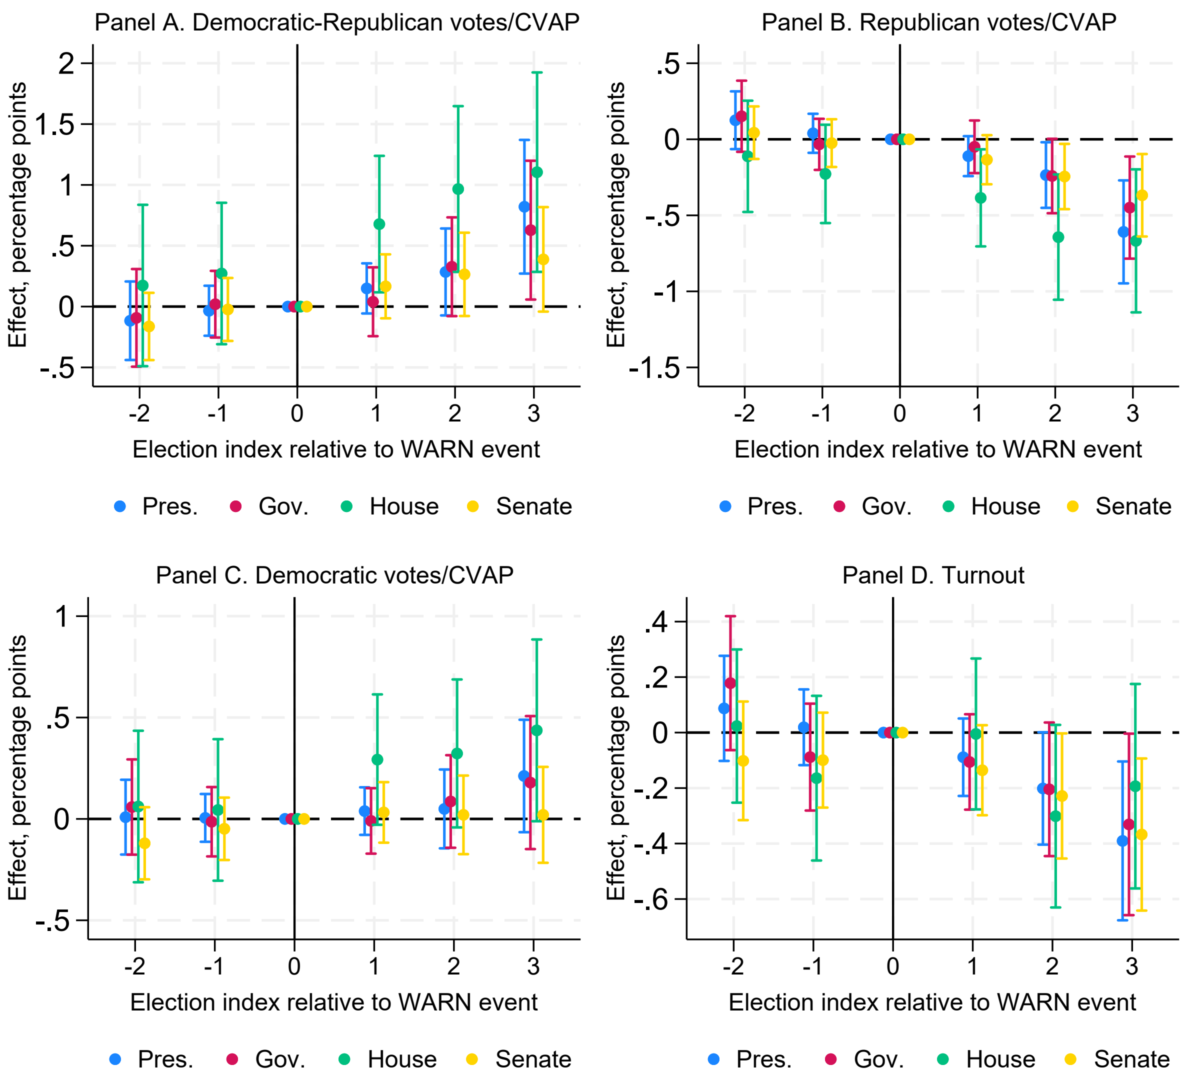

Note: This figure reports event-study estimates by election type for four electoral outcomes using Equation (3) in Section 2.4: Democratic-Republican votes per CVAP (Panel A), Republican votes per CVAP (Panel B), Democratic votes per CVAP (Panel C), and turnout (Panel D). All outcomes are measured in percentage points. The election types include presidential, gubernatorial, US House, and US Senate elections. The election index on the x axis follows the office-specific design discussed in Section 2.4. This design defines the h election index relative to the specific office contests. That is, $h = 0$ denotes the last election of a particular office before the sizable WARN event, $h = 1$ denotes the first election of the same office after the event, and so on. The omitted election index is $h = 0$, denoted by the vertical black line. We retain post-event indexes through $h = 3$, corresponding to a horizon of up to 12 years. The regressions use between 30,864 and 37,833 observations across offices and outcomes. Vertical whisker lines represent 95 percent confidence intervals based on standard errors two-way clustered by county and matched comparison group. Section 3.2 describes the results.

Online Appendix to "From Pink Slips to Polling Stations: The Electoral Consequences of Mass-Layoff Announcements"

Victor Hernandez Martinez[*] Pawel M. Krolikowski[†]
Andrew H. McCallum[‡]

August 11, 2026

A Data appendix

A.1 WARN Act details and data construction and coverage

The federal WARN Act was passed on August 4, 1988, and became effective on February 4, 1989. The act seeks to provide workers with sufficient time to begin new job searches or to obtain necessary training. The act requires employers with 100 or more employees to provide their workers with at least 60 days' written notice prior to layoff. WARN notification is triggered by large, permanent reductions in the labor force at individual employment sites, defined in two ways. First, a "plant closing" is when an employment site is shut down and 50 or more full-time workers lose their jobs. Second, a "mass layoff" happens when, during a 30-day period, employment loss affects 500 or more full-time workers at a given site, or 50 to 499 full-time workers if they make up at least one-third of the workforce at that site. The act also applies when an employer announces that a temporary layoff (less than 6 months) becomes a permanent one. The act requires notices to include the name and address of the affected employment site, the expected date of the first separation, and the anticipated number of affected employees. There are exceptions to issuing 60 days' notice. And some states and localities have passed legislation that expands the federal WARN Act so that covered employers and mass-layoff events may vary across geographies. See US Department of Labor (2026) and state legislation for more details.

We collect establishment-level WARN notices from state employment offices through their websites and by contacting state officials. We often extend publicly available WARN data by using digital archives of the internet and documents from state officials. The resulting panel is unbalanced across states because of data availability. See Krolikowski and Lunsford (2024, Appendix A.2) for more details about the data collection process.

Our WARN data begin as early as 1990 and extend through the end of 2025. Michigan is the first state to enter our sample in January 1990, and by November 1996, we have 5 states; by November 2006, we have 26 states; and by November 2016, we have 37 states. Table A1 presents the list of states in our WARN data and some summary statistics about WARN layoffs in these states.

We map addresses from establishment-level notices to counties. As mentioned in Section 1.1, each WARN notice typically includes the name and address of the employer,

[*]victor.hernandezmartinez@clev.frb.org, Federal Reserve Bank of Cleveland, 1455 E 6th St, Cleveland, OH 44114

[†]Corresponding author, pawel.krolikowski@clev.frb.org, Federal Reserve Bank of Cleveland, 1455 E 6th St, Cleveland, OH 44114

[‡]andrew.h.mccallum@frb.gov, Board of Governors of the Federal Reserve System, Washington, DC 20551

the date the notice was filed, the anticipated layoff date, and the estimated number of affected workers. We clean all employer addresses using a combination of Bing searches and manual cleaning. We use a Census geocoder to map street addresses to counties (US Census Bureau, 2026). Some notices include an associated city or ZIP code, but no street address. For these notices, we use a crosswalk to map cities and ZIP codes to counties (US Department of Housing and Urban Development, 2026). If a city or ZIP code maps to more than one county, we use county-specific business and address counts for that city or ZIP code to select a unique county.

A.2 Election returns, office calendar, party harmonization, CVAP, and candidate links

We construct a county-office-year election panel from US Elections Atlas returns for presidential, gubernatorial, US House, and US Senate elections from 1996 through 2024. We align the returns to an election calendar that records office, election year, Senate class, and House district. We retain regular general elections held from November 1 through November 8 and exclude special and recall elections. These restrictions remove 35 of 7,436 calendar events, or 0.47 percent; another 32 events, or 0.43 percent, lack usable county-level returns or positive CVAP, leaving 7,369 events represented in the county panel.

Table A2 reports election events and geographic and CVAP coverage by year and office. The sample includes eight presidential elections, 490 Senate elections, 361 gubernatorial elections, and 6,510 House district elections. The state and county columns report the geographic coverage of each office-year cell. The CVAP column displays the total unique citizen voting-age population called to vote in that type of election in a given year.

From the raw election data, we harmonize the two inputs used to construct the main outcomes: party vote totals and county CVAP. For party vote, presidential files identify Democratic and Republican votes directly. For Senate, House, and gubernatorial elections, we classify vote columns using candidate or party labels: Democratic, DFL, Farmer-Labor, and analogous labels count as Democratic votes; Republican, GOP, and analogous labels count as Republican votes; and independent candidates, write-ins, state-specific parties, and residual categories count as other-candidate votes.

We construct annual county CVAP from the 1990 and 2000 Censuses and annual five-year ACS estimates beginning in 2009. Between official observations, we interpolate the county CVAP-to-voting-age-population ratio and apply it to annual intercensal estimates of the adult population. For November elections, we assign the population estimate from the concurrent calendar year.

After merging CVAP into the election panel, turnout equals total votes divided by CVAP, multiplied by 100. Democratic, Republican, and other-candidate vote rates use each party group's votes in the numerator and therefore sum to turnout, as discussed in Section 1.2. Within each election year, we trim the top and bottom 0.25 percent of turnout and apply the resulting common sample to all electoral variables.

While House elections are determined at the congressional district level, our raw election-outcome data provide only aggregate party votes across the House contests represented within a county. That is, in our data, we observe one county-House-year outcome for turnout and each party-specific vote rate. Table A2 counts the underlying district contests, while its states, counties, and CVAP columns describe the county-level

House outcomes. This aggregation preserves county-level participation measures but limits the cases in which an individual House incumbent can be linked unambiguously to the outcome.

For the accountability outcomes, we separately identify incumbent candidates and the party holding each office before the election. Incumbent-candidate information for presidential, Senate, and House elections comes from Federal Election Commission records; gubernatorial information comes from a manually assembled set of candidate and incumbent records. We harmonize party identifiers to Democratic and Republican categories and construct an incumbent-candidate vote share when a major-party incumbent is running.[1]

The pre-election office-holder-party outcome is defined whether or not the office holder is running. For presidential elections, we use the sitting president's party. For gubernatorial and Senate elections, we infer the party from earlier election results, supplemented by the incumbent candidate's party when prior-election history is unavailable. We then measure the vote share received by the pre-election office-holder's party.

For House elections, we assign an incumbent candidate or pre-election office-holder party only when every regular House contest in the state-year has a resolved Democratic or Republican incumbent and all incumbents belong to the same party; otherwise, the identifiers remain undefined.

Table A3 reports identifier availability among county-office-year observations. Once we observe an incumbent candidate, the candidate's Democratic or Republican affiliation is available for every office. We observe incumbent candidates in 50.0 percent of presidential observations, reflecting that the sitting president runs in four of the eight elections in the sample, 58.1 percent of gubernatorial observations, and 86.5 percent of Senate observations. The pre-election office-holder party is available more broadly: for all presidential observations, 94.1 percent of gubernatorial observations, and 96.8 percent of Senate observations. This variable is more widely available than the incumbent candidate identifier because the party holding office remains defined in open-seat elections. House identifiers are available for 14.4 percent of observations because we require an unambiguous assignment when county returns aggregate votes across congressional districts, rather than because Democratic or Republican vote totals are missing.

A.3 Auxiliary data and controls

We supplement the WARN and election data with county-level measures of labor-market conditions, income, population composition, migration, and COVID exposure. The Bureau of Labor Statistics Local Area Unemployment Statistics provide monthly county labor force, employment, unemployment, and unemployment rates from 1994 through 2024. We use labor force to scale WARN-affected workers; pre-event labor-force size, unemployment, and changes in unemployment enter the matching procedure. The Bureau of Economic Analysis Regional Economic Accounts provide annual county per capita personal income.

[1] We restrict our sample of incumbents to Republican and Democratic incumbents due to data limitations. This is not a stringent restriction. In more than 95 percent of the elections in our sample, the party receiving the most votes is either Republican or Democratic.

The baseline controls include the county unemployment rate, labor-force size, per capita personal income, and cumulative COVID cases per 1,000 residents. Labor-market conditions are measured in the election month. County COVID data come from Killeen et al. (2020), with cumulative cases measured through October 31 of the election year and set to zero before 2020.

Specifications with richer controls and matching on a broader set of county characteristics draw on additional annual data. The Quarterly Census of Employment and Wages provides county covered employment, employment by sector, average weekly wages, and establishment counts. These data yield manufacturing and goods-producing employment shares, average weekly wages, and average establishment size, while Census land area and BEA resident population are used to construct population density.

The Census Bureau's Quarterly Workforce Indicators (QWI) provide quarterly county-level measures of employment flows and firm job destruction. We use these data to examine whether WARN events are followed by realized local job destruction. Because these data cover all private-sector firms in the county, they do not isolate losses at establishments named in WARN notices.

Annual demographic and migration data allow us to examine changes in county population composition. SEER and Census estimates provide population counts by age and race or ethnicity, from which we construct the shares age 65 or older, age 25 to 54, and nonwhite or Hispanic. Census and ACS data provide voting-age population and the ratio of CVAP to voting-age population. Internal Revenue Service Statistics of Income files report county inflows and outflows, from which we construct net and gross migration rates relative to prior population.

To study mechanisms, we use the Cooperative Election Study, a large national survey of political attitudes and voting behavior conducted around federal elections. Its cumulative common-content files from 2006 through 2024 provide respondent county, survey weights, demographics, party identification, pre-election vote intent, reported vote choice, and validated turnout, allowing us to distinguish changes in electoral composition from vote switching. The Campaign Visits Database (Devine, 2020) records presidential and vice-presidential candidate appearances by county and date from 2008 through 2024.

County transfer payments through 2022 come from the Bureau of Economic Analysis Regional Economic Accounts, where they are reported as personal current transfer receipts. We group the detailed components into shock-responsive safety-net transfers, including unemployment insurance, SNAP, the Earned Income Tax Credit, and other income-maintenance programs; education, training, and other actively administered assistance; slower-moving entitlement transfers, including retirement, disability, Medicare, veterans, and Social Security programs; and transfers to nonprofit institutions and transfers to individuals from businesses.

Finally, we classify WARN notices as manufacturing or nonmanufacturing using state-reported industry codes when available and manual and AI-assisted assignments based on the employer and information in the notice.

B Empirical strategy appendix

B.1 Treatment definition and election-year windows

Section 2.1 defines sizable WARN events using the number of workers covered by WARN notices over three consecutive months, scaled by the county labor force. We construct this measure only when WARN and labor-force data are observed in all three months. A county-month qualifies when the resulting exposure is at or above the 85th percentile of positive three-month exposure windows, and the notices cover at least 50 workers. In the baseline sample, the percentile cutoff corresponds to approximately 0.6 percent of the county labor force. The percentile cutoff selects events that are large relative to the local labor market, while the worker floor ensures they are also meaningful in absolute terms.

For an election in year y, the corresponding election-year window runs from November of year $y-1$ through October of year y. A qualifying event during this window is assigned to the regular November election that follows. If a county crosses the threshold in more than one month within the same election-year window, we treat these crossings as a single event dated by the first qualifying month.

To distinguish new shocks from ongoing large layoff activity, we require the four preceding election-year windows to be observed and contain no other qualifying event. Counties may still experience smaller WARN announcements during the lookback period; conditional on complete coverage, only an earlier event that also meets the treatment threshold makes the county ineligible.

Treatment is therefore defined at the county-event level rather than as a permanent county status. The same county can enter the treatment sample again after four consecutive, fully observed election-year windows without another qualifying event. Figure A1 summarizes the construction: the three-month window measures the size of the event, the first threshold-crossing month dates the event, and the November-October window links it to the election that follows. The June treatment month in the figure is illustrative; a treatment event can occur in any month from November through October of the following calendar year.

B.2 Control eligibility, matching cells, calipers, and weights

The matched design compares each treated county-event with counties observed in the same calendar and electoral setting and with similar political and economic conditions before the WARN event. Potential controls may experience smaller WARN announcements, but they cannot experience a sizable event over the comparison window. Conceptually, the comparison is therefore between a sizable local layoff event and the absence of a similarly sizable event, rather than between any WARN activity and none. In practice, however, average control-group WARN exposure remains near zero throughout the event window, while treated-county exposure falls close to that level after the qualifying month, as Figure A2 shows.

For a treated event in election year y, a potential control must have complete WARN coverage and no qualifying event in the treatment election-year window, the four preceding election-year windows, and four following election-year windows. For events near the end of the sample, the four following election-year windows may not be available. In those cases, the forward restriction applies only to every subsequent window observed through 2024. The backward restriction rules out controls whose pre-event histories may already reflect another sizable WARN event, while the forward restriction keeps controls from receiving the same treatment during the four following election-year windows.

We first search for controls in the same state, the same treatment election year, and the same treatment month. These cells compare counties facing the same broad political environment and observed at the same point in the election cycle and national economy. If no county in this exact cell survives all of the matching restrictions, we broaden the search within the same state and treatment election year. The fallback preserves the office used to define the pre-event political history but relaxes exact treatment-month alignment. It is used only when the exact cell produces no acceptable control. Only 3 percent of matched treated events rely on this fallback.

We compare counties on prior turnout, two prior Republican vote shares, log labor-force size, unemployment, the change in unemployment, and two measures of prior WARN exposure. The political history is anchored in the November election immediately preceding treatment. If several offices are contested on that date, we use a common ordering (presidential, gubernatorial, Senate, and House) so that every county in the cell is evaluated relative to the same election. Prior turnout and the first Republican-share measure come from that election; the second Republican-share measure comes from the preceding election for the selected office.

The economic measures summarize conditions before treatment. We calculate log labor-force size as the natural logarithm of the average monthly labor force over the preceding 12 months and average the unemployment rate over the same period. The unemployment change compares the month immediately before treatment with the same month one year earlier. We measure WARN history over two successive 12-month periods preceding the three-month treatment window. The more recent measure captures WARN activity immediately before treatment, while the earlier measure helps rule out counties with persistently different histories of announced layoffs.

We standardize these variables in the pool of match-ready treated events and eligible controls. We then impose calipers requiring treated and control counties to be within one standard deviation on the political and labor-market measures and within 1.5 standard deviations on each WARN-history measure. Among the counties that satisfy these restrictions, we rank controls by their standardized Euclidean distance in prior turnout, the two Republican-share measures, log labor-force size, unemployment, the unemployment change, and recent WARN exposure. The earlier WARN-history measure is used only to rule out poor matches; it does not affect the ranking of the remaining controls.

We retain up to four of the closest controls and use all that remain when fewer than four satisfy the restrictions. Controls that more closely resemble the treated county receive greater weight. If d_{sj} denotes the distance between treated event s and control j, the control receives weight

$$w_{sj} = \frac{(1 + d_{sj})^{-1}}{\sum_{\ell \in \mathcal{C}_s} (1 + d_{s\ell})^{-1}},$$

in which $\mathcal{C}_s$ is the set of controls selected for event s. The control weights sum to one within each comparison, and the treated county receives a weight of one. Each treated event and its weighted controls then form one comparison in the stacked event study.

B.3 Treatment characteristics, matched support, and balance

B.3.1 Treated-event characteristics and matched support

Panel A of Table 1 describes the WARN events represented in the matched sample. The median event covers 206 workers over the qualifying three-month window, equivalent to 1.1 percent of the county labor force. The distribution also includes substantially larger local shocks: at the 90th percentile, an event covers 864 workers and 3.4 percent of the local labor force. The median event occurs 10 months before the next election. These statistics show that the treatment captures economically meaningful announcements while spanning considerable variation in both their absolute size and their importance relative to the local labor market.

Figure A3 reports the calendar year of the first qualifying month for each matched event. The first matched events occur in 2000, once the design can establish the required four-year clean history, and the sample extends through 2024. Events occur every year over this period rather than being concentrated in a single election cycle. Their frequency nonetheless varies over time, with increases in 2008–09 and especially in 2020, coinciding with the 2008–2009 recession and the COVID-19 pandemic. Together, however, these periods account for only 29.3 percent of matched events. The design, therefore, draws on WARN events observed across a wide range of national economic conditions and electoral environments.

Panel C of Table 1 summarizes support for constructing the matched comparisons. The matching procedure is applied to 1,703 treated events with the pre-event information needed to impose the matching and caliper restrictions. Of these, 1,508 obtain at least one acceptable control, for a matched share of 88.5 percent. Matched treatment events have, on average, 3.46 controls.

The 1,508 matched events form the event-level sample for the chronological design. Linking these events to the election calendar yields office-specific outcomes at the observed election dates around treatment. Appendix B.4 defines this election ordering and reports the resulting support by office and event-time index.

Overall, the matching procedure retains the large majority of treated events with the required pre-event information. The resulting estimates describe sizable WARN events for which the data contain at least one comparable same-state county.

B.3.2 Balance and common-support diagnostics

Panel D of Table 1 reports the balance in the political, labor-market, and WARN histories used to construct or assess the matched comparisons. The before-matching statistics compare treated events with the broader pool of eligible clean controls, while the after-matching statistics compare treated events with the controls selected for each event using the inverse-distance weights described in Appendix B.2. We report absolute standardized mean differences, which express each treated-control difference in pooled standard-deviation units.

Before matching, the largest differences are in labor-market size, prior Republican voting, changes in unemployment, and prior WARN exposure. Matching substantially reduces each of these differences. The standardized difference in log labor force falls from 0.362 to 0.014, the difference in prior Republican vote share from 0.239 to 0.041, and the difference in the unemployment-rate change from 0.194 to 0.029. The standardized differences in the two prior WARN measures fall from 0.224 and 0.178 to 0.051 and 0.019, respectively. After matching, every displayed standardized difference is 0.051 or smaller. These results show that the selected controls closely resemble treated counties in the observed pre-event characteristics most relevant to the comparison.

Table A4 shows how the matching restrictions narrow the pool of potential controls. Among events matched within the exact state-year-month cell, an average of 62.4 distinct control counties are available before the calipers, 47.6 remain after the prior-WARN restrictions, and 17.6 remain after all calipers; the corresponding medians are 52, 40, and 13. Thus, although the calipers remove many potential controls, the typical exact-cell comparison retains a sizable set of acceptable counties from which to select up to four. Support is thinner among events requiring the fallback: on average, only 1.6 controls remain after all calipers, with a median of 1.

The final weights are generally dispersed across the selected controls in exact-cell comparisons. The median largest control weight is 0.277, close to the 0.25 weight that each county would receive if four controls were weighted equally. The effective number of controls, which equals four when four controls receive equal weight and one when a single control receives all the weight, averages 3.51 for exact-cell matches. One county receives more than half of the control weight in 16.4 percent of these comparisons. Fallback comparisons are more concentrated: their median largest weight is one, and their mean effective number of controls is 1.19. Across all matched events, the effective number of controls averages 3.43, although one county accounts for more than half the weight in 19.0 percent of comparisons. Most matched counterfactuals therefore combine information from several counties, while a minority rely heavily on a single control.

One hundred and ninety-five treated events fall outside the design's common support: no sufficiently comparable control can be found, so they are not included in the matched estimation sample. For 61 events, no potential control satisfies the restrictions on prior WARN history; for the other 134, no potential control satisfies the remaining political and labor-market calipers. These events are therefore excluded because the data do not contain a sufficiently comparable same-state county, rather than because their election outcomes are unavailable.

Table A5 compares the treated events represented in the matched sample with those outside common support. Unmatched events occur in substantially larger labor markets (the mean county's labor force for unmatched events is over three times that for matched events) and involve much larger announcements in absolute terms: they cover 2,958 workers on average, compared with 431 among matched events. Their size relative to the local labor force is much more similar, averaging 1.7 percent for unmatched events and 1.8 percent for matched events. Unmatched events also occur in counties with more prior WARN activity and different pre-event political and labor-market characteristics.

Therefore, the common-support restriction primarily excludes unusually large absolute shocks in larger and more WARN-intensive labor markets, rather than events that are unusually large relative to the local economy. The matched estimates describe the large majority of treated events for which the data contain at least one credible same-state comparison.

B.4 Event-time ordering and the office-specific design

Because presidential, gubernatorial, House, and Senate elections follow different schedules, any pooled design must decide how to place elections from different office cycles on a common event-time scale. Our main design orders elections chronologically. This ordering follows elections in the sequence in which counties face them, treats offices contested on the same date as part of the same electoral moment, and allows the pooled estimates to trace responses across successive election dates without privileging the cycle of a particular office. Because other pooled orderings are also reasonable, Appendix C.1.1 describes and reports the main results under an alternative target-office design.

For the chronological design, we order all included presidential, gubernatorial, House, and Senate election dates within each state. We define $h^C = 0$ as the last included election date before the WARN event and $h^C = 1$ as the first included election date afterward. Earlier and later election dates receive indexes relative to these two elections, and the analysis retains $h^C = -3, \ldots, 4$. A one-unit change in h^C, therefore, represents the next included election date rather than a fixed number of calendar years.

When several offices are contested on the same date, they receive the same chronological index but remain separate county-office outcomes. For example, if presidential, House, and Senate elections occur on the first election date after a WARN event, the event contributes a separate outcome for each office at $h^C = 1$. The sequence is defined from the common state election calendar, so a missing county outcome for one office or year does not change the indexes assigned to later elections. Panel A of Figure A4 illustrates this chronological ordering.

Table A6 reports the number of matched events contributing the four main outcomes for each office and chronological election index. Each cell counts distinct treated events with a usable outcome for that office at that index. The cells are not mutually exclusive: an event can appear in several office rows when multiple offices are contested on the same date and can contribute to several indexes as later elections occur.

Because h^C orders election dates across offices rather than repeated elections for one fixed office, office-specific support need not decline steadily across columns. For example, presidential support is greater at $h^C = -3$ than at $h^C = 0$ because many events have a presidential election three chronological election dates before treatment, even when the immediately preceding election date contains only other offices. The same calendar structure explains why support can rise or fall between adjacent indexes. Nevertheless, all four offices remain well represented throughout the event window. At $h^C = 4$, the sample still includes 526 events with presidential outcomes, 401 with gubernatorial outcomes, 930 with House outcomes, and 600 with Senate outcomes.

To examine whether the pooled response differs across election types, we also reorganize event-time separately within each office. For office o, we define $h_o = 0$ as the last election of that office before the WARN event and $h_o = 1$ as the first election of the same office afterward. Earlier and later values follow the sequence of elections for that office. A House election occurring between two Senate elections therefore changes the chronological index but not the Senate-specific index. Panel B of Figure A4 illustrates this distinction.

The office-specific design changes the ordering of elections rather than the underlying treatment definition. The same universe of qualifying WARN events remains eligible for each office-specific analysis, and no election is discarded merely because another office is contested sooner. Instead, presidential elections are assigned positions within the presidential sequence, gubernatorial elections within the gubernatorial sequence, and likewise for House and Senate elections. The office-specific comparisons are matched separately, so their exact samples can differ when the required election history or outcome is unavailable.

Table A7 reports support across these office-specific indexes. At $h_o = 0$ and $h_o = 1$, the four office-specific samples contain between 1,341 and 1,444 treated comparisons. Support declines at more distant indexes as earlier and later elections fall outside the available sample period, but it remains substantial: at $h_o = 3$, the samples contain 819 presidential, 692 gubernatorial, 1,241 House, and 889 Senate comparisons. The House retains the greatest support because House elections occur every two years, whereas reaching the same office-specific index for the other offices requires a longer span of calendar time.

We estimate a separate event-study path for each office using these office-specific indexes. Presidential elections are therefore compared with earlier and later presidential elections, gubernatorial elections with gubernatorial elections, and likewise for House and Senate races. This construction isolates heterogeneity across election types while preserving the treatment definition, matched-comparison logic, outcomes, controls, and fixed effects used in the main design.

C Main results appendix

C.1 Alternative matched design, party control, residualized WARN exposure, and job destruction

C.1.1 Pooled target-office event-time design

The chronological design orders elections in the sequence in which counties face them. This is a natural way to study the response at successive election dates, but the office observed at $h^C = 0$ need not be the same as the office observed at a later chronological index. The pooled chronological coefficients therefore average treated-control changes across the different pre- and post-WARN office combinations generated by each state's election calendar. Although this design answers the natural question of what happens at the next elections a county faces, it is not the only reasonable way to pool outcomes across offices. We therefore estimate an alternative target-office design in which every pre- and post-WARN comparison is made within the same office before the office-specific comparisons are stacked.

The two designs involve a trade-off. Chronological ordering follows the actual sequence of electoral opportunities and does not privilege the cycle of a particular office, but its pooled path can reflect changes in the mix of offices represented at each index. Target-office ordering holds the office fixed within each comparison, making the pre- and post-WARN elections more directly comparable. Its event indexes, however, span different amounts of calendar time across offices, and the separately matched office samples become thinner at distant horizons. We therefore use chronological ordering as the baseline and the target-office design to assess whether the results depend on the changing office composition of the chronological sequence.

For each treated event and office, we define $h^O = 0$ as the last election of that office before the WARN event and $h^O = 1$ as the first election of the same office afterward. Earlier and later elections are indexed within the same office, and the analysis retains $h^O = -2, \ldots, 3$. Elections for other offices that occur between two target-office elections do not change this index. A one-unit change in h^O, therefore, represents one additional election of the same office rather than the next election date of any office or a fixed number of calendar years.

Figure A5 illustrates the construction. Each row follows one office's election sequence, with $h^O = 0$ marking its last pre-WARN election and $h^O = 1$ its first post-WARN election. We emphasize that the target-office sample is not obtained by simply assigning new indexes to the chronological matched sample. We repeat the matching procedure separately for each office using the treatment definition, matching variables, calipers, number of controls, and inverse-distance weighting rule described in Appendix B.2. Because each office has its own next-election timing, the matching cells and eligible control pools are defined separately across presidential, gubernatorial, House, and Senate comparisons and may differ because the matching election variables vary across offices.

We then stack the separately matched office comparisons at each common target-office index. The pooled coefficient at $h^O = 1$, for example, combines the first post-WARN presidential, gubernatorial, House, and Senate elections from their respective matched comparisons. The outcomes, regression controls, and fixed-effect structure otherwise remain the same as in the chronological design.

Figure A6 reports the target-office event-study estimates, while Table A8 compares the pooled pre- and post-WARN estimates under the two event-time orderings. For each ordering, we re-estimate the baseline specification after replacing the horizon-specific treatment interactions with one Pre interaction and one Post interaction. In the chronological design, Pre combines $h^C = -3, -2, -1$ and Post combines $h^C = 1, \ldots, 4$; in the target-office design, Pre combines $h^O = -2, -1$ and Post combines $h^O = 1, 2, 3$.

The pre-WARN coefficients remain close to zero in the target-office design. The target-office pooled Post estimates are 0.446 percentage points for the Democratic-Republican margin, -0.314 points for Republican voting, 0.132 points for Democratic voting, -0.005 points for other-candidate voting, and -0.188 points for turnout. The corresponding chronological estimates are 0.501, -0.318, 0.182, 0.000, and -0.136 percentage points. Target-office ordering therefore produces a somewhat larger net turnout decline and a smaller Democratic offset, while leaving the Republican vote-rate decline and the absence of an other-candidate response nearly unchanged. Because the pooled coefficients cover $h^C = 1, \ldots, 4$ in the chronological design and $h^O = 1, \ldots, 3$ in the target-office design, these magnitudes are best viewed as design-specific summaries rather than estimates over identical calendar horizons.

The target-office point estimates generally become larger in absolute value at later same-office elections. These indexes cover longer periods of calendar time than the corresponding chronological indexes, especially for presidential, gubernatorial, and Senate elections. Among comparisons contributing the main outcomes, the average elapsed time from the WARN month to $h^C = 1, 2, 3,$ and 4 is 12 months, 34 months (2.8 years), 57 months (4.8 years), and 80 months (6.7 years), respectively, compared with 20 months (1.7 years), 56 months (4.7 years), and 93 months (7.8 years) to $h^O = 1, 2,$ and 3. Support also becomes thinner at later target-office indexes because more recent treatment cohorts do not yet have sufficiently distant same-office elections in the available data. The larger estimates at those indexes should therefore be interpreted alongside their longer calendar horizons and should not be compared mechanically with chronological coefficients carrying the same numerical index.

C.1.2 WARN events and electoral winners

The main results show that sizable WARN events reduce Republican votes per CVAP and increase Democratic votes per CVAP, but they do not establish whether those changes are large enough to alter which party wins an election. We examine that question using three complementary exercises. The county exercise remains closest to the paper's matched design but measures which party leads within the county rather than which party wins the office. The House and state exercises measure actual office winners, but use more selected or aggregated samples and different empirical designs. We therefore view the exercises as complementary evidence on party control, not as estimates of a common parameter.

C.1.2.1 The Democratic county two-party-lead indicator

The county exercise replaces the paper's vote-rate outcomes with an indicator equal to one when Democratic votes exceed Republican votes within the county and zero when Republican votes exceed Democratic votes. We exclude exact ties and require positive votes for both parties. The indicator records the local Democratic-Republican lead; it does not identify the winner of a statewide or district-level office. For House elections, it aggregates the Democratic and Republican votes from the contests represented within the county, as in the main analysis.

The unrestricted specification retains the paper's matched comparisons, controls, fixed effects, weights, and two-way clustering. We estimate the chronological and target-office designs, with the pooled Pre and Post coefficients measured relative to the omitted election at $h = 0$. We also examine elections preceded by a close contest. The relevant geography is the state for presidential, gubernatorial, and Senate elections and the congressional district for House elections. The close-contest samples require the preceding Democratic-Republican margin at that geography to lie within 10 or 5 percentage points. In the target-office design, this is the last same-office election; in the chronological design, it is the lag election at $h = 0$ that supplies the matching covariates. The most restrictive samples also require the county's own preceding margin to lie within the same bandwidth. House observations must map unambiguously to one district, and state-year cells with multiple Senate contests are excluded because the relevant race cannot be identified uniquely.

Table A9 reports the pooled results. Panel A presents the chronological design and Panel B the target-office design. Column (1) retains the paper's matched sample. Columns (2) and (4)-(6) impose the close-race restrictions before applying the paper's matching procedure, while column (3) applies the 10-point contest screen to the existing matches and renormalizes the retained control weights.

Figure A7 reports the unrestricted event-study paths. The pre-event estimates are small and statistically insignificant in both designs. The post-event point estimates are positive at every horizon, but none is statistically significant at the 5 percent level. Pooling these horizons in column (1) of Table A9, the probability that Democrats outpoll Republicans within the county rises by 0.800 percentage points in the chronological design, with a standard error of 0.686, and by 0.716 points in the target-office design, with a standard error of 0.603. The corresponding Pre estimates are -0.234 and -0.008 points.

Columns (2)-(6) provide no stronger evidence in previously close elections. Under the 10-point contest screen, the chronological Post estimate is 1.067 points after rematching and 1.066 after pruning; the target-office estimates are 0.076 and 0.179 points. Narrowing the rematched bandwidth to 5 points raises the chronological estimate to 2.285 points but leaves the target-office estimate at 0.216 points. Requiring both the electoral-jurisdiction contest and the county contest to have been close sharply reduces support. The 10-point Post estimates are based on 103 chronological and 378 target-office treated stacks; at 5 points, only 29 and 82 remain. The estimates from these samples are very imprecise and change sign across designs. None of the pooled Pre or Post estimates in the table is statistically significant.

The unrestricted estimates point in the same direction as the paper's Democratic vote-rate response, but the close-contest estimates do not show a stable change in the county two-party lead. This remains an intermediate answer to the winner question because the outcome measures local partisan balance rather than control of an electoral jurisdiction.

C.1.2.2 Winners of US House elections

The House exercise measures an actual office winner. The binary outcome equals one when a Democrat wins the district and zero when a Republican wins, while the continuous outcome is the Democratic-Republican two-party district margin; this sample requires positive Democratic and Republican votes in both elections. We link each regular House contest to the next regular contest in the same district and retain only election pairs within the same apportionment era and under a stable county-to-district plan. A district interval is treated when at least one qualifying WARN event occurs in a constituent county assigned unambiguously to the district. Because this is a known positive, treatment does not require complete coverage elsewhere in the district and is not invalidated by an overlapping notice in a split county. A potential control must instead have complete WARN coverage for every intersecting county and no qualifying event in either a wholly assigned or split county. Ambiguous intervals are excluded rather than classified as untreated.

We impose these restrictions before rematching eligible district intervals using the paper's matching variables, calipers, and inverse-distance weights. Each treated episode is matched to at most four distinct control districts; if several eligible counties represent the same control district, we retain the closest. Within each comparison, control weights are renormalized to sum to one, while repeated treated county episodes mapping to the same district-election pair jointly receive unit weight. Of 1,703 match-ready treated county episodes, 1,112 satisfy the election and geographic restrictions. The final winner sample contains 404 comparison stacks representing 298 distinct treated district-election pairs; the two-party-margin sample contains 309 comparison stacks. We estimate

$$Y_{dgh} = \beta \left(D_{dg} \times \mathrm{Post}_h \right) + \alpha_{dg} + \lambda_h + \kappa_{sy} + \varepsilon_{dgh}, \tag{C1}$$

where d indexes districts, g comparison stacks, and $h \in \{0, 1\}$ the elections immediately before and after the event. Here, D_{dg} indicates the treated district unit and Post_h equals one for the first post-event election. The specification includes district-by-comparison, election-index, and state-by-election-year fixed effects. We report estimates with and without the paper's four contemporaneous county controls. Standard errors are two-way clustered by comparison stack and congressional-district map era.

Panel A of Table A10 reports the district outcomes. The Democratic-victory estimates are -1.488 percentage points without the four controls and 0.149 points with them, with standard errors of 2.689 and 2.539. The corresponding district-margin estimates are -1.946 and -1.799 points, with standard errors of 1.085 and 1.127. Panel B separates Democratic gains in previously Republican districts from retention in previously Democratic districts. The gain estimates are -2.722 and -1.701 points, with standard errors of 1.959 and 1.718. The retention estimates are also negative but are based on only 26 distinct treated district-election pairs and have standard errors above 13 points.

These estimates provide no evidence that WARN events change House control, but they are not directly comparable to the county results. The exercise uses only the first post-event election and a geographically selected sample. More fundamentally, the event and matching variables are measured for the event county, while the outcome is measured for the entire district. County-level balance therefore does not ensure balance in district electoral history or economic conditions.

C.1.2.3 Winners of statewide elections

The state exercise examines presidential, gubernatorial, and US Senate elections. The binary outcome equals one when a formally Democratic candidate wins the ending election and zero otherwise; the continuous outcome is the statewide Democratic-Republican two-party margin, observed only when the returns are complete and both parties receive positive votes. The same-office design links consecutive regular elections within each presidential, gubernatorial, or Senate-class series. The chronological design instead orders all such election dates within a state and measures exposure between consecutive distinct dates, regardless of office. When several offices are contested on the ending date, each contributes an outcome.

For each interval, we aggregate reported WARN-affected workers over months whose assigned date--the fifteenth of the month, as elsewhere in the paper--falls strictly between the two elections. The primary exposure is the annualized number of affected workers per 1,000 members of the state's average labor force. We annualize by dividing the interval total by the number of included months and multiplying by 12 before scaling by the average labor force. Alternative measures use annualized WARN notices per 100,000 labor-force participants or the inverse hyperbolic sine of the primary worker rate. These measures include notices that can be assigned to the state even when they lack a county assignment. We retain only intervals with documented WARN coverage and complete labor-force data throughout the exposure window, so zero exposure does not reflect missing coverage.

The pooled specifications absorb office-series and office-by-ending-election-year fixed effects and cluster standard errors by state. The regressions are unweighted, so each ending-election outcome receives equal weight. The lagged-margin specifications additionally control for the Democratic-Republican margin in the preceding election of the same office series, allowing its coefficient to differ by office. The pooled models estimate one common WARN-exposure coefficient across offices. Panels A and B of Table A11 report the same-office and chronological models, respectively, for the primary and alternative exposures. Panel C reports the lagged-margin same-office specification for the primary exposure separately by office.

For one additional annualized WARN-affected worker per 1,000 labor-force participants, the same-office lagged-margin estimates are a 0.485-percentage-point increase in the probability of a Democratic victory and a 0.332-point decline in the Democratic margin, with standard errors of 1.600 and 0.678. The corresponding chronological estimates are a 0.528-point increase in the Democratic-win probability and a 0.058-point decline in the margin, with standard errors of 0.815 and 0.346. Omitting the lagged margin or using either alternative exposure measure yields the same broad conclusion: statewide WARN exposure is not consistently associated with either outcome.

The office-specific Democratic-win coefficients in Panel C range from -0.989 points for gubernatorial elections to 3.905 points for Senate elections; none is statistically significant. Among the office-specific margin estimates, only the presidential estimate is conventionally significant, at -0.526 points with a standard error of 0.235. Because it appears among multiple outcomes and specifications and is not accompanied by a change in the presidential winner, we treat it as exploratory rather than evidence of a systematic effect.

Unlike the county and House exercises, the state analysis is not a matched event study of discrete WARN events. It estimates a state-panel association between cumulative exposure and the outcome of the ending election. Statewide aggregation aligns the outcome with the jurisdiction awarding the office, but combines many notices over intervals of different lengths and leaves only 34 or 35 state clusters in the pooled models. The estimates may also reflect state-level changes correlated with WARN exposure. They are therefore complementary evidence on statewide party control, not an alternative estimate of the paper's county-level causal effect.

C.1.2.4 Interpretation

Across the three exercises, we find no consistent evidence that WARN exposure changes the local party lead or the party winning an office. The unrestricted county estimates are positive but imprecise, and the close-contest estimates are not stable across designs. The House and state analyses, which measure actual office winners, likewise yield no consistent increase in Democratic victories. These results do not establish that the effects are exactly zero, particularly in the smaller House and close-county samples.

This conclusion does not conflict with the main results. A vote shift changes the winner only when the initial margin is sufficiently small, and an effect within one county may be diluted when votes are aggregated over a district or state. The actual-winner exercises also use different samples and counterfactuals from the county-level matched event study. The evidence therefore supports a narrower conclusion: WARN events change the partisan composition of votes in affected counties, but the available data do not show that these changes systematically alter party control.

C.1.3 Political accountability

The turnout and party-specific vote-rate results do not reveal whether WARN events also reduce support for the candidates or parties in power at the time of the announcement. We examine this possibility using two accountability outcomes: the incumbent candidate's vote share and the pre-election vote share of the office-holder's party. Unlike the candidate outcome, the office-holder-party outcome remains defined in open-seat elections because it follows the party controlling the office before the election.

For presidential, Senate, and House elections, we identify incumbent candidates and their parties using Federal Election Commission candidate records. For gubernatorial elections, we use a manually assembled set of candidate and incumbent records. We link these records to the election calendar, which identifies each election by office and year and, where relevant, Senate class or House district. Exact election identifiers provide the primary link, but, when necessary, we use state-district-year links for House races and state-year links for gubernatorial elections.[2]

Once an incumbent candidate is identified, we map the candidate's party to the corresponding Democratic or Republican vote total in the county returns. We restrict this outcome to Democratic and Republican incumbents because some election files combine independent and minor-party candidates into a single "other-candidate" category and therefore do not allow us to isolate the votes received by an individual nonmajor-party incumbent. We assign the pre-election office-holder party separately. For presidential elections, we use the sitting president's party. For Senate and gubernatorial elections, we use the party that won the previous election for the same Senate class or gubernatorial office. For these two offices, when the previous election falls outside the period covered

[2] We use state-year links for gubernatorial elections only when the contest is unique, and the assignment is unambiguous.

by our returns data, we use the party of the major-party incumbent as the pre-election office-holder party when an incumbent can be identified.

House elections require additional restrictions because the county returns aggregate votes across all House contests represented within a county. We construct the House accountability outcomes only when every regular House contest in the state-year has a resolved Democratic or Republican incumbent, and all incumbents belong to the same party. This condition excludes state-years with an open seat, an unresolved or non-major-party incumbent, or incumbents from both parties. When the condition is satisfied, votes for the common incumbent party can be identified unambiguously in the aggregate county returns; otherwise, both House accountability outcomes remain undefined. Under this construction, the incumbent-candidate and pre-election office-holder-party outcomes coincide for House elections.

Table A3 reports the support of our incumbent-candidate and office-holder-party data. We observe incumbent candidates in 50.0 percent of presidential observations, 58.1 percent of gubernatorial observations, 86.5 percent of Senate observations, and 14.4 percent of House observations. The office-holder party is observed for all presidential observations and for 94.1 and 96.8 percent of gubernatorial and Senate observations, respectively. For the House, only 14.4 percent of observations include information on the office-holder party because of the aggregation restrictions described above.

We estimate the accountability effects using the same chronological matched event-study specification as for the main outcomes, restricting each regression to county-office-election observations for which the relevant accountability outcome is available. Figure A8 reports the resulting paths for the incumbent-candidate and pre-election office-holder-party vote shares.

The post-WARN incumbent-candidate estimates are negative at all four post-WARN periods, ranging from -0.28 to -0.11 percentage points. In a specification that replaces the horizon-specific post-WARN interactions with a single post-WARN interaction, the incumbent-candidate share falls by 0.18 percentage points, with a standard error of 0.16. The estimates, therefore, point to some decline in incumbent-candidate support, but they are imprecise. The office-holder-party results are weaker: both the horizon-specific estimates and the estimate pooling all post-WARN periods are close to zero and statistically insignificant.[3]

Taken together, the accountability results provide suggestive evidence that incumbent candidates may lose support after a sizable WARN event, but no evidence of a broader response against the party controlling the office.

C.1.4 Residualized WARN exposure

The main text asks whether WARN events contain electorally relevant information beyond that summarized by standard labor-market indicators. We examine this by reconstructing treatment from the portion of WARN exposure unexplained by observed pre-event economic conditions. If the electoral response is driven by the component of WARN exposure predictable from these conditions, it should weaken when treatment is instead defined from residual exposure.

[3]The target-office specification leads to similar conclusions: the pooled post-WARN incumbent-candidate estimate is -0.25 percentage points, with a standard error of 0.17, while the office-holder-party estimate is virtually zero.

We begin with the three-month WARN exposure measure defined in Equation (2). At the county-month level, we regress this measure on the levels, 12-month changes, and annualized five-year changes in log labor force, the employment-to-population ratio, the unemployment rate, per capita income, manufacturing, goods-producing, and private-employment shares, average weekly wages, average establishment size, and log population. The regression also includes state, calendar-year, and month-of-year fixed effects. Monthly predictors are measured through the month immediately preceding the three-month WARN window, while annual predictors use the latest completed pre-window year.

Residualized WARN exposure equals observed exposure minus the full fitted value from this regression, including the fixed effects. The included variables explain 1.2 percent of the variation in WARN exposure. Standard labor-market conditions therefore predict little of the variation in WARN exposure; the exercise removes the component they do predict.

We reconstruct the matched design using residualized exposure in place of observed exposure. A county-month qualifies when residual exposure is at or above the 85th percentile among positive residuals, a cutoff of 0.22 percentage points of the labor force, and implies at least 50 affected-worker equivalents.[4] Residual exposure defines the first qualifying month, clean-treatment histories, control eligibility, and the prior-exposure variables used in matching. All other matching and estimation choices follow the baseline design. Treatment therefore represents WARN exposure that is unusually high relative to observed local labor-market conditions. Because the residualized design can be constructed only where all first-stage predictors are observed, we also report an actual-WARN design rebuilt on exactly the same county-month support. For this series, we recompute the 85th-percentile cutoff among positive observed exposure on that support, retain the baseline floor, and reconstruct treatment timing, clean histories, control eligibility, prior-exposure matching variables, and matches. This comparison holds first-stage support fixed and shows that the partisan paths remain close to the full baseline, so the support restriction does not account for the residualized-WARN results.

The residualized procedure produces 1,604 match-ready events, of which 1,389, or 86.6 percent, obtain at least one acceptable control. Matched events receive 3.36 controls on average, and 98.1 percent are matched in the exact state-year-month cell. Although the continuous observed and residualized measures are highly correlated, reconstructing treatment produces related but distinct samples: 800 of the 1,604 residualized events, or 49.9 percent, also qualify as baseline WARN treatments in the same month. Of the remaining 50.1 percent, around 10 percent qualify in the same county and election-year window but in a different month, around 50 percent occur in baseline-treated counties but in different election windows, and the remaining approximately 40 percent occur in counties not treated under the baseline definition.

Figure A9 compares the full baseline, the actual-WARN design rebuilt on the same sample used to build the residualized measure, and the residualized-WARN event-study paths, while Row 2 of Table 2 reports the pooled Pre and Post estimates. Pooling the four

[4]This count equals residual exposure, measured in percentage points, multiplied by the county labor force and divided by 100. It allows us to retain the baseline absolute-size floor but is not an observed worker count.

post-event indexes, residualized WARN events increase the Democratic-Republican vote-rate margin by 0.73 percentage points, reduce Republican votes per CVAP by 0.45 points, increase Democratic votes per CVAP by 0.28 points, leave other-candidate voting close to zero, and reduce turnout by 0.17 points. The corresponding baseline estimates are 0.50, -0.32, 0.18, approximately zero, and -0.14 percentage points. The pooled Republican pre-WARN estimate is small and marginally significant, but the index-specific estimates do not increase as treatment approaches; the remaining pre-WARN estimates are statistically insignificant.

The similar post-event paths show that the electoral response survives after removing the component of WARN exposure predicted by standard labor-market indicators. WARN events therefore contain information relevant to subsequent electoral outcomes that these measures do not summarize. Residualization does not isolate a pure announcement effect, however, because residual exposure may also contain omitted local shocks, measurement error, or subsequently realized job destruction that the included conditions fail to predict. Appendix C.1.5 examines the last possibility.

C.1.5 WARN events and realized job destruction

Because WARN notices generally precede covered layoffs, the electoral response may reflect the announcement itself, the job destruction that follows, or both. The preceding residualized-exposure exercise shows that standard pre-event labor-market conditions do not account for the response, but it cannot distinguish between these channels. We therefore use county-level firm job-destruction rates (JDR) from the Quarterly Workforce Indicators, described in Appendix A.3, to examine whether WARN events are followed by realized job destruction and whether that destruction, rather than the announcement itself, accounts for the partisan response.

The JDR measures jobs lost at contracting private-sector firms as a share of average county employment. For each treated county and its matched controls, we define the pre-WARN JDR as the county's average during the two complete quarters ending before the qualifying WARN window begins and the subsequent JDR as its average during the WARN-trigger quarter and the following quarter. Their difference defines the change in the JDR. Controls use the same calendar quarters as their matched treated county. We require the following quarter to end before the first post-WARN election, a restriction satisfied by 84.3 percent of the baseline matched comparisons. Requiring complete QWI data for every county in a matched comparison leaves 1,255 comparisons, for which we retain the original controls and weights.

Because the trigger quarter may include months before the WARN trigger, we also measure subsequent job destruction using the first two complete quarters after that quarter. This stricter definition provides clearer temporal separation between the WARN event and measured destruction but reduces the sample to 955 comparisons.

Panel A of Table A12 shows that WARN events are followed by substantial realized job destruction. Under the primary definition, treated counties experience a 0.84-percentage-point larger increase in the JDR than their matched controls, with a standard error of 0.07. The corresponding complete-quarter estimate is 0.69 percentage points, also with a standard error of 0.07.

We next reestimate the baseline specification on each QWI sample, first without controlling for the JDR and then allowing the electoral path to vary with the change in the JDR. All other elements of the specification remain unchanged. Figure A10 compares the primary same-sample baseline and JDR-adjusted event-study paths, while Panel B of Table A12 reports the pooled estimates. On the primary sample, the post-WARN estimates for the Democratic-Republican vote-rate margin, Republican votes per CVAP, and Democratic votes per CVAP are 0.53, -0.32, and 0.21 percentage points. Accounting for the change in the JDR moves them to 0.46, -0.28, and 0.18 points, respectively. Thus, each partisan estimate declines by approximately 12 percent but remains statistically significant. Other-candidate voting remains close to zero, and all pooled pre-WARN estimates remain statistically insignificant. The stricter complete-quarter estimates reported in Panel B yield the same conclusion.

Panel C of Table A12 reports how the association between changes in the JDR and electoral outcomes differs between the pooled Post and Pre periods. Under the primary definition, the association of a 1-percentage-point larger increase in the JDR is 0.05 percentage points more negative in the Post period than in the Pre period for Republican voting and turnout, with no discernible shift for Democratic voting. The corresponding 0.06-point Post-minus-Pre increase in the association with the Democratic-Republican margin is imprecisely estimated. Under the stricter definition, the party-specific shifts are statistically insignificant, whereas the turnout shift remains similar and statistically significant at -0.06 percentage points. We therefore view lower turnout as the only consistent Post-minus-Pre shift associated with realized job destruction; its partisan implications are less clear.

Overall, WARN events are followed by substantial realized countywide job destruction, but measured destruction accounts for little of the subsequent partisan response. Because the JDR is measured after treatment and covers all private-sector firms in the county, the exercise does not directly identify a causal announcement effect. Nevertheless, the limited attenuation leaves a substantial partisan response beyond observed job destruction, consistent with an announcement component of the WARN effects.

C.2 Timing, industry heterogeneity, and geographic scope

C.2.1 Treatment-to-election timing

The chronological design orders elections according to when they occur around a WARN event, but the interval between the event and the first post-WARN election varies across counties. To examine this dimension, we divide the matched sample into two groups. We set $\tau^C = 0$ when the first post-WARN election occurs in the treatment election year, so the event occurs 1-12 months before that election. We set $\tau^C = 1$ when the first post-WARN election occurs in the following election year, placing the event 13-24 months before it. In both cases, the relevant election is the first subsequent presidential, gubernatorial, House, or Senate election in the state's chronological election calendar. This division leaves the event-time index unchanged: within either group, $h^C = 1$ remains the first election after treatment, $h^C = 2$ the next election, and so forth.

Figure A11 reports the event-study paths separately for the two groups, while Panel B of Table 2 reports the corresponding pooled Pre and Post estimates. The paths follow the same broad pattern, although the declines in Republican votes per CVAP and turnout are somewhat larger when the WARN event occurs closer to the next election. In the pooled Post specifications, the Democratic-Republican vote-rate margin rises by 0.530 percentage points for $\tau^C = 0$ and 0.446 points for $\tau^C = 1$. Republican votes per CVAP fall by 0.336

percentage points for $\tau^C = 0$, compared with 0.277 points for $\tau^C = 1$. Democratic votes per CVAP rise by 0.194 and 0.169 percentage points in the two groups, and turnout falls by 0.145 and 0.104 percentage points, respectively. Other-candidate votes per CVAP remain close to zero in both groups.

These differences should be interpreted cautiously because dividing the sample by treatment-to-election timing reduces support, especially at later horizons. The turnout specifications contain 902 and 605 treated comparisons at $h^C = 1$ for $\tau^C = 0$ and $\tau^C = 1$, respectively; by $h^C = 4$, support falls to 569 and 458 comparisons. The later-horizon estimates are therefore less precise, and the differences between the two paths do not identify a sharp timing threshold. Overall, the timing estimates point in the same direction across both groups. The declines in Republican votes per CVAP and turnout are somewhat larger when WARN events occur closer to the next election, but the more robust conclusion is that the response is qualitatively similar whether the event occurs within one year or one to two years before the election.

C.2.2 Nonmanufacturing results

This exercise asks whether the electoral response to sizable WARN events is confined to manufacturing-related announcements. We classify notices as manufacturing or nonmanufacturing using state-reported industry codes when available and manual or AI-assisted assignments based on the employer and notice information otherwise. Manual verification of more than 300 notices shows agreement above 94 percent between the human and AI classifications. More than 99 percent of notices can be classified as manufacturing or nonmanufacturing. Nonmanufacturing accounts for about 70 percent of both classified notices and affected workers.

We reconstruct treatment and the matched comparisons using nonmanufacturing notices, rather than simply restricting the baseline matched sample. Specifically, we replace total WARN workers in Equation (2) with workers covered by nonmanufacturing notices and recompute the 85th-percentile cutoff among positive three-month nonmanufacturing exposure windows. This cutoff equals 0.39 percent of the county labor force, compared with 0.63 percent in the all-sector design. Control eligibility is likewise based on nonmanufacturing exposure: a control must have complete coverage and no qualifying nonmanufacturing event in the treatment window, the four preceding election-year windows, or the observed portion of the four subsequent election-year windows. All other dimensions of the matched empirical design remain unchanged from the main analysis.

Of 1,448 match-ready nonmanufacturing WARN treatment events, 1,246 (86 percent) can be matched to at least one control. Conditional on matching, events receive about 3.5 controls on average, and covariate balance is strong: standardized differences for all matching variables are 0.08 or less. We then estimate the same chronological event-study specification as in the main analysis.

Figure A12 shows little evidence of differential pre-event trends. Pooling $h = 1, \ldots, 4$, the post-WARN estimates are -0.25 percentage points for Republican votes per CVAP, 0.23 points for Democratic votes per CVAP, and -0.02 points for turnout, with the last estimate statistically insignificant. The corresponding baseline estimates are -0.32, 0.18, and -0.14, respectively. The party-specific response is therefore close to the baseline pattern, but the aggregate turnout effect is more attenuated because the increase in

Democratic votes offsets more of the decline in Republican votes.

Because manufacturing and nonmanufacturing layoffs may occur in the same county and period, we examine whether the sector-specific treatment could still capture manufacturing WARN activity. The overlap is limited. Among county-months with a nonmanufacturing notice, 82.2 percent contain no manufacturing notice, and the correlation between manufacturing and nonmanufacturing WARN-exposure rates, each measured relative to county labor force, is only 0.007. Among matched nonmanufacturing treatment events, just 10.4 percent contain a manufacturing notice in the same three-month treatment window; this share falls to 7.9 percent after excluding 2020 notices, and the median number of manufacturing workers in that window is zero. These patterns indicate that the nonmanufacturing treatment is not simply a proxy for manufacturing WARN activity.

Taken together, the exercise extends the main result beyond the manufacturing decline and trade-exposure settings emphasized in the related literature. Mass-layoff announcements outside manufacturing generate a shift in party-specific vote rates similar to that in the full sample, although the net turnout effect is more attenuated.

C.2.3 Geographic scope and spillovers

Our geographic-scope exercise tests for electoral spillovers from WARN events into neighboring counties. We do so by comparing the baseline design, which allows counties in the same local labor market to serve as controls, with a specification that excludes them from the potential control set. Because the main matched design can select controls from the treated county's local labor market, spillovers to those counties could attenuate the estimated treated-control differences. We implement the alternative specification as a donut design that excludes all potential controls in the treated county's state-by-commuting-zone unit.

After imposing this restriction, we repeat the baseline matching procedure. The treatment definition, matching cells, covariates, calipers, maximum number of controls, and inverse-distance weighting rule remain unchanged. We then estimate the same chronological specification as in the main design. The geographic restriction leaves 1,490 treated events with at least one acceptable control, compared with 1,508 in the baseline design, and provides the same average of 3.46 controls per matched event as the baseline.

Figure A13 compares the baseline and donut event-study paths, while Panel D of Table 2 reports the pooled Pre and Post estimates. The two sets of estimates are similar throughout the event window. In the donut design, the pooled post-WARN effects are 0.489 percentage points for the Democratic-Republican margin, -0.302 points for Republican voting, 0.187 points for Democratic voting, and -0.119 points for turnout, compared with 0.501, -0.318, 0.182, and -0.136 in the baseline design. Other-candidate voting remains close to zero in both designs.

Excluding nearby controls, therefore, does not materially change the vote-rate pattern. The exercise provides no evidence of electoral spillovers from WARN events into neighboring counties, consistent with a geographically localized political response to WARN shocks.

C.3 Pre-existing local deterioration and residualized WARN exposure

The central identification concern is that WARN events may occur in counties already on worsening economic or political trajectories, causing treated and matched-control outcomes to diverge even without the announcement. We address this concern in two ways. First, we reweight or rematch counties using longer pre-WARN economic, demographic, and sectoral histories. Second, we summarize pre-event conditions in a forecasted-deterioration score and allow counties with different predicted trajectories to follow different event-time paths.

C.3.1 Observable controls and long-run trajectories

We first vary the time-varying county controls while leaving the matched sample, weights, and fixed effects unchanged. The no-control specification omits these controls. The baseline specification, the main one used throughout the paper, includes unemployment, labor-force size, per capita income, and cumulative COVID cases per 1,000 residents. The richer specification additionally controls for the manufacturing employment share, log average weekly wages, and log population density. Comparing these specifications is useful because our goal is to understand the effects of a WARN shock net of labor-market conditions or other confounders. Therefore, large differences in estimates across specifications would suggest that the treatment is merely acting as a proxy for broader deterioration.

Figure A14 reports the resulting event-study paths. The pre-WARN estimates are similar across the three specifications. The party-specific movements are almost identical across specifications: Republican votes per CVAP fall similarly under all three control sets, while Democratic votes per CVAP rise. The turnout decline is slightly larger without controls and slightly more attenuated with the richer controls, particularly at later election indexes. The choice of controls, therefore, has little effect on our estimates and strengthens our interpretation that our estimated effects are not merely a proxy for the broader deterioration in labor-market conditions.

These specifications may still miss slower-moving differences between treated and control counties. We therefore construct longer pre-WARN histories using three-year changes in unemployment, labor force, per capita income, manufacturing and goods-producing employment shares, wages, establishment size, population, voting-age population, and CVAP. We also use five-year changes in unemployment, labor force, and the manufacturing and goods-producing employment shares.

We incorporate these histories in two ways. The trajectory-reweighted specification retains the original matched controls but downweights controls whose longer-run histories are less similar to that of the treated county. More specifically, we standardize the additional measures, calculate the distance between each treated county and its selected controls across the available trajectories, and then renormalize the adjusted control weights within each matched comparison. The trajectory-rematched specification instead adds the longer histories with adequate common support to the original matching distance and rebuilds the control group and its weights. Reweighting, therefore, asks whether residual differences among the original controls matter, while rematching asks whether accounting for longer-run decline from the outset selects a different comparison group.

Figure A15 compares the preferred estimates with the trajectory-reweighted and trajectory-rematched paths. Reweighting produces estimates that are nearly indistinguishable from the preferred specification. Rematching changes some point estimates and yields wider confidence intervals, but it preserves the main pattern: the pre-WARN estimates remain centered near zero, while Republican votes per CVAP fall, Democratic votes per CVAP rise, and turnout falls after the event. Observable differences in longer-run economic, demographic, and sectoral trajectories therefore do not account for the post-WARN response.

C.3.2 Forecast score

The trajectory exercises above assess individual dimensions of pre-WARN decline, but treated counties may still differ from controls in their combined conditions. We summarize this possibility with a forecasted-deterioration score constructed solely from pre-WARN observations.

The score combines the unemployment rate in the year before treatment with changes from three years before treatment to the year before treatment in unemployment, labor force, per capita income, manufacturing and goods-producing employment shares, population, and CVAP. We standardize each measure and orient it so that higher or rising unemployment, slower labor-force and income growth, declining manufacturing and goods-producing employment shares, and slower population or CVAP growth increase the score. Each available component receives equal weight, and we standardize the resulting index, so one unit represents one standard deviation. A higher score, therefore, identifies counties that appeared more likely to deteriorate before treatment. The score is an equal-weighted summary of pre-WARN conditions rather than a prediction estimated from post-WARN outcomes.

Treated counties have a somewhat higher average forecast score than their matched controls, approximately 0.07 compared with -0.04, for a standardized difference of 0.11.[5] To examine whether this residual difference accounts for the results, we re-estimate the main specification, retaining the baseline matched sample, weights, controls, and fixed effects, and add interactions between the forecast score and each election index. This allows counties with different pre-WARN scores to follow different event-time paths even in the absence of treatment.

Figure A16 compares the preferred and forecast-adjusted event-study paths, while Row 11 of Table 2 reports the pooled Pre and Post estimates. The two paths are nearly indistinguishable. Allowing counties with different observable risks of deterioration to follow different event-time paths, therefore, does not attenuate the turnout or party-specific vote-rate responses.

Finally, although its primary purpose is to test whether WARN events contain information beyond standard labor-market indicators, the residualized-exposure exercise in Appendix C.1.4 also bears on this identification concern. Its similar estimates show that the electoral response persists after removing the component of WARN exposure predicted by observed pre-event conditions, reinforcing the evidence above that observable pre-existing deterioration does not account for the main results.

C.4 Additional robustness of party-specific vote-rate and turnout results

This section provides additional details for the robustness exercises reported in Panel F of Table 2. We test whether the results are sensitive to the exposure-aggregation period, the treatment percentile threshold, the minimum-worker floor, the empirical specification and matched design, and the exclusion of the 2020 treatment cohort.

[5] The matched-control mean uses the baseline inverse-distance weights.

C.4.1 Alternative treatment definitions

The baseline treatment definition depends on the threshold used to identify a sizable WARN event, the period over which nearby notices are aggregated, and the minimum number of affected workers. We first replace the baseline 85th-percentile threshold with the 80th or 90th percentile. We then replace the baseline three-month exposure window with a one- or six-month window. Finally, we vary the absolute-size requirement for treatment, replacing the baseline 50-worker floor with floors of 25 or 75 affected workers. Each change can alter not only which county-months are treated, but also which counties satisfy the clean-window restriction and are eligible controls. We therefore reconstruct treatment and control eligibility and repeat the matching procedure for each specification. Everything else in the design follows the baseline matched design, including the use of an empirical specification identical to that shown in the main text.

Figure A17 shows that the dynamic paths remain similar across these definitions. Across the six alternatives, the pooled post-WARN estimates range from 0.41 to 0.60 percentage points for the Democratic-Republican vote-rate margin, from -0.37 to -0.26 points for Republican votes per CVAP, from 0.15 to 0.25 points for Democratic votes per CVAP, remain close to zero for other-candidate votes per CVAP, and range from -0.13 to -0.08 points for turnout. Thus, moderate changes in the percentile cutoff, minimum-worker floor, or exposure window used to define sizable WARN events do not materially alter the results.

C.4.2 Varying the number of control counties

The baseline design retains up to four acceptable control counties for each treated event. We use four weighted controls because this number provides the lowest standardized differences between treatment and control observations across matching variables. Emphasizing the closest matches improves comparability, but it may make the estimated counterfactual sensitive to a small set of selected counties. To assess this, we repeat the matching procedure while retaining up to ten controls, using the same eligibility restrictions, calipers, distance measure, and inverse-distance weighting rule. This broadens the control set without admitting counties that fail the baseline match-quality restrictions. We then estimate the effects using an empirical specification identical to that in the main text.

Figure A18 shows that allowing additional control counties in each matched comparison leaves the dynamic patterns largely unchanged. The pooled post-WARN estimates are -0.28 percentage points for Republican votes per CVAP, 0.18 points for Democratic votes per CVAP, and -0.10 points for turnout, compared with -0.32, 0.18, and -0.14 in the baseline. Thus, the main pattern does not depend on a counterfactual constructed from only the four closest control counties.

C.4.3 Varying the estimator: alternative fixed effects and weighting

We examine two additional specification choices. First, we replace the baseline event-index and state-by-office-by-year fixed effects with comparison-by-election-index-by-office fixed effects. These fixed effects absorb shocks shared by the treated county and its controls within each matched comparison, office, and election index, so the effects are identified entirely from within-comparison treated-control differences. This guards against comparison-specific election shocks that are not absorbed by the baseline state-by-office-by-year fixed effects. For example, a regional campaign or media-market mobilization may raise turnout in the counties belonging to one matched comparison but not in other counties in the same state and year. The baseline specification may not fully absorb this common movement, whereas the comparison-specific fixed effect does. This makes the specification more demanding because it discards variation across matched comparisons and relies only on the treated-control contrast within each comparison.

Second, we weight matched comparisons by the natural logarithm of the treated county's CVAP at the last pre-WARN election. The baseline assigns the same total weight to every treated event, regardless of the size of its electorate. This exercise tests whether the results are concentrated in smaller counties or change when events affecting larger electorates receive greater influence. Because these robustness exercises focus on changes to the empirical specification, the matched comparison groups remain the same as in the main design.

Figure A19 reports the estimates of the more granular fixed-effect specification and the specification using weights based on the log of citizen-voting-age population. Both produce estimates close to the baseline. The comparison-index specification yields pooled post-WARN effects of -0.31, 0.17, and -0.14 percentage points for Republican votes per CVAP, Democratic votes per CVAP, and turnout, respectively; the corresponding log-CVAP-weighted estimates are -0.33, 0.18, and -0.15. The results of the fixed effects exercise show that the estimates do not rely on pooling election-time shocks across matched comparisons. The log-CVAP-weighted estimates show that the vote-rate results are not driven by smaller counties alone.

C.4.4 Excluding 2020 WARN events

WARN activity increased sharply during the COVID-19 pandemic, and the 2020 election occurred under unusually disrupted economic and electoral conditions. Consequently, the full-sample estimates could be disproportionately influenced by an unusually large treatment cohort whose WARN events coincide with pandemic-specific economic and electoral shocks, despite the inclusion of COVID controls and state-by-office-by-year fixed effects. We therefore exclude matched comparisons whose treatment election year is 2020 and re-estimate the baseline specification for the remaining comparisons. This restriction removes the 2020 treatment cohort rather than all election outcomes observed in 2020. We do not repeat the matching procedure because the purpose of the exercise is to isolate the influence of this cohort while holding the counterfactuals for all remaining treated events fixed. We instead remove the corresponding comparison groups and retain the original matches for all other cohorts. Everything else follows the baseline design.

Figure A20 shows that excluding this cohort of treatments leaves the main patterns intact. The pooled post-WARN estimates are -0.32 percentage points for Republican votes per CVAP, 0.18 points for Democratic votes per CVAP, and -0.11 points for turnout. The virtually unchanged Republican and Democratic estimates show that the partisan participation response is not a feature of the exceptional 2020 cohort; the modest attenuation of aggregate turnout does not alter that conclusion.

C.4.5 Functional form and fitted-value bounds

The party-specific vote rates and turnout are bounded between zero and 100, whereas Equation (3) is linear.[6] We first assess whether this functional form produces inadmissible predictions using fitted values from the baseline event study, including the absorbed fixed effects. Negative predictions account for 0.1 percent of Republican-vote-rate observations (min = -12.2), 0.9 percent of Democratic-vote-rate observations (min = -9.8), and 7.8 percent of other-candidate-vote-rate observations (min = -2.4); no turnout prediction is negative (min = 2.4), and no fitted value exceeds 100. Out-of-range predictions are therefore very rare for the Republican and Democratic outcomes and absent for turnout.

The predicted-value examination limits concerns about inadmissible predictions but does not address the constant percentage-point effect across outcome levels imposed by the linear specification. Therefore, we reestimate the baseline chronological specification after replacing each included outcome $V/CVAP$ with its log-odds transformation $\ln[V/(CVAP - V)]$, keeping everything else in the design unchanged.[7] This allows the implied percentage-point effect to vary with the fitted outcome level and shrink near the bounds. We report weighted average derivatives on the original percentage-point scale. Because the transformation is undefined at the boundaries, we restrict the samples by outcome ($0 < V < CVAP$). This restriction excludes 0.6, 3.9, and less than 0.1 percent of the observations for Republican vote rates, Democratic vote rates, and turnout, respectively, without eliminating any treated stack from the baseline estimation sample.[8]

Figure A19 reports the corresponding estimates for the three transformed outcomes; Panel A contains no log-odds series because the signed margin is not transformed. Pooling ($h^C = 1, \ldots, 4$), the average marginal effects are -0.23 percentage points for Republican votes per CVAP (p = 0.03), 0.14 points for Democratic votes per CVAP (p = 0.06), and -0.07 points for turnout (p = 0.46). The corresponding linear estimates are -0.32, 0.18, and -0.14 percentage points. All pooled pre-WARN estimates from the transformed specification are close to zero and statistically insignificant. Thus, the log-odds specification preserves the opposite-signed party responses and a slightly closer to zero negative turnout estimate.

[6]The Democratic-Republican vote-rate margin is instead bounded between -100 and 100, and all of its fitted values lie within that interval.

[7]Because the Democratic-Republican margin is signed, it does not admit the same ordinary log-odds transformation.

[8]We do not apply the transformation to other-candidate voting because 14.3 percent of those observations equal zero and the resulting interior restriction would materially change the sample.

D Mechanisms and downstream policy outcomes

D.1 Population size, migration, and demographic composition

The county-level vote-rate responses could partly reflect changes in the size or composition of the resident population rather than changes in voting behavior among continuing residents. Migration may also reduce turnout by creating voter-registration frictions. We assess these possibilities by adding controls for migration and demographic composition.

Using SEER and Census population estimates, we construct the shares of the county population age 65 or older, ages 25 to 54, and nonwhite or Hispanic, together with the ratio of CVAP to voting-age population. Using IRS migration data, we construct net and gross migration rates, where gross migration is the sum of inflows and outflows, each relative to the previous year's population. We attach these annual measures by county and election year and add all six jointly to the baseline regression without reconstructing treatment or matching. The weights, remaining controls, fixed effects, and clustering are unchanged. Because these annual measures are unavailable for some county-years, the adjusted specification uses a smaller sample. Re-estimating the baseline specification on this common sample yields an event-study pattern very similar to that of the baseline.

Figure A21 shows that the demographic- and migration-adjusted event-study paths remain close to the baseline, although the Republican-vote and turnout declines are slightly attenuated. Pooling the four post-event indexes, the Democratic-Republican vote-rate margin rises by 0.46 percentage points, Republican votes per CVAP fall by 0.27 points, Democratic votes per CVAP rise by 0.19 points, other-candidate votes per CVAP change by 0.01 points, and turnout falls by 0.08 points. The pooled pre-event estimates are close to zero for all five outcomes.

D.2 CES evidence on turnout composition vs. vote switching

To distinguish changes in turnout composition from vote switching, we use the cumulative Cooperative Election Study (CES) for 2006-2024. Although the CES is fielded annually, we use its even-year election surveys because they report respondents' county, party identification and leaning, pre-election vote intention, reported vote choice, and validated turnout. We attach each respondent-office-year observation to the corresponding county, office, and election year in the matched chronological design.

Respondents inherit the treatment status and chronological election index of the county-office-year observation to which they are attached. We estimate respondent-level versions of the baseline event-study specification with the same county controls and fixed effects. All matched CES specifications also include age, age squared, and sex as common main effects. Within each matched comparison-county-office-year cell, we normalize the CES survey weights to sum to one and multiply them by the matched-comparison weight. This preserves the weight assigned to each county-office-year observation while retaining the relative CES weights across respondents. Standard errors are two-way clustered by county and matched comparison group. Across specifications, we retain between 1,335 and 1,499 of the 1,508 matched WARN events.[9]

[9] The switching samples are smaller than the turnout samples because they require a Democratic or Republican pre-election vote intention, self-reported turnout, and an observed response for the relevant office. The expanded observation counts exceed the number of distinct respondents because a respondent

We first measure switching relative to respondents' stated pre-election intentions. Among respondents who intend to vote Republican for the relevant office and later report voting in the election, the outcome equals one if they report a Democratic vote for that office and zero if they report a Republican or other-candidate vote or explicitly report not voting in that race. We construct the corresponding outcome symmetrically for respondents who intend to vote Democratic. These outcomes measure departures from stated intentions, not changes in long-run party identification. Panels A and B of Figure A22 show no systematic increase in either form of switching after a WARN event. Pooling the four post-event indexes, Republican-to-Democratic switching changes by -0.22 percentage points, with a standard error of 0.41 points, while Democratic-to-Republican switching changes by -0.34 points, with a standard error of 0.55 points. The pooled pre-event estimates are also small and statistically insignificant.

Table A13 expands this accounting to all realized outcomes conditional on respondents' office-specific pre-election intentions. Panel A shows that 86.8 percent of Democratic intenders vote Democratic and 90.7 percent of Republican intenders vote Republican; only 2.9 and 1.7 percent, respectively, vote for the opposing major party. Panel B shows no statistically significant Post change in any outcome among major-party intenders. These coefficients need not match the pooled Post coefficients in Panels A and B of Figure A22 because the switching regressions are conditioned on respondents who report voting in the election, whereas the table retains the full sample of classified intenders, including reported nonvoters, and assigns reported nonvoters to the neither-major-party outcome. For Republican intenders, the point estimates imply a 1.05-percentage-point decline in Republican voting and a 1.36-point increase in neither-major-party voting, but both are statistically insignificant and smaller in magnitude than the corresponding Pre estimates in Panel C. The other-party-intention estimates are larger but less precise and accompanied by sizable Pre differences. Overall, the expanded accounting provides no clear evidence of a systematic post-WARN reallocation conditional on intention. Neither-major-party voting combines nonvoting, other-party voting, and office undervoting, and stated intention need not precede WARN exposure; the estimates therefore do not identify causal transitions from a predetermined state.

We next examine the composition of turnout among respondents who identify with or lean toward a major party. We use the specification described above, but interact the treatment coefficients with an indicator for Republican identification or leaning.[10] The reported coefficient is therefore the Republican-minus-Democratic difference in the validated-turnout response. Panel C of Figure A22 shows a negative difference at each post-event index. Pooling these indexes, the Republican turnout response is 4.15 percentage points lower than the Democratic response, with a standard error of 2.32 points. The corresponding pooled pre-event difference is close to zero and statistically insignificant. We then restrict the sample to respondents who said before the election that they definitely or probably intended to vote or planned to vote early. Panel D shows a similar pattern: the pooled Republican-minus-Democratic difference is -5.11 percentage points, with a standard error of 2.94 points, while the pooled pre-event difference

may enter more than one matched comparison or contribute more than one office outcome.

[10]We also allow the county controls and fixed effects to differ across the two groups.

remains small and statistically insignificant. This result suggests that the turnout gap is not explained solely by weaker stated pre-election turnout intentions among Republican leaners and is slightly more pronounced among respondents who expressed an intention to vote before the election.

As an additional robustness exercise, we examine Democratic vote choice among major-party voters using a different specification. The outcome equals one for a Democratic vote and zero for a Republican vote. Rather than using the matched event-study design, this exercise relates vote choice to an indicator for a qualifying WARN event in the respondent's county during the election-year window. The repeated-cross-section specifications include state-by-year and office-by-year fixed effects, and then sequentially add respondents' demographics, party identification, and pre-election vote intention. The CES also contains two multi-wave panels that follow respondents across the 2010, 2012, and 2014 elections and across the 2020, 2022, and 2024 elections. We pool the two panels and estimate a final specification with respondent and office-by-year fixed effects. All specifications use CES survey weights and cluster standard errors by county.

Table A14 reports the results. Panel A uses all observations available for each specification. A WARN event is initially associated with a 2.9-percentage-point increase in the probability of a Democratic vote and a 2.3-point increase after adding demographic controls. The estimate falls to 0.3 points after controlling for party identification and to 0.2 points after adding pre-election vote intention; neither estimate is statistically significant. Because party identification captures the partisan mix of voters, this attenuation indicates that the positive association is accounted for by differences in partisan composition rather than vote-choice changes within partisan types. The respondent-fixed-effect estimate is -0.6 points and is also statistically insignificant. Panel B re-estimates the first four specifications on a common sample with nonmissing values for all variables, keeping the same sample across all columns, and yields very similar results.

Taken together, the absence of systematic switching in both the binary and the full-transition estimates, the lower post-WARN validated-turnout response among Republican than Democratic leaners both overall and among respondents who intended to vote, and the broader vote-choice results are consistent with a change in turnout composition rather than with voters switching between parties.

D.3 Presidential campaign visits

Presidential campaigns may respond to local layoff announcements by shifting candidate travel toward affected counties. If Democratic campaigns devote more attention than Republican campaigns, differential campaign attention could contribute to the Democratic vote offset. We use candidate visits, one observable but narrow form of campaign activity, to examine this possibility.

The Campaign Visits Database (CVD) covers appearances by presidential and vice-presidential candidates in the five presidential elections from 2008 through 2024. We construct two county-presidential-year outcomes: an indicator for any Democratic visit minus an indicator for any Republican visit, and the number of Democratic visits minus the number of Republican visits. The visit-count outcome measures distinct campaign stops, so joint appearances by candidates from the same party at the same event are counted once. Visits are rare: among 15,710 county-presidential-year observations over the period covered by the CVD, 390 include a Democratic visit, and 541 include a Republican visit. Because visits occur only in presidential election years, we apply the presidential target-office matched design described in Section 2.4. Panel C of Table 3 shows no meaningful relative increase in Democratic visits after a WARN event. The pooled Post estimate is 0.34 percentage points for the any-visit outcome, with a standard error of 0.86, and 0.007 visits for the visit-count outcome, with a standard error of 0.012. The pooled Pre estimates are also close to zero.

The conclusion is also unchanged in a separate county-presidential-year specification that defines exposure using qualifying WARN events before September of the election year, a window that gives campaigns more time to respond before Election Day. The exercise, therefore, weakens a candidate-travel explanation for the Democratic offset but does not rule out differential advertising, field offices, local surrogates, media attention, or other campaign responses.

D.4 Local transfer payments

We use annual county transfer payments from the Bureau of Economic Analysis through 2022, the latest year for which program-level data are available. The data measure transfers received by individuals and institutions located in the county, not spending chosen by the county government. We group the underlying programs according to how they are likely to respond to local economic distress. Shock-responsive safety-net transfers include unemployment insurance, SNAP, the Earned Income Tax Credit, and other income-maintenance and public-assistance benefits. Active and administratively mediated support includes education and training assistance and other government transfers to individuals. Slow-moving categorical entitlements include Social Security, retirement and disability, Medicare, Supplemental Security Income, and veterans' benefits. The final category, private and community transfers, combines current transfers received by nonprofit institutions and transfers to individuals from businesses. For each category, we sum the component programs, convert the total to real 2022 dollars per capita, and use the natural logarithm of one plus that amount.

Because transfer payments are annual, we reorganize the matched design around calendar years rather than elections. We begin with each treated event and its matched controls and retain one observation per county within the comparison, so the same event is not repeated merely because several electoral offices appear in the chronological design. We define $k = 0$ as the calendar year containing the first qualifying WARN month. Because the event can occur at any point during that year, $k = 0$ is a transition year that may contain both pre- and post-WARN months. The preceding calendar year, $k = -1$, is the last complete pre-WARN year and serves as the omitted period; $k = -3$ and $k = -2$ provide earlier pre-event observations, while $k = 1, 2, 3$ are complete post-WARN years. We estimate a no-controls specification and a baseline specification that, paralleling Equation (3), adds annual measures of unemployment, labor force, per capita income, and COVID cases. The no-controls specification captures the full change in payments following a WARN event, including any portion attributable to broader contemporaneous deterioration. Comparing it with the baseline specification shows how much of that response remains after accounting for measured local economic conditions and is therefore more specific to a WARN event. Both retain the original inverse-distance matching weights and include county-by-comparison, relative-year, and state-by-calendar-year fixed effects, with standard errors two-way clustered by county and matched comparison group. Thus, both compare the change in payments from $k = -1$ in treated counties with the corresponding change in their matched controls, net of common state-year shocks, while the baseline additionally conditions on observed annual economic conditions. Across specifications, the analysis retains at least 1,445 of the 1,508 matched events, or about 96 percent.

Figure A23 shows that shock-responsive safety-net payments rise in the WARN year and remain elevated through $k = 3$. In the baseline specification, pooling $k = 0, 1, 2, 3$ yields an increase of approximately 0.5 percent (Table 3, Panel D, Column 1). The no-controls estimates are larger, indicating that part of the total increase is associated with broader deterioration in local economic conditions. The remaining positive baseline response, however, shows that the measured economic conditions do not fully account for the increase in shock-responsive transfers. Active and administratively mediated support, as well as slow-moving categorical entitlements, shows little systematic response. Private and community transfers rise more modestly, by 0.3 percent (Table 3, Panel D, Column 4). The pre-WARN estimates are close to zero across all four categories. Among public programs, the concentration of the response in the shock-responsive safety net, together with the absence of a response in the more administratively mediated and potentially discretionary category, is suggestive of automatic stabilizers rather than transfer allocations responding to post-WARN electoral changes.

This category pattern is suggestive, but it does not directly test whether the electoral response accounts for later changes in transfer payments. We therefore ask more directly whether the transfer responses are associated with the electoral changes that follow WARN events. For each county, matched comparison, and transfer year, we identify the most recent regular post-WARN election completed before January 1 of that year. Restricting attention to elections completed before the transfer year ensures that the electoral outcome precedes all payments measured during that year.[11] For the selected office, we measure the cumulative changes in turnout and Democratic two-party vote share, defined as Democratic votes divided by the sum of Democratic and Republican votes, relative to the last pre-WARN election of the same office. Before calculating these changes, we remove the corresponding state-by-office-by-election-year averages, so the measures capture county-specific electoral movements rather than statewide changes in a particular race. The most recent changes remain in place until a newer election becomes available. Before any post-WARN election, both changes are set to zero, and a separate indicator records whether such an election has occurred. This distinguishes years before an electoral response can be observed from post-election years in which the measured change happens to equal zero.

The voting-adjusted specification adds these two electoral changes and the post-election indicator to the baseline regression. To ensure that the comparison reflects only the addition of the electoral measures rather than a change in the sample, we re-estimate the baseline specification using the same observations as the voting-adjusted specification. The voting-adjustment gap is this common-sample baseline coefficient minus the voting-adjusted coefficient at each relative year. If the observed electoral

[11] When several offices are contested on the same date, we select one using a fixed presidential, gubernatorial, Senate, and House ordering.

changes accounted for a positive transfer response, adding them would attenuate the WARN coefficient and produce a positive gap. Instead, the baseline and voting-adjusted paths are nearly identical, and the estimated gaps remain close to zero at every relative year and for every transfer category.

The concentration of the response in shock-responsive programs, together with the negligible voting-adjustment gaps, is consistent with transfer payments rising primarily because existing programs respond to greater local economic need rather than because electoral changes lead to increased economic support. While the results favor an automatic-stabilizer interpretation, they do not rule out political effects operating through other policies, administrative decisions, changes in representation, or over longer time horizons.

Appendix references

Devine, Christopher J. (2020). "Campaign Visits Database." http://www.christopherjdevine.com/data.html. Accessed July 2026.

Killeen, Benjamin D., Jie Ying Wu, Kinjal Shah, Anna Zapaishchykova, Philipp Nikutta, Aniruddha Tamhane, Shreya Chakraborty, Jinchi Wei, Tiger Gao, Mareike Thies, and Mathias Unberath (2020). "A county-level dataset for informing the United States' response to COVID-19." *arXiv preprint arXiv:200400756*. doi:10.48550/arXiv.2004.00756.

Krolikowski, Pawel M. and Kurt G. Lunsford (2024). "Advance Layoff Notices and Aggregate Job Loss." *Journal of Applied Econometrics*, 39(3), pp. 462--480. doi:10.1002/jae.3032.

US Census Bureau (2026). "Census Geocoder." https://geocoding.geo.census.gov/geocoder/geographies/addressbatch?form. Accessed July 2026.

US Department of Housing and Urban Development (2026). "HUD-USPS ZIP Code Crosswalk Data." https://www.huduser.gov/portal/datasets/usps_crosswalk.html. Accessed July 2026.

US Department of Labor (2026). "WARN Act Compliance Assistance." https://www.dol.gov/agencies/eta/layoffs/warn. Accessed July 2026.

Table A1: WARN notice coverage by state

State	Coverage	Notices	Affected Workers	Avg. Affected Workers per Notice	Available County Info (%)
AK	Feb 2006--Jun 2025	73	9,800	134.2	86.3
AL	Feb 1998--Dec 2025	1,009	179,233	177.6	97.8
AZ	Apr 2010--Dec 2025	639	101,373	158.6	93.7
CA	Nov 2005--Dec 2025	23,556	2,226,650	94.5	99.6
CO	Jan 2006--Dec 2025	1,126	133,405	118.5	54.1
CT	Jan 1998--Nov 2025	1,228	123,996	101.0	97.1
FL	Jan 1998--Dec 2025	4,911	647,255	131.8	97.2
HI	Jan 2019--Oct 2021	259	40,358	155.8	92.3
IA	Jul 2005--Dec 2025	1,032	79,652	77.2	99.7
ID	Jan 2009--Oct 2025	182	21,699	119.2	96.2
IL	Jan 1999--Dec 2025	3,561	535,142	150.3	99.6
IN	Jan 1997--Dec 2025	1,977	284,545	143.9	97.4
KS	Aug 1998--Oct 2025	840	122,456	145.8	97.4
KY	Oct 1998--Dec 2025	1,095	144,125	131.6	89.6
LA	Jan 2007--Dec 2025	593	81,241	137.0	96.0
MD	Jan 2000--Dec 2025	1,673	187,379	112.0	85.7
MI	Jan 1990--Dec 2025	2,915	452,255	155.1	99.2
MN	Jun 2000--Nov 2025	773	99,832	129.1	97.8
MO	Jan 2000--Dec 2025	1,406	200,705	142.7	96.7
MS	Jul 2010--Dec 2025	296	48,629	164.3	96.6
NC	May 1996--Dec 2025	3,050	419,387	137.5	99.4
ND	Jul 2015--Oct 2025	53	5,331	100.6	75.5
NJ	Jan 2004--Dec 2025	2,196	327,468	149.1	94.9
NM	Jan 2016--Sep 2025	117	15,792	135.0	98.3
NV	Jul 2012--Dec 2019	522	174,643	334.6	97.3
NY	Sep 2001--Dec 2025	7,141	769,750	107.8	99.6
OH	Jun 1996--Dec 2025	3,375	500,701	148.4	98.9
OK	Nov 1999--Jan 2025	507	76,522	150.9	87.6
OR	Jan 2005--Dec 2025	788	82,280	104.4	95.2
PA	Jan 2001--Dec 2025	3,139	428,194	136.4	99.5
RI	Jan 2009--Sep 2025	122	15,438	126.5	90.2
SC	Jun 2008--Dec 2013	650	32,993	50.8	98.6
TN	Jan 2011--Dec 2025	1,139	123,702	108.6	99.2
TX	Jan 1999--Dec 2025	6,766	682,864	100.9	100.0
UT	Jan 2009--Dec 2025	274	35,520	129.6	95.3
VA	Jul 1994--Dec 2025	2,231	316,528	141.9	86.4
WA	Jan 2004--Dec 2025	1,414	211,360	149.5	86.6
WI	Jan 1996--Dec 2025	2,843	268,405	94.4	99.6
WV	Mar 2011--Aug 2025	281	45,445	161.7	98.6

Note: Each row describes the WARN records available for one state. Coverage reports the first and last months for which the assembled administrative records are available. Notices is the number of establishment-level WARN notices; Affected Workers sums the workers reported in those notices; and Avg. Affected Workers per Notice is their ratio. Available County Info reports the percentage of notices with sufficient geographic information to assign the affected worksite to a county. The table contains 85,752 notices from 39 states. Coverage varies across states because historical administrative records are not uniformly available. See Appendix A.1 for further details.

Table A2: Election sample by year and office

	Presidential				Senate				Gubernatorial				House			
Year	Elec.	States	Counties	CVAP	Elec.	States	Counties	CVAP	Elec.	States	Counties	CVAP	Elec.	States	Counties	CVAP
1996	1	50	3,112	184.6m	32	32	2,288	97.1m	11	11	566	23.3m	434	49	3,111	184.2m
1997	--	--	--	--	--	--	--	--	2	2	156	10.5m	--	--	--	--
1998	--	--	--	--	33	33	2,056	136.4m	35	35	2,147	146.8m	434	49	3,111	188.0m
1999	--	--	--	--	--	--	--	--	2	2	202	5.0m	--	--	--	--
2000	1	50	3,113	193.0m	33	33	1,879	144.0m	11	11	566	24.3m	434	49	3,111	192.5m
2001	--	--	--	--	--	--	--	--	2	2	155	10.8m	--	--	--	--
2002	--	--	--	--	32	32	2,287	103.2m	35	35	2,148	154.3m	434	49	3,111	197.1m
2003	--	--	--	--	--	--	--	--	2	2	202	5.1m	--	--	--	--
2004	1	50	3,112	201.6m	33	33	2,057	146.5m	11	11	566	25.4m	434	49	3,111	201.1m
2005	--	--	--	--	--	--	--	--	2	2	155	11.2m	--	--	--	--
2006	--	--	--	--	33	33	1,878	154.0m	35	35	2,148	161.2m	434	49	3,111	205.6m
2007	--	--	--	--	--	--	--	--	2	2	202	5.3m	--	--	--	--
2008	1	50	3,112	210.3m	32	32	2,287	109.8m	11	11	566	26.8m	434	49	3,111	209.8m
2009	--	--	--	--	--	--	--	--	2	2	155	11.3m	--	--	--	--
2010	--	--	--	--	33	33	2,058	152.6m	35	35	2,149	164.1m	434	49	3,112	209.3m
2011	--	--	--	--	--	--	--	--	2	2	202	5.4m	--	--	--	--
2012	1	50	3,113	214.6m	33	33	1,878	160.2m	11	11	566	27.4m	434	49	3,112	214.2m
2013	--	--	--	--	--	--	--	--	2	2	154	11.8m	--	--	--	--
2014	--	--	--	--	32	32	2,287	114.5m	35	35	2,149	171.9m	434	49	3,111	219.1m
2015	--	--	--	--	--	--	--	--	2	2	202	5.5m	--	--	--	--
2016	1	50	3,112	224.1m	33	33	2,057	162.9m	11	11	566	28.5m	434	49	3,110	223.6m
2017	--	--	--	--	--	--	--	--	2	2	154	12.2m	--	--	--	--
2018	--	--	--	--	33	33	1,877	170.7m	35	35	2,148	179.1m	434	49	3,110	227.9m
2019	--	--	--	--	--	--	--	--	2	2	202	5.6m	--	--	--	--
2020	1	50	3,112	232.9m	32	32	2,287	121.4m	11	11	566	29.7m	434	49	3,111	232.4m
2021	--	--	--	--	--	--	--	--	2	2	154	12.6m	--	--	--	--
2022	--	--	--	--	33	33	2,059	172.5m	35	35	2,150	186.6m	434	49	3,112	237.1m
2023	--	--	--	--	--	--	--	--	2	2	202	5.6m	--	--	--	--
2024	1	50	3,113	240.3m	33	33	1,878	179.8m	11	11	566	30.7m	434	49	3,112	239.8m

Note: Each row is an election year. Within each office group, Elec. counts regular election events, States counts represented states including the District of Columbia, Counties counts represented counties in the county-office-year outcome panel, and CVAP reports total citizen voting-age population in millions. CVAP is summed once per county-year-office cell, so a county is not counted twice when a state holds two Senate contests in the same year. Presidential, Senate, and gubernatorial counts are distinct election events; House counts are district contests in state-years represented in the usable county panel. The sample retains elections held from November 1 through November 8, excludes special and recall elections, and requires usable county returns and positive CVAP. Of 7,436 calendar events from 1996 through 2024, 35 are excluded by the date or special/recall restrictions and another 32 lack usable county-level outcome or CVAP support, leaving 7,369 election events. Dashes denote no election in the year-office cell. See Appendix A.2 for further details.

Table A3: Candidate and office-holder identifier support by office

Office	Years	Obs.	Incumbent candidate in election	Incumbent-candidate party identified (as % of prev. column)	Pre-election office-holder party identified
Presidential	1996--2024	24,817	12,408 (50.0%)	12,408 (100.0%)	24,817 (100.0%)
Governor	1996--2024	22,005	12,787 (58.1%)	12,787 (100.0%)	20,702 (94.1%)
House	1996--2024	46,415	6,677 (14.4%)	6,677 (100.0%)	6,677 (14.4%)
Senate	1996--2024	31,006	26,818 (86.5%)	26,818 (100.0%)	30,011 (96.8%)

Note: The table reports identifier availability in the county-office-year election panel. Obs. is the number of county-office-year observations for the indicated office. Incumbent Candidate in Election reports the number and share for which a major-party incumbent candidate can be identified. Incumbent Candidate Party Identified reports party availability conditional on identifying such a candidate. Pre-election Office-Holder Party Identified reports whether the party holding the office before the election can be assigned, whether or not the office holder is running. House identifiers are retained only when the county-level return can be linked to an unambiguous major-party assignment. See Appendix A.2 for further details.

Table A4: Support and concentration of the matched comparisons

Panel A. Control availability through the matching stages

			Control counties per treated episode, mean (median)			
Matching cell	Episodes	Share (%)	Before calipers	After WARN-history	After all	Selected
Exact state-year-month	1,462	96.9	62.4 (52)	47.6 (40)	17.6 (13)	3.53 (4)
Fallback state-year	46	3.1	51.0 (44)	21.1 (8)	1.6 (1)	1.22 (1)
All matched episodes	1,508	100.0	62.1 (52)	46.8 (39)	17.1 (13)	3.46 (4)

Panel B. Concentration of inverse-distance weights

Matching cell	Largest control weight		Effective controls	Largest weight above 0.50
	Median	90th pct.	Mean	Share (%)
Exact state-year-month	0.277	0.564	3.51	16.4
Fallback state-year	1.000	1.000	1.19	100.0
All matched episodes	0.278	1.000	3.43	19.0

Note: Panel A reports the number and share of matched treated events by matching cell and the number of distinct potential control counties remaining at each stage. Entries in the control-count columns are means, with medians in parentheses. The fallback state--year cell is used only when no control in the exact state--year--month cell survives all calipers. Selected is the number of controls retained after keeping up to the four nearest acceptable counties. Panel B describes concentration in the final inverse-distance weights. The effective number of controls is $1/\sum_j w_{sj}^2$, which equals one when a comparison is effectively based on one control and four when four controls receive equal weight. Weights assigned to repeated records for the same actual control county are combined before calculating Panel B. The table covers all 1,508 matched treated events: 1,462 exact-cell matches and 46 fallback matches. See Appendix B.3.2 for further details.

Table A5: Characteristics of matched and unmatched treated episodes

Variable	Matched episodes ($N = 1,508$)	Unmatched episodes ($N = 195$)	Std. difference
Treatment year	2014.0	2015.6	-0.252
Affected workers	431	2,958	-0.383
WARN exposure (% of labor force)	1.811	1.700	0.050
Prior turnout	49.254	50.107	-0.069
Prior Republican share	0.592	0.497	0.587
Second-lag Republican share	0.571	0.495	0.467
Change in prior Republican share	0.021	0.002	0.164
Log labor force	9.786	11.009	-0.877
Pre-event unemployment rate	5.666	6.426	-0.248
Change in unemployment rate	0.508	1.798	-0.320
Prior WARN exposure	0.005	0.025	-1.093
Prior WARN history	0.004	0.022	-1.073

Note: The table compares treated events that enter the chronological matched sample with treated events outside common support. Entries are unweighted means across treated events. WARN exposure sums, over the three-month window, each month's affected-worker count divided by that month's county labor force and multiplied by 100. Standardized differences are signed and equal the matched mean minus the unmatched mean divided by the pooled standard deviation. The table covers all 1,703 match-ready treated events: 1,508 matched and 195 unmatched. Of the unmatched events, 61 lose all potential controls at the WARN-history calipers and 134 at the remaining political and labor-market calipers. See Appendix B.3.2 for further details.

Table A6: Chronological support by office and election index

Office	$h^C = -3$	-2	-1	0	1	2	3	4
Presidential	786	609	772	585	810	570	693	526
Gubernatorial	595	706	642	735	604	665	555	401
House	1,311	1,396	1,355	1,361	1,396	1,259	1,186	930
Senate	883	845	964	925	865	893	802	600

Note: Entries are distinct matched treated events with usable Republican votes per CVAP, Democratic votes per CVAP, other votes per CVAP, turnout, and preferred county controls for the indicated office and chronological election index. The chronological index orders included election dates in the sequence counties face them; offices contested on the same date share an index but remain separate county-office outcomes. The same WARN event may therefore appear under more than one office at a shared election date and under more than one election index over time, so the cells are not mutually exclusive. The underlying matched sample contains 1,508 treated events. See Appendix B.4 for further details.

Table A7: Office-specific support by election index

Office	$h_o = -2$	-1	0	1	2	3
Presidential	1,306	1,443	1,443	1,443	1,246	819
Gubernatorial	1,138	1,300	1,341	1,341	997	692
House	1,443	1,444	1,444	1,443	1,339	1,241
Senate	1,308	1,399	1,399	1,399	1,277	889

Note: The four office-specific designs are matched separately. Entries are distinct treated comparisons with usable Republican votes per CVAP, Democratic votes per CVAP, other votes per CVAP, turnout, and preferred county controls at the indicated office-specific election index. Every included comparison has a usable same-office baseline election at $h_o = 0$. The same universe of qualifying WARN events is eligible for each design, but the realized samples differ when the required election history or outcome is unavailable. Differences across columns show how support changes as earlier or later contests for the same office fall outside the sample period. See Appendix B.4 for further details.

Table A8: Pooled pre- and post-WARN effects under alternative event-time orderings

	Democratic-Republican votes/CVAP		Republican votes/CVAP		Democratic votes/CVAP		Other votes/CVAP		Turnout	
Event-time ordering	Pre	Post	Pre	Post	Pre	Post	Pre	Post	Pre	Post
	0.041	0.501***	-0.032	-0.318***	0.010	0.182**	0.001	0.000	-0.021	-0.136
Chronological	(0.133)	(0.169)	(0.079)	(0.105)	(0.080)	(0.088)	(0.028)	(0.025)	(0.083)	(0.093)
	0.010	0.446***	-0.012	-0.314***	-0.002	0.132*	-0.011	-0.005	-0.026	-0.188**
Target-office	(0.121)	(0.152)	(0.067)	(0.094)	(0.072)	(0.080)	(0.022)	(0.021)	(0.066)	(0.083)

Note: The table reports separately estimated pooled Pre and Post coefficients from the preferred specification under two event-time orderings. The Democratic-Republican outcome is Democratic votes minus Republican votes, divided by CVAP. In the chronological design, Pre combines $h^C = -3, -2, -1$, Post combines $h^C = 1, 2, 3, 4$, and $h^C = 0$ is omitted. In the target-office design, Pre combines $h^O = -2, -1$, Post combines $h^O = 1, 2, 3$, and $h^O = 0$ is omitted. The chronological ordering follows the sequence of included election dates, while the target-office ordering compares treated and control outcomes within the same office before pooling across offices. Both designs use matched treated-control comparisons, inverse-distance weights, preferred county controls, county-by-comparison and event-index fixed effects, and state-by-office-by-year fixed effects. The chronological regressions use 122,698 stack-expanded county-office-election observations; the target-office regressions use 138,499. All outcomes are measured in percentage points. Standard errors in parentheses are two-way clustered by county and matched comparison group. $*$, $**$, and $***$ denote significance at the 10, 5, and 1 percent levels. See Appendix C.1.1 for further details.

Table A9: County-level Democratic two-party-lead estimates

Period	All matched events (1)	Previously close contest, ± 10 rematched (2)	Previously close contest, ± 10 pruned (3)	Previously close contest, ± 5 (4)	Previously close contest and county, ± 10 (5)	Previously close contest and county, ± 5 (6)
Panel A. Chronological design						
Pre	-0.234	-0.054	-0.083	0.289	-0.573	4.657
	(0.633)	(0.967)	(0.965)	(1.344)	(3.296)	(8.317)
Post	0.800	1.067	1.066	2.285	1.275	0.540
	(0.686)	(1.103)	(1.105)	(1.471)	(3.593)	(5.364)
Panel B. Target-office design						
Pre	-0.008	-1.106	-1.003	-0.275	-2.656	-0.589
	(0.561)	(0.993)	(1.003)	(1.258)	(3.389)	(7.501)
Post	0.716	0.076	0.179	0.216	-0.914	-1.506
	(0.603)	(0.945)	(0.945)	(1.217)	(3.385)	(6.762)

Note: The outcome equals one when Democratic votes exceed Republican votes within the county and zero when Republican votes exceed Democratic votes; exact ties and results without positive votes for both parties are excluded. It measures the county two-party lead, not the winner of the electoral jurisdiction. Column (1) retains the paper's matched sample. Columns (2) and (4)--(6) reconstruct the matches after requiring the preceding electoral-jurisdiction margin to lie within the stated bandwidth; columns (5) and (6) also require the preceding county margin to lie within that bandwidth. Column (3) instead applies the 10-point contest screen to the existing matches and renormalizes the retained control weights. Chronological Pre pools $h^C = -3, -2, -1$ and Post pools $h^C = 1, 2, 3, 4$; target-office Pre pools $h^O = -2, -1$ and Post pools $h^O = 1, 2, 3$; $h = 0$ is omitted. Across columns (1)--(6), the chronological regressions contain 113,729, 42,143, 41,915, 24,427, 6,754, and 1,689 observations; the target-office regressions contain 129,380, 41,245, 39,916, 21,125, 8,059, and 1,525. All specifications retain the baseline controls, fixed effects, weights, and two-way clustering by county and matched comparison group. Estimates are percentage points; standard errors are in parentheses. $*$, $**$, and $***$ denote significance at the 10, 5, and 1 percent levels. See Appendix C.1.2.1 for further details.

Table A10: WARN events and US House election outcomes

Outcome	No controls	Four controls
Panel A. District outcomes		
Democratic victory	-1.488 (2.689)	0.149 (2.539)
Democratic--Republican district margin	-1.946* (1.085)	-1.799 (1.127)
Panel B. Gains and retention		
Democratic gain	-2.722 (1.959)	-1.701 (1.718)
Democratic retention	-2.155 (13.196)	-5.778 (14.266)

Note: Democratic victory equals one when the Democratic candidate wins the US House district and zero when the Republican candidate wins; the district margin is $100(D-R)/(D+R)$. Democratic gain and retention refer to districts previously held by Republicans and Democrats, respectively. The first column omits the four contemporaneous county controls, while the second includes them. Treated districts require a qualifying event in an unambiguously assigned county; control districts additionally require complete WARN coverage and no qualifying exposure in any intersecting county. Eligible districts are then rematched using the paper's matching rule. For the no-control/four-control specifications, observations are 2,264/2,178 for victory, 1,550/1,496 for the margin, 1,486/1,438 for gains, and 116/110 for retention. All specifications use the matched weights and absorb matched district-unit, election-index, and state-by-election-year fixed effects. Standard errors in parentheses are two-way clustered by matched comparison group and by congressional district within an apportionment period. Estimates are percentage points for the binary outcomes and district-margin points for the margin. $*$, $**$, and $***$ denote significance at the 10, 5, and 1 percent levels. See Appendix C.1.2.2 for further details.

Table A11: Statewide WARN exposure and election outcomes

	Panel A. Same-office intervals		*Panel B. Chronological intervals*	
Specification	Democratic victory	Democratic--Republican margin	Democratic victory	Democratic--Republican margin
Affected workers per 1,000				
	0.857	-0.246	0.558	-0.022
No lagged margin	(1.315)	(0.695)	(0.749)	(0.392)
	0.485	-0.332	0.528	-0.058
Lagged margin	(1.600)	(0.678)	(0.815)	(0.346)
Annualized notices per 100,000				
	-0.881	-0.733	-0.676	-0.372
No lagged margin	(1.248)	(0.830)	(1.050)	(0.498)
	-1.653	-0.793	-0.895	-0.394
Lagged margin	(1.471)	(0.800)	(0.986)	(0.408)
IHS affected workers per 1,000				
	4.149	-0.425	3.251	0.908
No lagged margin	(6.108)	(2.935)	(4.136)	(1.767)
	3.260	-0.092	3.816	1.311
Lagged margin	(6.677)	(2.836)	(4.267)	(1.682)

Panel C. Same-office estimates by office		
	Democratic victory	Democratic--Republican margin
	-0.989	0.368
Gubernatorial	(2.198)	(1.192)
	0.670	-0.526**
Presidential	(1.261)	(0.235)
	3.905	-1.699
US Senate	(3.095)	(1.640)

Note: Democratic victory equals one when the formally listed Democratic candidate wins the ending election and zero otherwise; the margin is $100(D-R)/(D+R)$. Same-office intervals join consecutive regular elections in the same presidential, gubernatorial, or Senate-class series; chronological intervals join consecutive distinct regular election dates within a state, regardless of office. The primary exposure is annualized WARN-affected workers per 1,000 members of the average state labor force; the alternatives use annualized notices per 100,000 labor-force participants or the inverse hyperbolic sine of the primary rate and use the corresponding primary-model samples. Lagged-margin specifications allow the coefficient on the preceding same-office margin to differ by office. Panel C reports this specification separately by office; no chronological counterpart is estimated. For the no-lag/lag specifications, observations for victory/margin are 490/482 and 482/475 in Panel A, and 578/570 and 560/550 in Panel B. Panel C contains 146/146 observations for gubernatorial elections, 165/165 for presidential elections, and 171/164 for US Senate elections. Pooled models absorb office-series and office-by-ending-election-year fixed effects. Standard errors in parentheses are clustered by state. Victory estimates are percentage points; margin estimates are margin points. $*$, $**$, and $***$ denote significance at the 10, 5, and 1 percent levels. See Appendix C.1.2.3 for further details.

Table A12: WARN events, realized job destruction, and electoral outcomes

Panel A. WARN events and realized job destruction					
Definition	Treated mean	Control mean	Difference	Comparisons	Counties
Primary definition	0.520	-0.316	0.836*** (0.075)	1,255	1,944
Complete-quarter definition	0.490	-0.202	0.692*** (0.072)	955	1,696

Panel B. Pooled electoral estimates					
Specification	D-R votes/CVAP	Rep./CVAP	Dem./CVAP	Other/CVAP	Turnout
Primary definition					
Same-sample baseline: Pre	0.114 (0.141)	-0.078 (0.086)	0.036 (0.083)	0.004 (0.031)	-0.037 (0.088)
Same-sample baseline: Post	0.528*** (0.180)	-0.321*** (0.114)	0.207** (0.093)	-0.003 (0.027)	-0.117 (0.100)
JDR-adjusted: Pre	0.103 (0.143)	-0.082 (0.087)	0.020 (0.086)	0.011 (0.031)	-0.051 (0.090)
JDR-adjusted: Post	0.464** (0.184)	-0.282** (0.115)	0.182* (0.094)	0.006 (0.027)	-0.094 (0.100)
Complete-quarter definition					
Same-sample baseline: Pre	0.090 (0.162)	-0.066 (0.097)	0.023 (0.094)	0.013 (0.036)	-0.030 (0.098)
Same-sample baseline: Post	0.521*** (0.202)	-0.273** (0.120)	0.248** (0.108)	0.011 (0.032)	-0.015 (0.105)
JDR-adjusted: Pre	0.086 (0.167)	-0.071 (0.100)	0.015 (0.097)	0.008 (0.037)	-0.049 (0.100)
JDR-adjusted: Post	0.520** (0.209)	-0.259** (0.123)	0.261** (0.112)	0.003 (0.033)	0.005 (0.105)

Panel C. Association between job destruction and electoral outcomes					
Definition	D-R votes/CVAP	Rep./CVAP	Dem./CVAP	Other/CVAP	Turnout
Primary definition	0.065* (0.038)	-0.054** (0.023)	0.011 (0.021)	-0.003 (0.006)	-0.046** (0.020)
Complete-quarter definition	-0.004 (0.049)	-0.028 (0.029)	-0.032 (0.027)	0.004 (0.007)	-0.056** (0.026)

Note: JDR is firm job destruction as a percentage of average county employment. The primary definition compares the two complete quarters ending before the qualifying WARN window with the WARN-trigger quarter and the following quarter. The complete-quarter definition instead uses the first two complete quarters after the trigger quarter. Panel A reports matched-weighted treated and control means and the within-comparison difference in the change in JDR. Panel B reports pooled Pre and Post WARN coefficients from the baseline specification on each QWI-compatible sample and from the corresponding specification that allows the electoral path to vary with the change in JDR. Pre combines $h^C = -3, -2, -1$, Post combines $h^C = 1, 2, 3, 4$, and $h^C = 0$ is omitted. Panel C reports the Post-minus-Pre change in the association between JDR and each electoral outcome. Electoral outcomes and JDR are measured in percentage points. Panels B and C use the QWI-compatible electoral regression samples, containing 1,252 and 953 matched comparisons under the primary and complete-quarter definitions, respectively, after the baseline outcome-data and fixed-effect restrictions; within each definition, the baseline and JDR-adjusted estimates use identical observations and the original matched weights. Standard errors in parentheses are two-way clustered by county and matched comparison group. $*$, $**$, and $***$ denote significance at the 10, 5, and 1 percent levels. Because JDR is measured after treatment and covers all private-sector firms in the county, the adjusted estimates are conditional comparisons rather than a causal decomposition. See Appendix C.1.5 for further details.

Table A13: CES pre-election intention and realized voting behavior

Panel A. Transition counts and weighted row shares			
Pre-election intention	Democratic vote	Republican vote	Neither major-party vote
Democratic intention: unique count	164,566	4,154	18,174
Weighted row share (%)	86.8	2.9	10.4
Republican intention: unique count	3,557	176,769	15,786
Weighted row share (%)	1.7	90.7	7.6
Other-party intention: unique count	2,829	3,627	7,897
Weighted row share (%)	20.7	33.3	46.0
Panel B. Pooled Post effects			
Pre-election intention	Democratic vote	Republican vote	Neither major-party vote
Democratic intention	-0.592	-0.281	0.873
	(1.240)	(0.491)	(1.143)
Republican intention	-0.301	-1.055	1.356
	(0.374)	(0.976)	(0.918)
Other-party intention	-7.387*	-1.358	8.746*
	(4.056)	(5.022)	(5.055)
Panel C. Pooled Pre effects			
Pre-election intention	Democratic vote	Republican vote	Neither major-party vote
Democratic intention	0.273	-0.313	0.040
	(1.290)	(0.521)	(1.194)
Republican intention	0.283	-2.137**	1.855*
	(0.387)	(0.997)	(0.958)
Other-party intention	-9.645**	2.255	7.390
	(4.377)	(5.366)	(5.242)

Note: Panel A cross-tabulates office-specific pre-election intention and realized voting behavior within the same election in the matched chronological CES sample. Counts are unique respondent--office--year observations after de-duplicating the stack-expanded sample; the table contains 397,359 such observations and 768,912 stack-expanded rows. Weighted row shares combine the original matched-comparison weights with CES survey weights normalized within each stack--county--office--election-year cell and sum to 100 within intention. The sample requires a classified Democratic, Republican, or unambiguous other-party intention and an internally consistent realized outcome; within each intention group, the three outcomes use identical observations. Neither-major-party vote comprises reported nonvoters, unambiguous other-party voters, and explicit office undervoters. Panels B and C report pooled Post and Pre coefficients, respectively, from separate linear probability models for the realized-outcome indicator shown in each column, estimated within the intention group shown in each row. Post combines $h^C = 1, 2, 3, 4$, Pre combines $h^C = -3, -2, -1$, and $h^C = 0$ is omitted. Regressions include the baseline county controls and respondent age, age squared, and sex; absorb county-by-comparison, election-index, and state-by-office-by-year fixed effects; and use the same weights as Panel A. Because same-election intentions need not precede WARN exposure at positive h^C, Panel B decomposes realized behavior conditional on reported intention and is not interpreted as a causal effect on transitions from a predetermined intention. Panels B and C entries are percentage points. Standard errors in parentheses are two-way clustered by county and matched comparison group. $*$, $**$, and $***$ denote significance at the 10, 5, and 1 percent levels. See Appendix D.2 for further details.

Table A14: CES vote choice among voters

Panel A. Complete Sample	Base	+ demographics	+ party ID	+ intent	Panel FE
WARN exposure	0.029*** (0.009)	0.023*** (0.008)	0.003 (0.004)	0.002 (0.002)	-0.006 (0.010)
Observations	806,785	686,715	683,817	557,783	58,297
Panel B. Common Sample	Base	+ demographics	+ party ID	+ intent	
WARN exposure	0.028*** (0.011)	0.018** (0.008)	0.004 (0.004)	0.002 (0.002)	
Observations	557,783	557,783	557,783	557,783	

Note: The dependent variable equals one for a Democratic vote and zero for a Republican vote among major-party voters. WARN exposure is an indicator for a qualifying event in the respondent's county during the election-year window. Panel A uses all observations available to each specification. The repeated-cross-section columns include state-by-year and office-by-year fixed effects and then sequentially add respondent demographics, party identification, and pre-election vote intention. The final column pools the 2010--2014 and 2020--2024 CES panels and includes respondent and office-by-year fixed effects. The repeated-cross-section specifications use between 557,783 and 806,785 observations; the respondent-fixed-effect specification uses 58,297. Panel B re-estimates the first four specifications on the common sample with nonmissing information for all included variables, holding the sample fixed at 557,783 observations in every column. All specifications use CES survey weights, and standard errors in parentheses are clustered by county. $*$, $**$, and $***$ denote significance at the 10, 5, and 1 percent levels. See Appendix D.2 for further details.

Fig. A1: Qualifying WARN event and election-year window

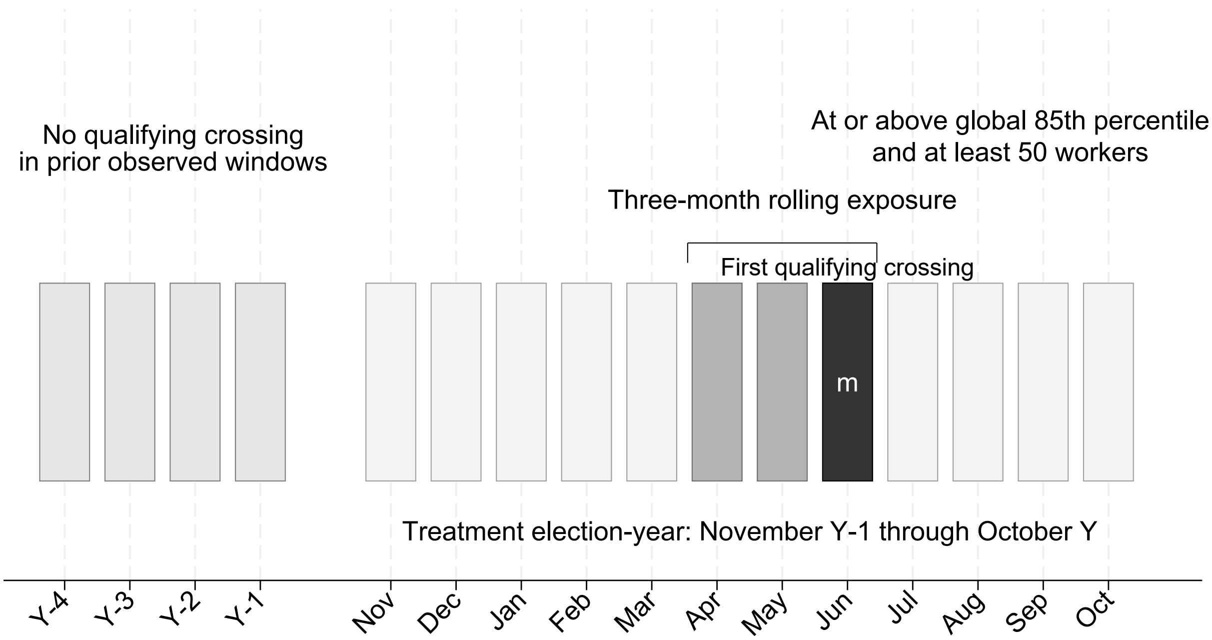

Note: The figure illustrates the treatment definition and the assignment of a qualifying county-month to an election-year window. Three-month WARN exposure sums exposure in the qualifying month and the two preceding months. A county-month qualifies when this complete rolling window is at or above the global 85th percentile among positive windows and contains at least 50 affected workers. The first qualifying month dates the event. The four preceding election-year windows must be fully observed and contain no other qualifying event, after which a county may be treated again following four clean windows. The displayed treatment month is illustrative; a qualifying event may occur in any month from November through the following October. The four darker bars on the left each represent an entire preceding election-year window, whereas the twelve bars on the right represent individual months in the treatment election-year. See Appendix B.1 for further details.

Fig. A2: Three-month WARN exposure around matched treatment events

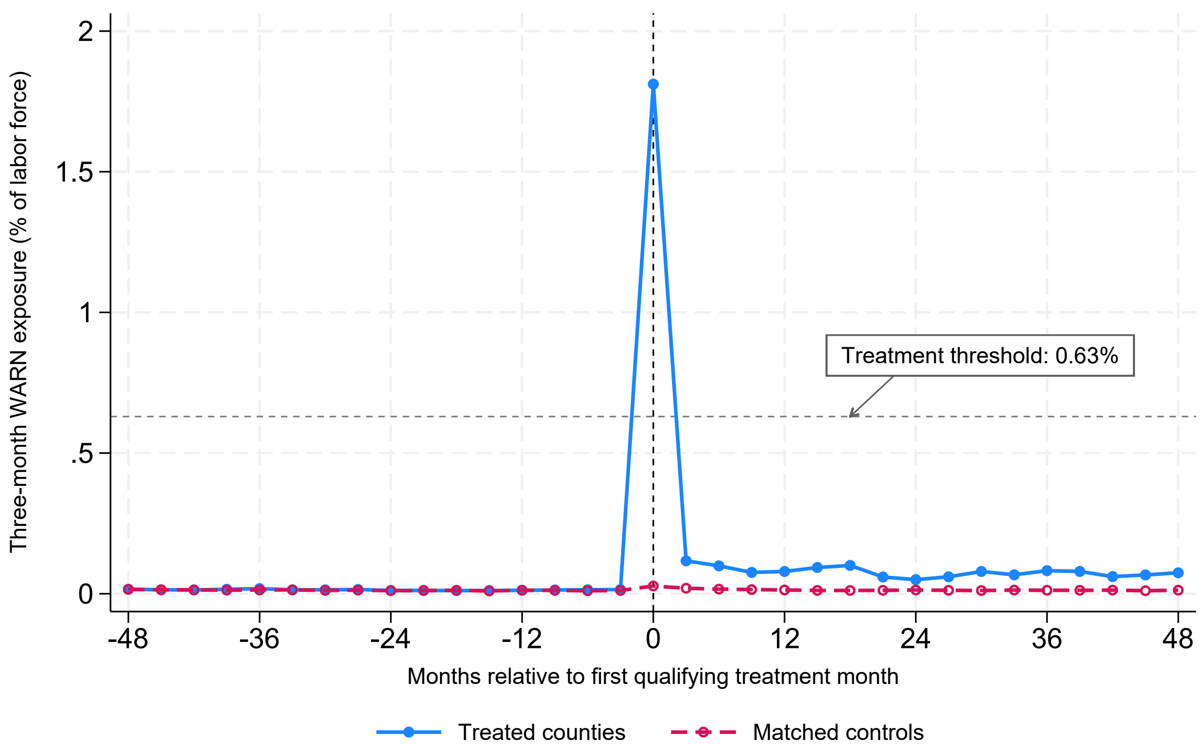

Note: The figure plots average three-month WARN exposure for treated counties and their matched controls from 48 months before through 48 months after the first qualifying treatment month. Exposure sums, over the three-month window, each month's affected-worker count divided by that month's county labor force and multiplied by 100. Month 0 is the first qualifying month, and the horizontal reference line marks the preferred 0.63-percent treatment threshold. Matched-control exposure is first averaged within each treated event using the inverse-distance control weights, after which treated events receive equal weight. Depending on distance from treatment and data availability, monthly support ranges from 1,312 to 1,508 of the 1,508 matched events. See Appendix B.2 for further details.

Fig. A3: Matched treated WARN events by calendar year

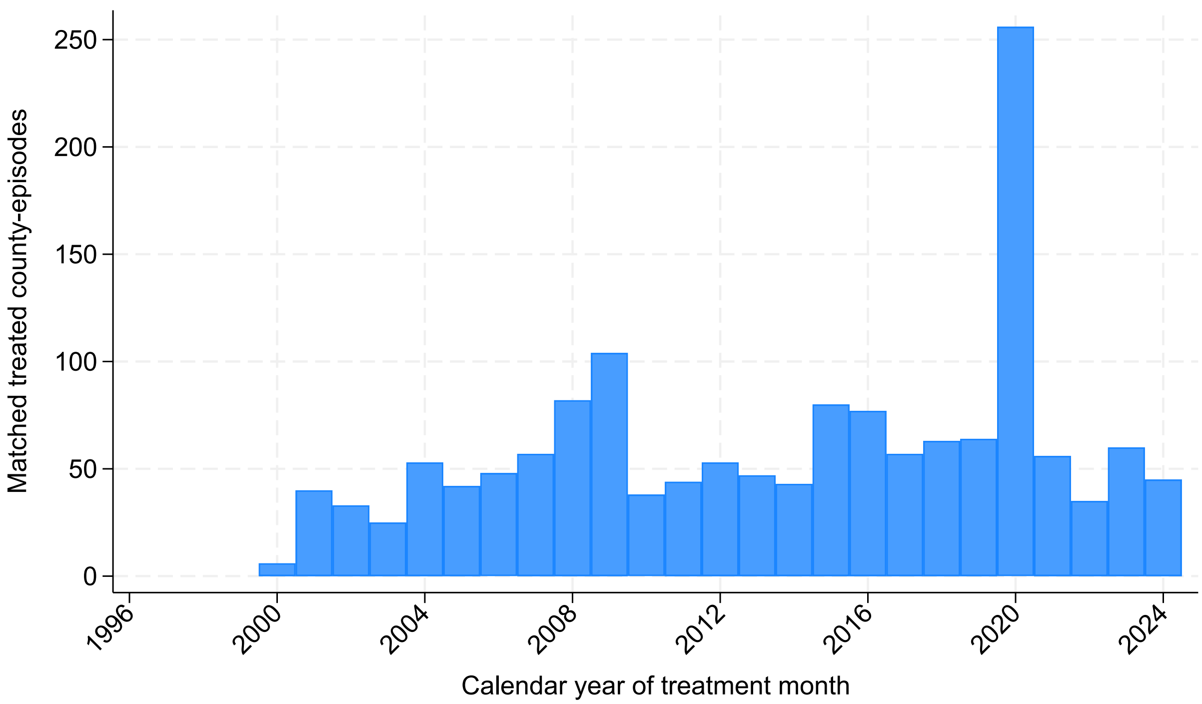

Note: The bars report the number of matched treated WARN events by the calendar year of the first qualifying treatment month. The figure covers all 1,508 events in the preferred chronological matched sample. The first matched events occur in 2000, once the design can establish the required four-year clean pre-event history, and the sample extends through 2024. See Appendix B.3.1 for further details.

Fig. A4: Chronological and office-specific election ordering

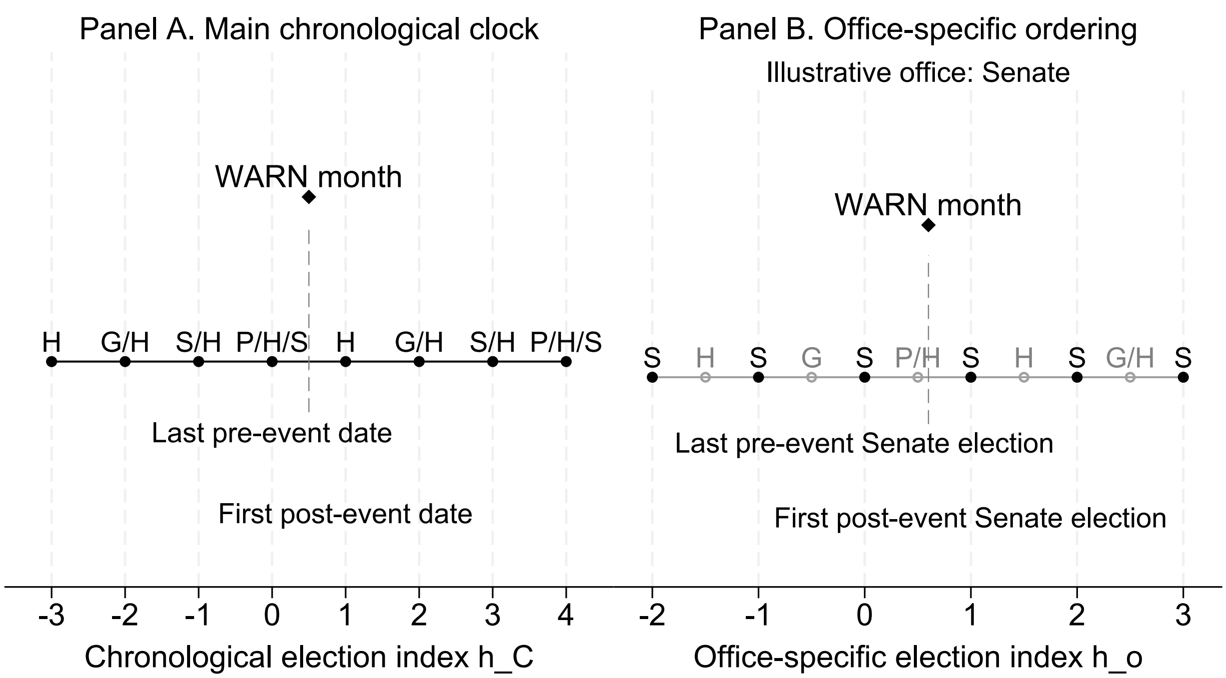

Note: Panel A illustrates the chronological election index, which orders included election dates in the sequence a county faces them. The final included date before the WARN event is $h^C = 0$, and the first date afterward is $h^C = 1$. When several offices are contested on the same date, they share the same chronological index and each office contributes a separate county-office outcome; the design does not select one office's turnout measure. Panel B illustrates the office-specific index using Senate elections. Only contests for the relevant office advance that office's index, while intervening contests for other offices leave it unchanged. The diagrams are illustrative and do not represent a particular comparison or sample count. See Appendix B.4 for further details.

Fig. A5: Pooled target-office design

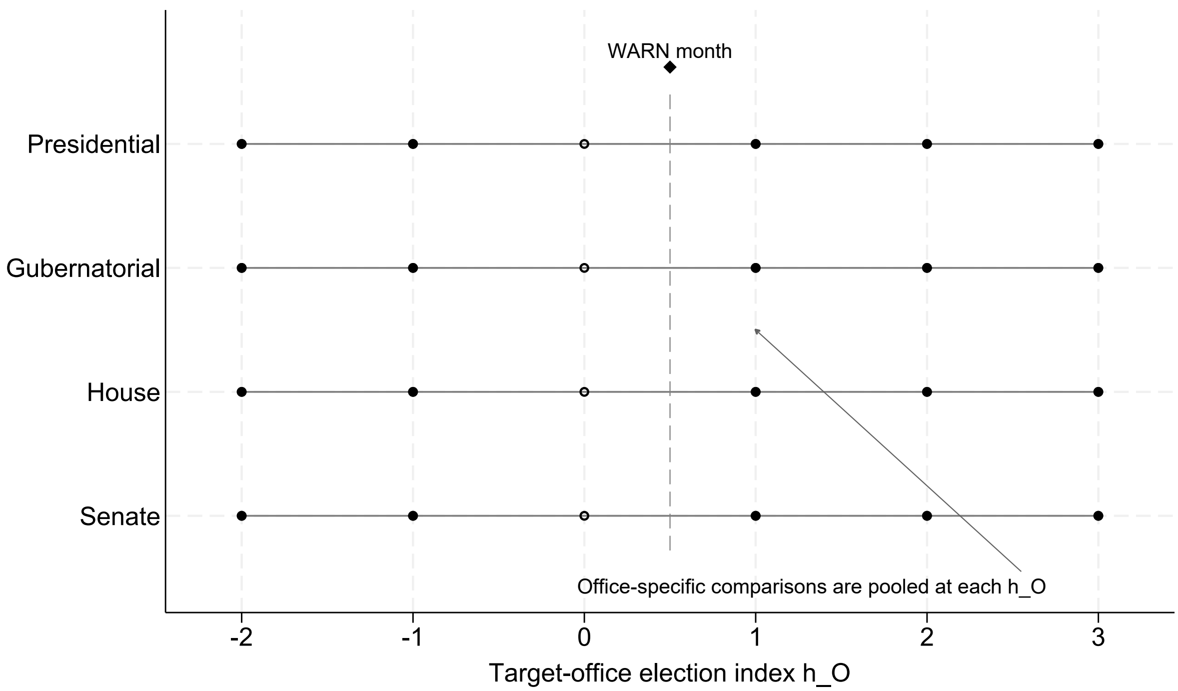

Note: The figure illustrates the target-office event-time design. Presidential, gubernatorial, House, and Senate comparisons are ordered separately around the last pre-WARN election of the same office, $h^O = 0$. The four office-specific designs are matched separately. At each office-specific index, the pooled target-office estimate combines the available comparisons across offices. The diagram is illustrative and does not represent a particular comparison or sample count. See Appendix C.1.1 for further details.

Fig. A6: Target-office effects on turnout and party-specific vote rates

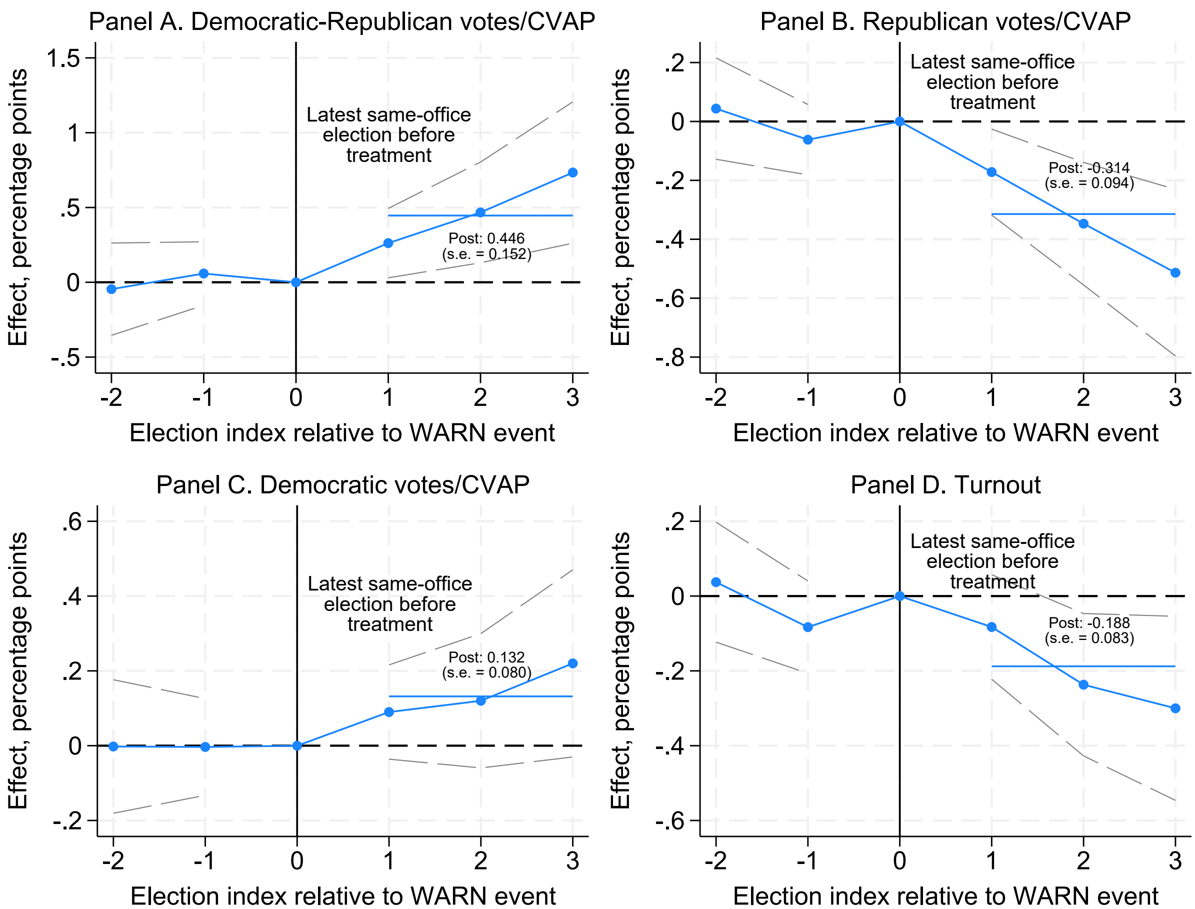

Note: Panels A-D report preferred event-study estimates under the pooled target-office design for Democratic-Republican votes per CVAP, Republican votes per CVAP, Democratic votes per CVAP, and turnout, respectively. The office-specific index sets $h^O = 0$ to the last election of the same office before the WARN event and $h^O = 1$ to the first election of that office afterward. The omitted index is $h^O = 0$, plotted at zero without a confidence interval. The estimates pool separately matched presidential, gubernatorial, House, and Senate comparisons and use 138,499 stack-expanded county-office-election observations. All outcomes are measured in percentage points. The specification includes preferred county controls, county-by-comparison and event-index fixed effects, state-by-office-by-year fixed effects, and inverse-distance weights. Dashed gray lines show 95 percent confidence intervals based on standard errors two-way clustered by county and matched comparison group. The text in each panel reports the pooled Post coefficient over $h^O = 1, 2, 3$ and its standard error. See Appendix C.1.1 for further details.

Fig. A7: Democratic county two-party-lead event-study estimates

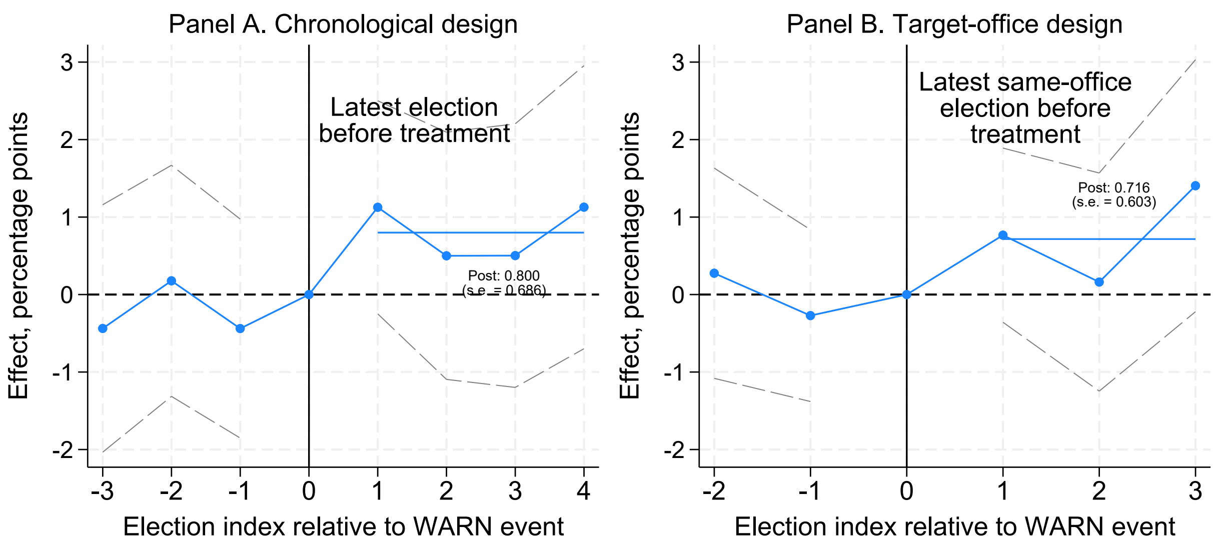

Note: The figure reports unrestricted event-study estimates for the Democratic county two-party-lead indicator. The indicator equals one when Democratic votes exceed Republican votes within the county and zero when Republican votes exceed Democratic votes; exact ties are excluded, and the indicator is defined only when both Democratic and Republican votes are positive. It is a local two-party-lead measure, not the winner of the electoral jurisdiction. Panels A and B report the chronological and target-office designs, respectively. The chronological estimates cover $h^C = -3, \ldots, 4$, while the target-office estimates cover $h^O = -2, \ldots, 3$. In each panel, $h = 0$ is omitted and plotted at zero. The regressions use 113,729 and 129,380 stack-expanded county-office-election observations, respectively, and retain the paper's baseline county controls, fixed effects, matched-comparison weights, and two-way clustering by county and matched comparison group. Dashed gray lines show 95 percent confidence intervals. The text in each panel reports the separately estimated pooled Post coefficient over $h^C = 1, 2, 3, 4$ in Panel A and $h^O = 1, 2, 3$ in Panel B and its standard error. See Appendix C.1.2.1 for further details.

Fig. A8: Effects of WARN events on electoral accountability

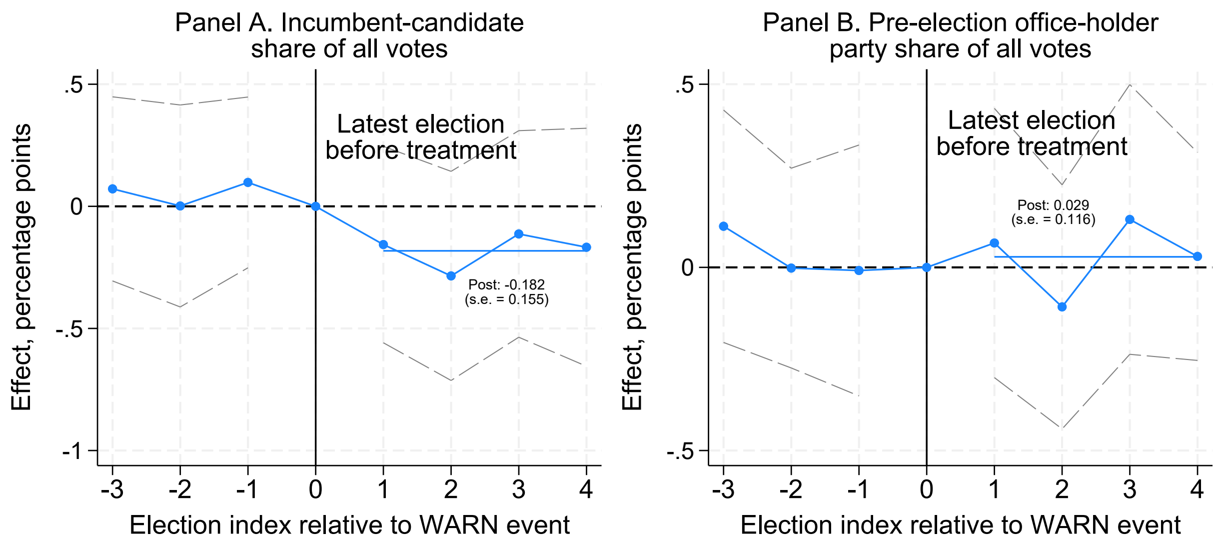

Note: The figure reports chronological event-study estimates for two accountability outcomes. Panel A uses the vote share received by a major-party incumbent candidate when such a candidate is running. Panel B uses the vote share received by the party holding the office before the election, whether or not the office holder is running. Both shares use all candidate votes in the denominator. The incumbent-candidate regression uses 53,659 stack-expanded county-office-election observations, and the office-holder-party regression uses 78,091. The omitted index is $h^C = 0$, plotted at zero without a confidence interval. Estimates use the preferred matched design, county controls, fixed effects, and inverse-distance weights. Dashed gray lines show 95 percent confidence intervals based on standard errors two-way clustered by county and matched comparison group. The text in both panels reports the pooled Post coefficient over $h^C = 1, 2, 3, 4$ and its standard error. See Appendix C.1.3 for further details.

Fig. A9: Observed vs. residualized WARN exposure

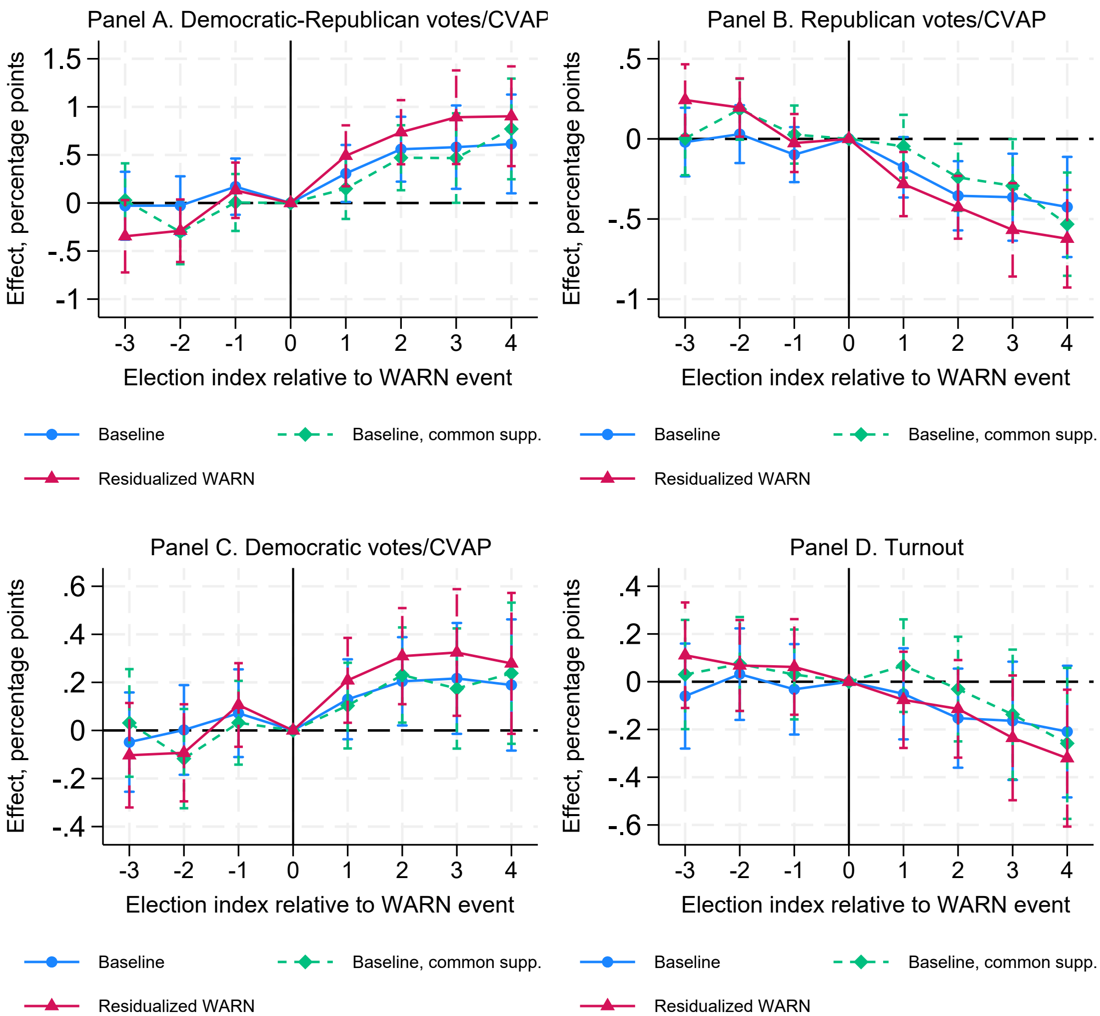

Note: The figure compares the full baseline actual-WARN design, an actual-WARN design rebuilt on the exact county-month support of the residual-exposure first stage, and the residualized-WARN design. Residualized exposure is observed three-month exposure minus the fitted value from a county-month regression on pre-event levels, 12-month changes, and annualized five-year trajectories in broad labor-market measures, together with state, year, and month fixed effects. A residualized county-month qualifies at the 85th percentile among positive residuals, a cutoff of 0.224 percentage points of county labor force, and must imply at least 50 affected-worker equivalents. For the common-support actual-WARN series, the 85th percentile is recomputed among positive observed exposure on that support, a cutoff of 0.481 percentage points, and a qualifying county-month must contain at least 50 observed affected workers. Each restricted-support design uses its own exposure measure to reconstruct treatment timing, clean histories, control eligibility, prior-exposure matching variables, and matches. The baseline, common-support actual-WARN, and residualized-WARN regressions use 122,698, 96,574, and 113,331 stack-expanded county-office-election observations and contain 1,508, 1,182, and 1,389 matched events, respectively. Panels A-D report Democratic-Republican votes per CVAP, Republican votes per CVAP, Democratic votes per CVAP, and turnout, respectively, all in percentage points. The omitted index is $h^C = 0$, plotted at zero. Vertical bars show 95 percent confidence intervals based on standard errors two-way clustered by county and matched comparison group. See Appendix C.1.4 for further details.

Fig. A10: Electoral effects of WARN events with and without adjustment for realized job destruction

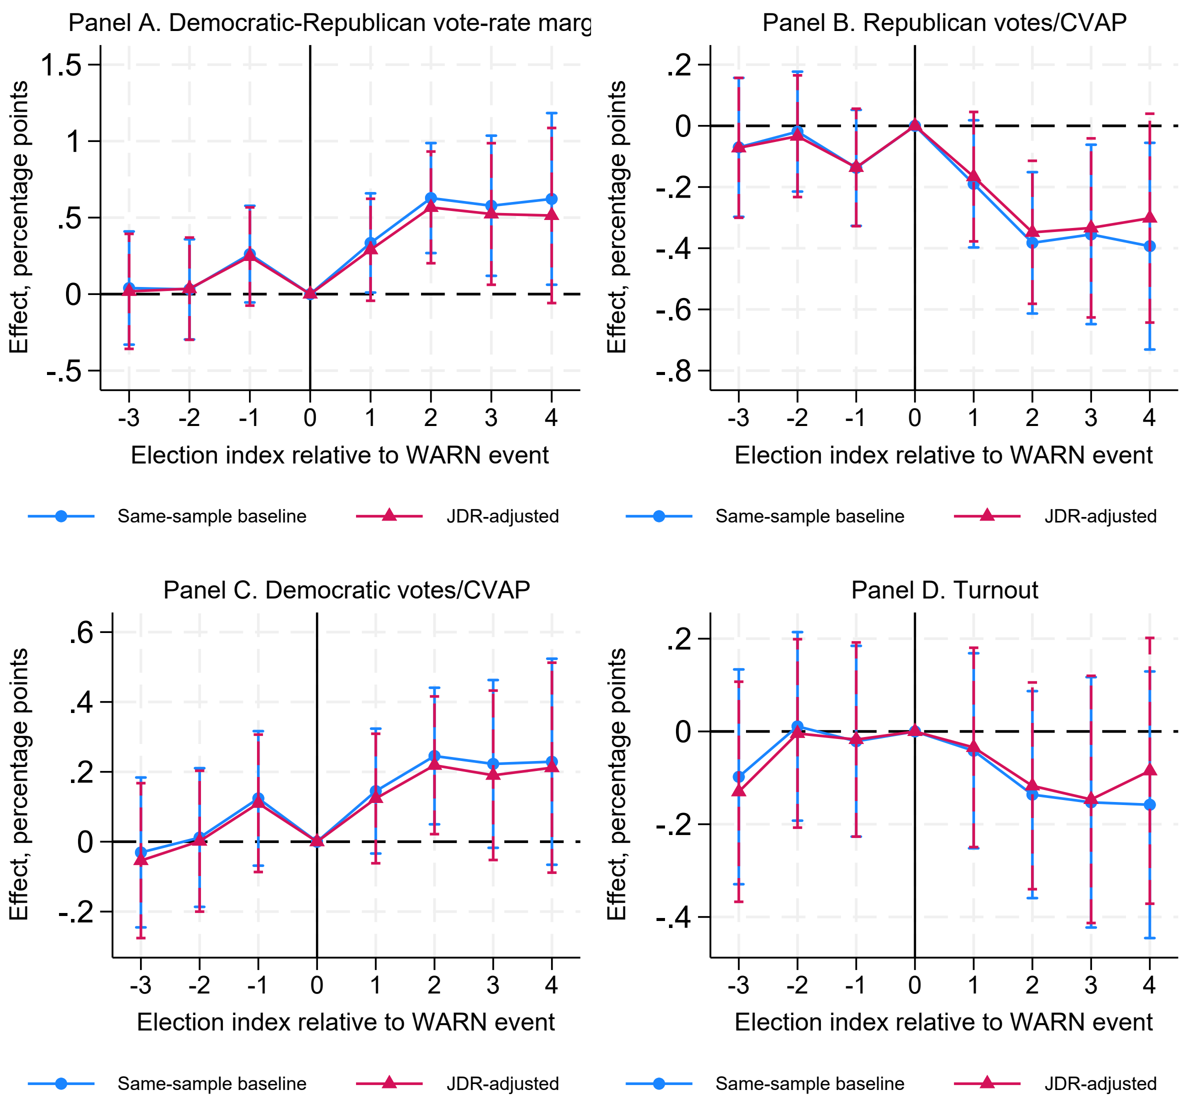

Note: The figure compares chronological event-study estimates from the same-sample baseline and JDR-adjusted specifications on the primary QWI-compatible sample. The change in the county firm job-destruction rate (JDR) is its average during the WARN-trigger quarter and the following quarter minus its average during the two complete quarters ending before the qualifying WARN window begins. The following quarter must end before the first post-WARN election. Panels A-D report the Democratic-Republican vote-rate margin, Republican votes per CVAP, Democratic votes per CVAP, and turnout, respectively, in percentage points. Both specifications use the same 102,884 stack-expanded observations in 1,252 matched comparisons and retain the baseline controls, fixed effects, matched weights, and clustering. The omitted index is $h^C = 0$, plotted at zero. Vertical bars show 95 percent confidence intervals based on standard errors two-way clustered by county and matched comparison group. Because JDR is measured after treatment and covers all private-sector firms in the county, the adjusted estimates are conditional comparisons rather than a causal decomposition. See Appendix C.1.5 for further details.

Fig. A11: Effects by treatment-to-election timing

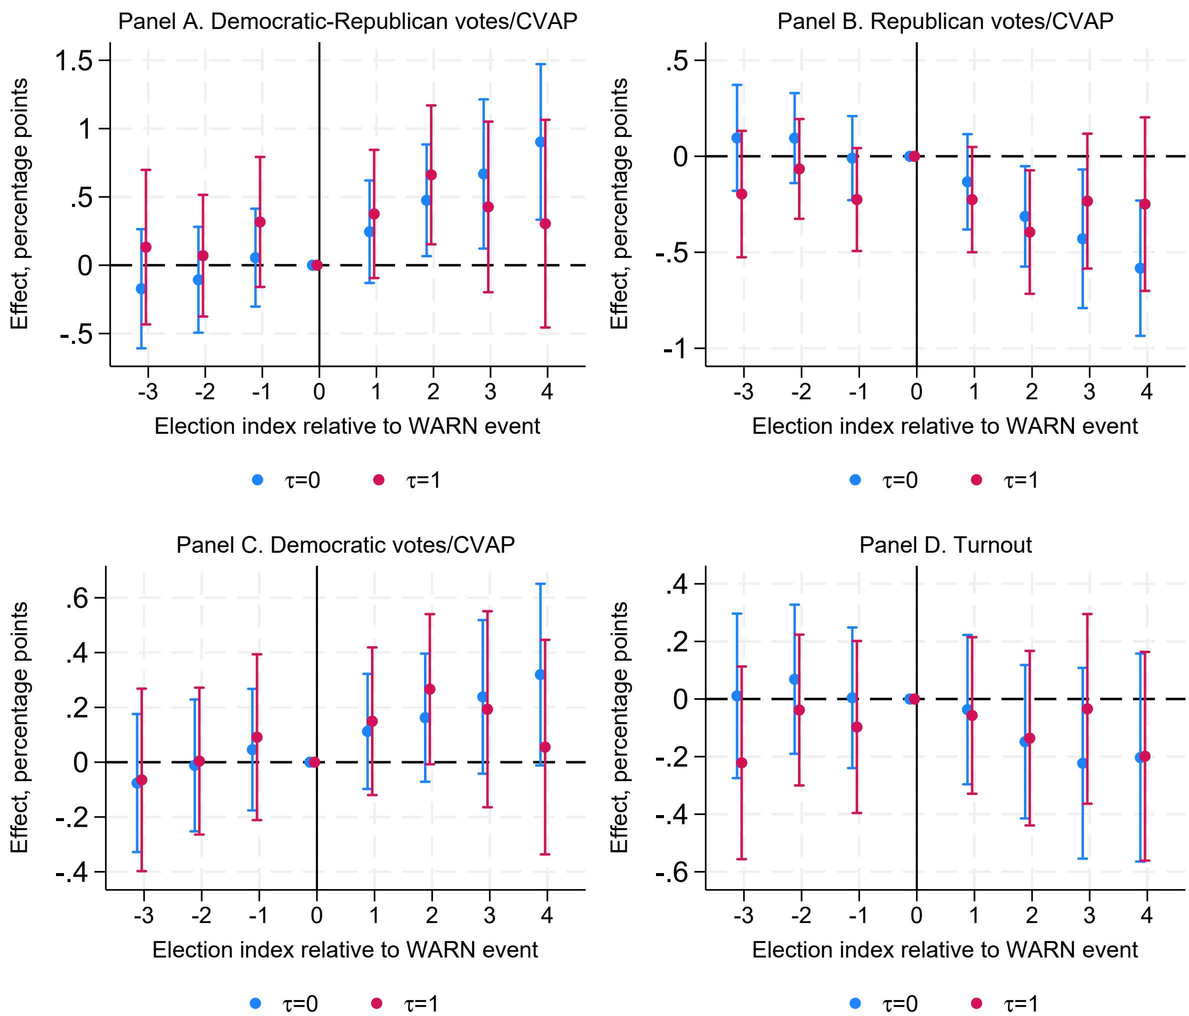

Note: The figure reports chronological event-study estimates separately by the distance from the first qualifying WARN month to the next included election. The $\tau = 0$ group contains events occurring 1-12 months before that election, while the $\tau = 1$ group contains events occurring 13-24 months beforehand. Panels A-D report Democratic-Republican votes per CVAP, Republican votes per CVAP, Democratic votes per CVAP, and turnout, respectively, all in percentage points. The $\tau = 0$ regressions use 71,508 stack-expanded county-office-election observations and the $\tau = 1$ regressions use 51,190. The omitted chronological index is $h^C = 0$, plotted at zero. Estimates retain the preferred controls, fixed effects, and matched-comparison weights. Standard errors are two-way clustered by county and matched comparison group. Vertical bars show 95 percent confidence intervals. See Appendix C.2.1 for further details.

Fig. A12: Effects of nonmanufacturing WARN events

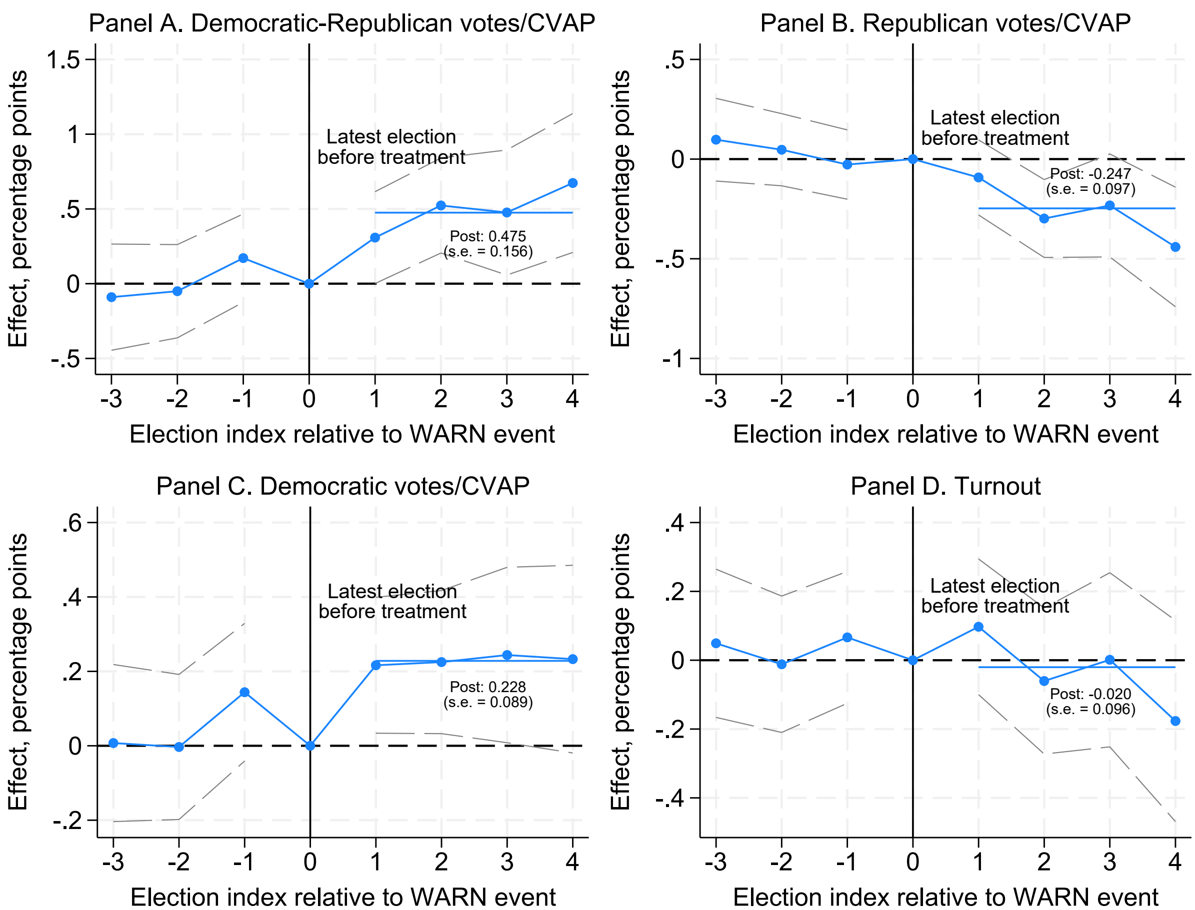

Note: The figure reports chronological event-study estimates after reconstructing treatment, control eligibility, and the matched comparisons using nonmanufacturing WARN notices. A county-month qualifies when three-month nonmanufacturing exposure is at or above the 85th percentile among positive windows, a cutoff of 0.39 percent of county labor force, and covers at least 50 workers. Of 1,448 match-ready events, 1,246 obtain at least one acceptable control. The regressions use 101,160 stack-expanded county-office-election observations for every displayed outcome. Panels A-D report Democratic-Republican votes per CVAP, Republican votes per CVAP, Democratic votes per CVAP, and turnout, respectively, all in percentage points. The omitted index is $h^C = 0$, plotted at zero without a confidence interval. Estimates otherwise retain the preferred chronological specification. Dashed gray lines show 95 percent confidence intervals based on standard errors two-way clustered by county and matched comparison group. The text in each panel reports the pooled Post coefficient over $h^C = 1, 2, 3, 4$ and its standard error. See Appendix C.2.2 for further details.

Fig. A13: Geographic scope: excluding nearby controls

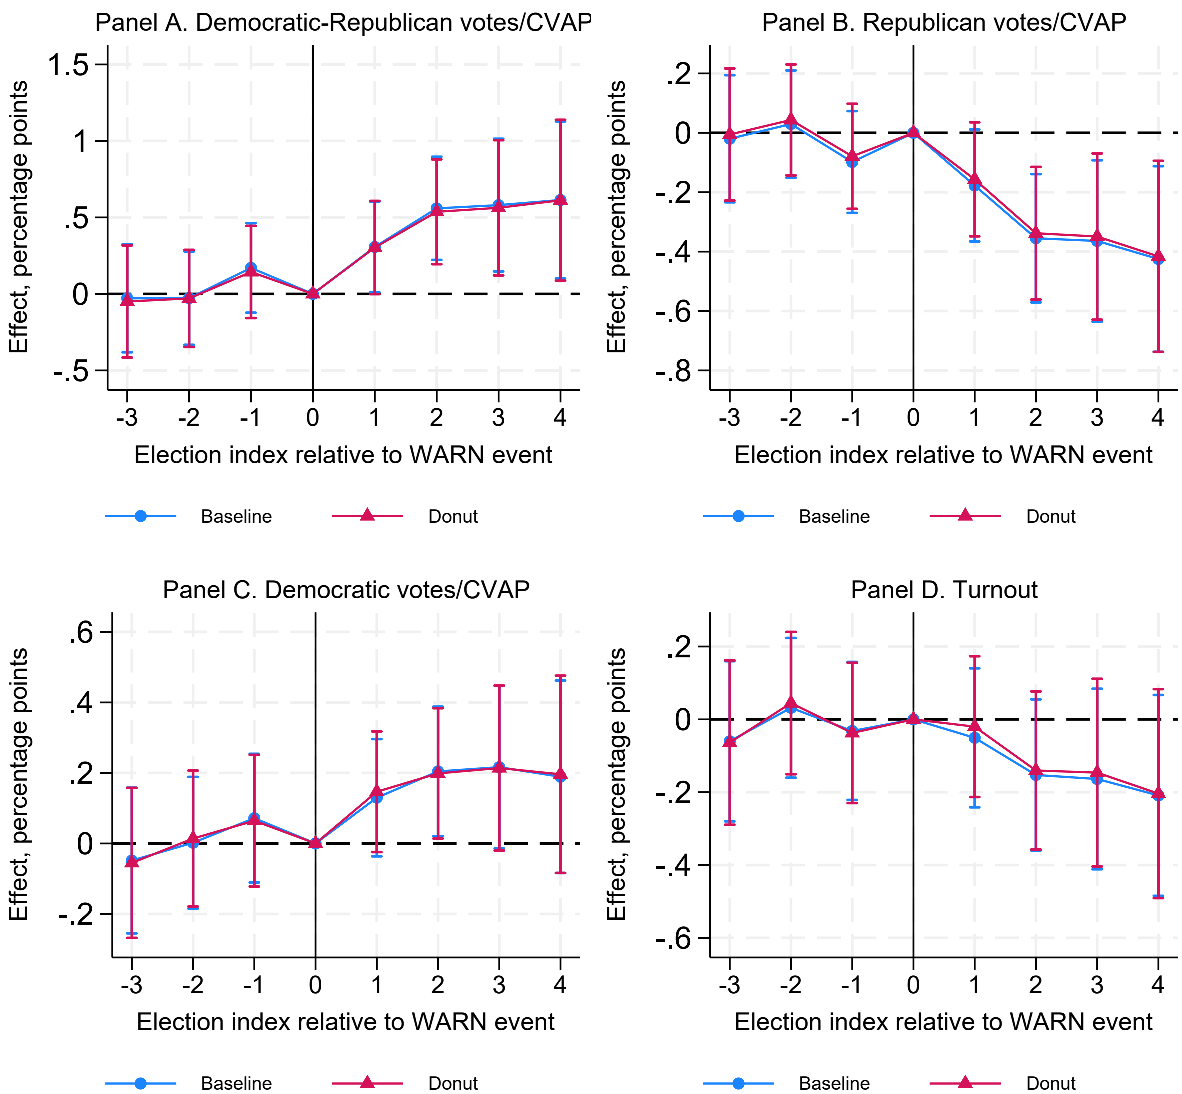

Note: The figure compares the preferred chronological estimates with a donut design that excludes every potential control in the treated county's state-by-commuting-zone unit before repeating the matching procedure. The preferred design matches 1,508 treated events and uses 122,698 stack-expanded county-office-election observations. The donut design matches 1,490 events and uses 120,460 observations. Panels A-D report Democratic-Republican votes per CVAP, Republican votes per CVAP, Democratic votes per CVAP, and turnout, respectively, all in percentage points. The omitted index is $h^C = 0$, plotted at zero. Both specifications use the same treatment definition, matching variables, calipers, preferred county controls, fixed effects, weighting rule, and use standard errors two-way clustered by county and matched comparison group. Vertical bars show 95 percent confidence intervals. See Appendix C.2.3 for further details.

Fig. A14: Sensitivity to alternative county control sets

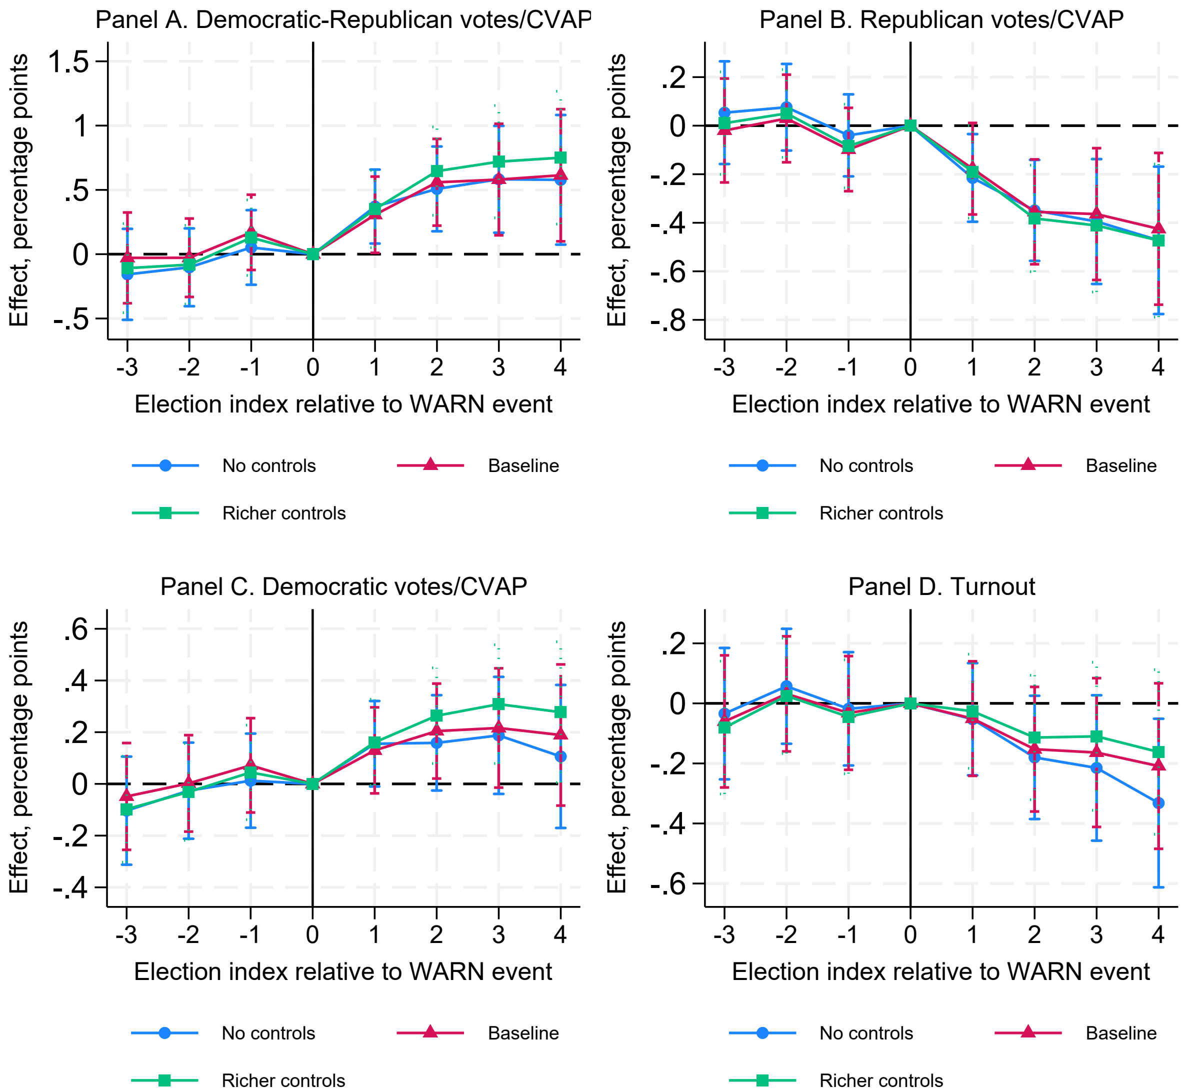

Note: The figure compares chronological event-study paths under three county-control sets. No controls omits time-varying county controls. Preferred includes the unemployment rate, labor-force size, per capita income, and cumulative COVID cases per 1,000 residents. Richer controls additionally include manufacturing employment share, log average weekly wages, and log population density. The matched sample, inverse-distance weights, and fixed effects are otherwise held fixed. Panels A-D report Democratic-Republican votes per CVAP, Republican votes per CVAP, Democratic votes per CVAP, and turnout, respectively, all in percentage points. The no-controls, preferred-controls, and richer-controls specifications use 125,113, 122,698, and 122,265 stack-expanded county-office-election observations, respectively. The omitted chronological index is $h^C = 0$, plotted at zero. Vertical bars show 95 percent confidence intervals based on standard errors two-way clustered by county and matched comparison group. See Appendix C.3.1 for further details.

Fig. A15: Long-run trajectory adjustments

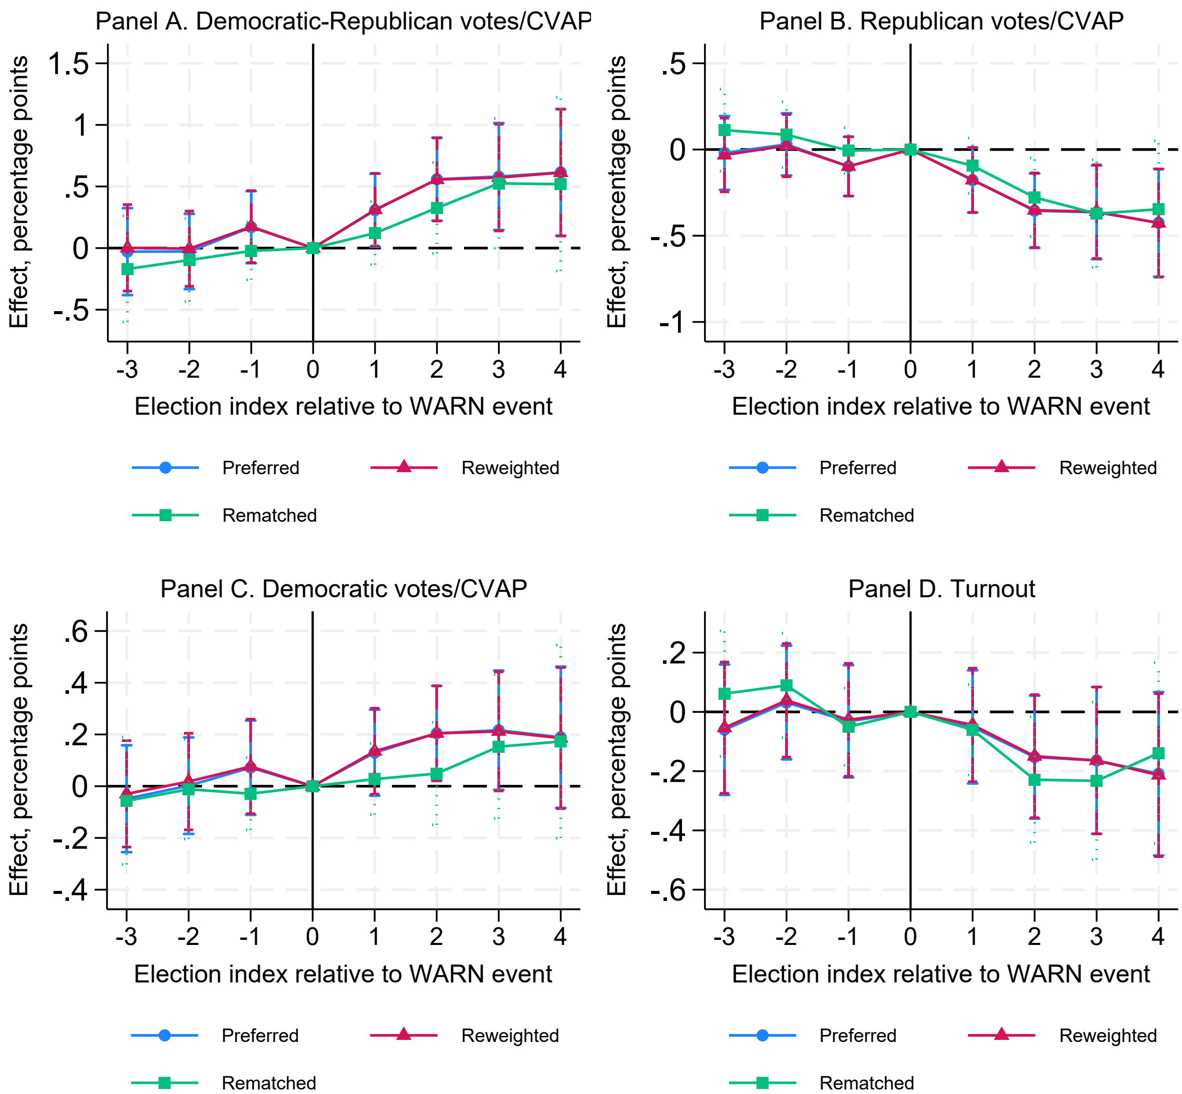

Note: The figure compares the preferred chronological estimates with two adjustments for longer pre-WARN economic, demographic, and sectoral trajectories. Reweighted retains the preferred matched controls but modifies their weights according to distance in the longer-run trajectory variables and renormalizes the control weights within each comparison. Rematched rebuilds the treated-control comparisons using those additional histories. Panels A-D report Democratic-Republican votes per CVAP, Republican votes per CVAP, Democratic votes per CVAP, and turnout, respectively, all in percentage points. Preferred and Reweighted each use 122,698 stack-expanded county-office-election observations; Rematched uses 137,660. The omitted index is $h^C = 0$, plotted at zero. Vertical bars show 95 percent confidence intervals based on standard errors two-way clustered by county and matched comparison group. See Appendix C.3.1 for further details.

Fig. A16: Adjustment for forecasted local deterioration

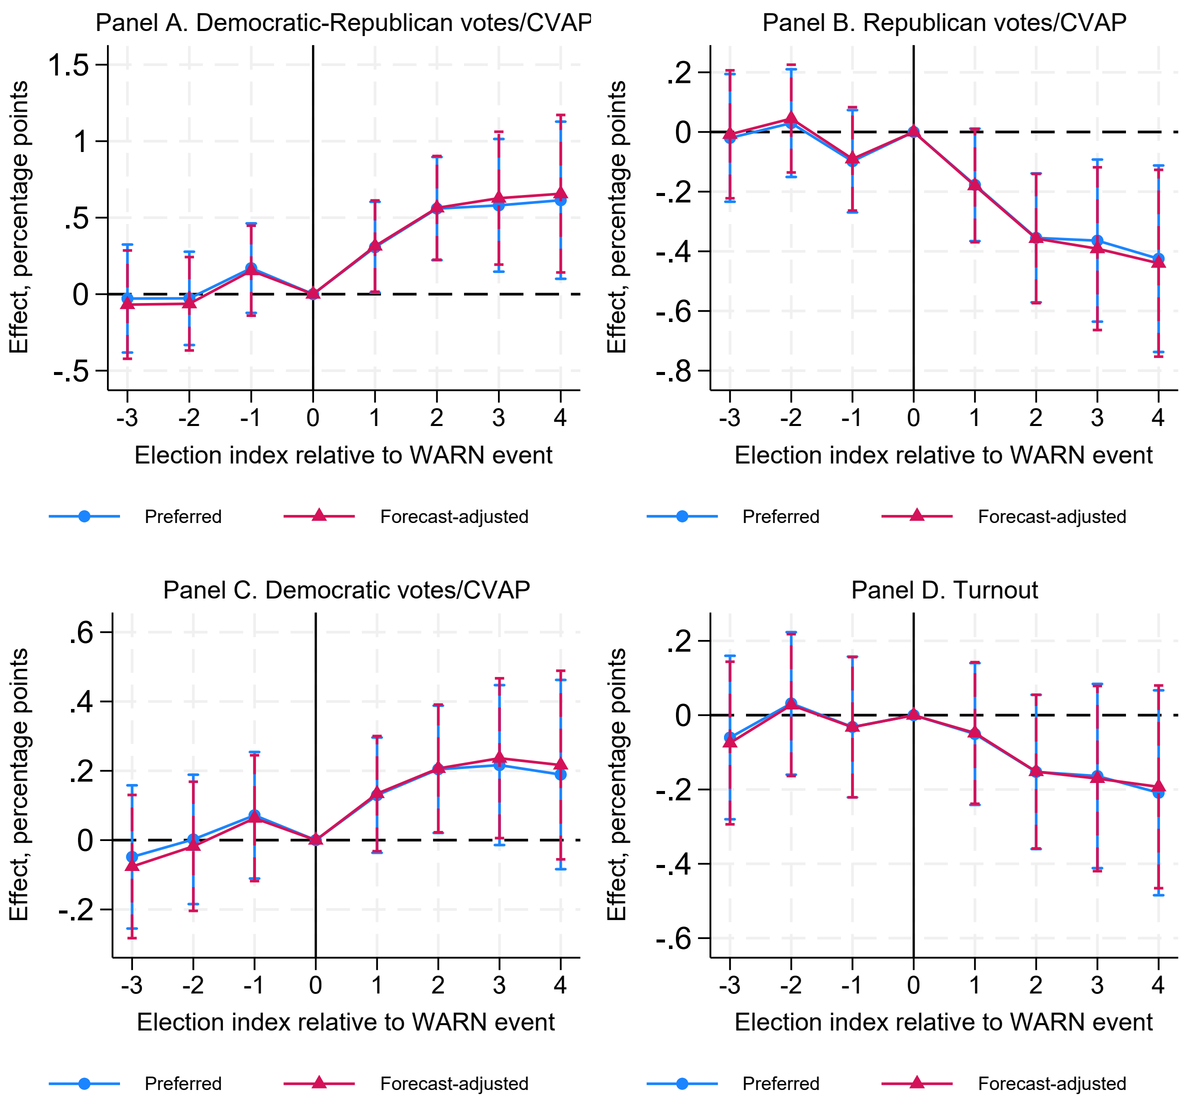

Note: The figure compares the preferred chronological specification with a forecast-adjusted specification. The forecast score is an equal-weight index of predetermined local deterioration measures, standardized and oriented so that higher values indicate worse predicted conditions. The adjusted specification interacts this score with every chronological election index while retaining the preferred matched sample, weights, county controls, and fixed effects. Panels A-D report Democratic-Republican votes per CVAP, Republican votes per CVAP, Democratic votes per CVAP, and turnout, respectively, all in percentage points. Both displayed specifications use 122,698 stack-expanded county-office-election observations. The omitted index is $h^C = 0$, plotted at zero. Vertical bars show 95 percent confidence intervals based on standard errors two-way clustered by county and matched comparison group. See Appendix C.3.2 for further details.

Fig. A17: Alternative treatment definitions

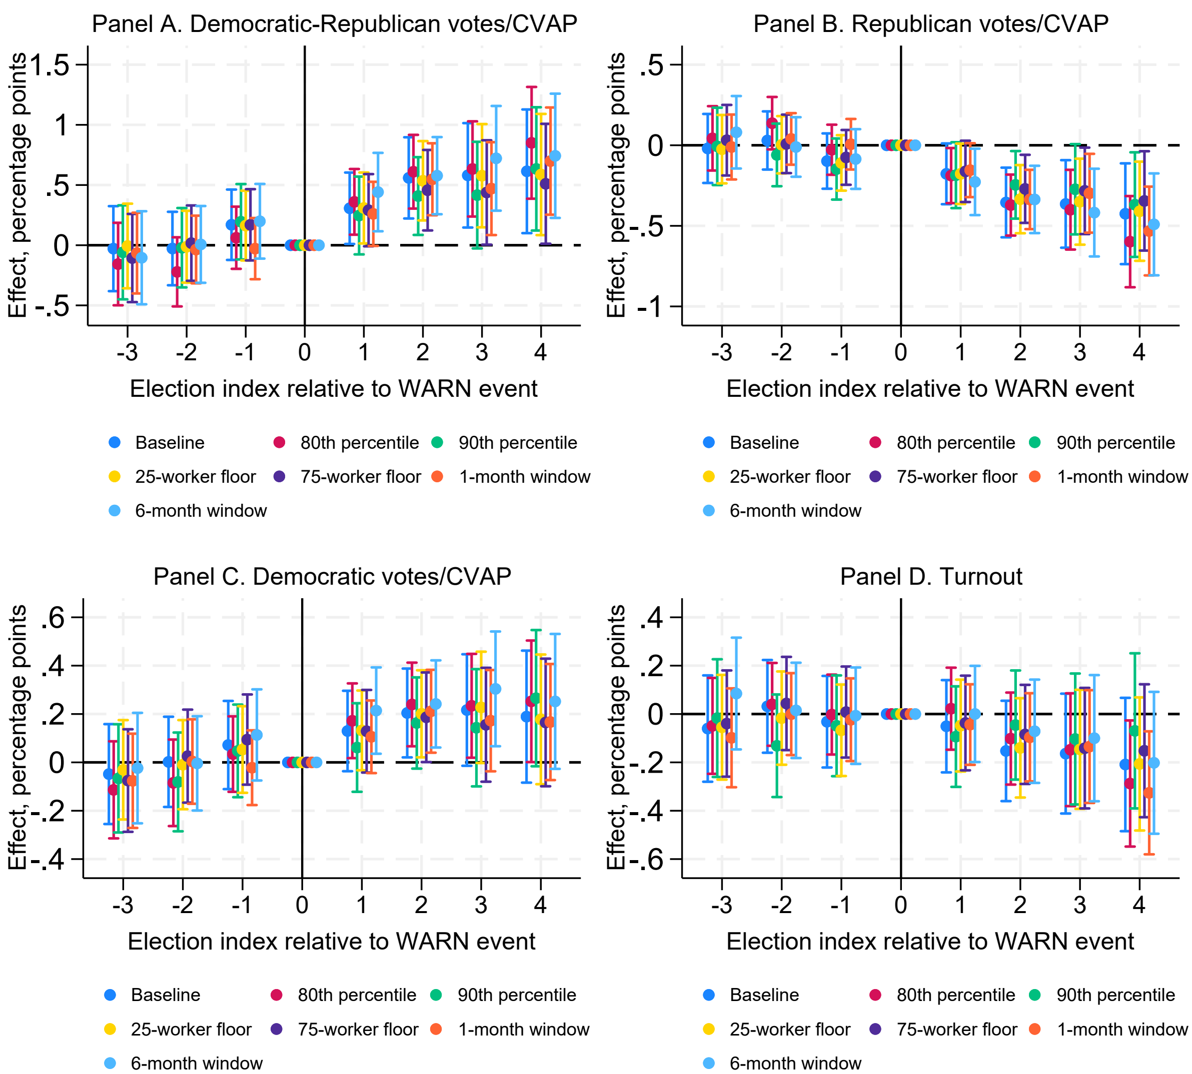

Note: The figure compares the preferred chronological design with alternatives that change one component of the treatment definition: the percentile cutoff for positive WARN exposure (80th or 90th rather than 85th), the minimum-worker floor (25 or 75 rather than 50), or the exposure window (one or six months rather than three). Each exercise retains the other baseline choices and reconstructs treatment, clean histories, control eligibility, and the matched comparisons. Panels A-D report Democratic-Republican votes per CVAP, Republican votes per CVAP, Democratic votes per CVAP, and turnout, respectively, all in percentage points. The omitted index is $h^C = 0$, plotted at zero. The preferred, 80th-percentile, 90th-percentile, 25-worker-floor, 75-worker-floor, one-month-window, and six-month-window designs use 122,698, 138,863, 97,889, 123,134, 118,802, 143,898, and 106,299 stack-expanded county-office-election observations, respectively. Vertical bars show 95 percent confidence intervals based on standard errors two-way clustered by county and matched comparison group. See Appendix C.4.1 for further details.

Fig. A18: Varying the number of matched controls

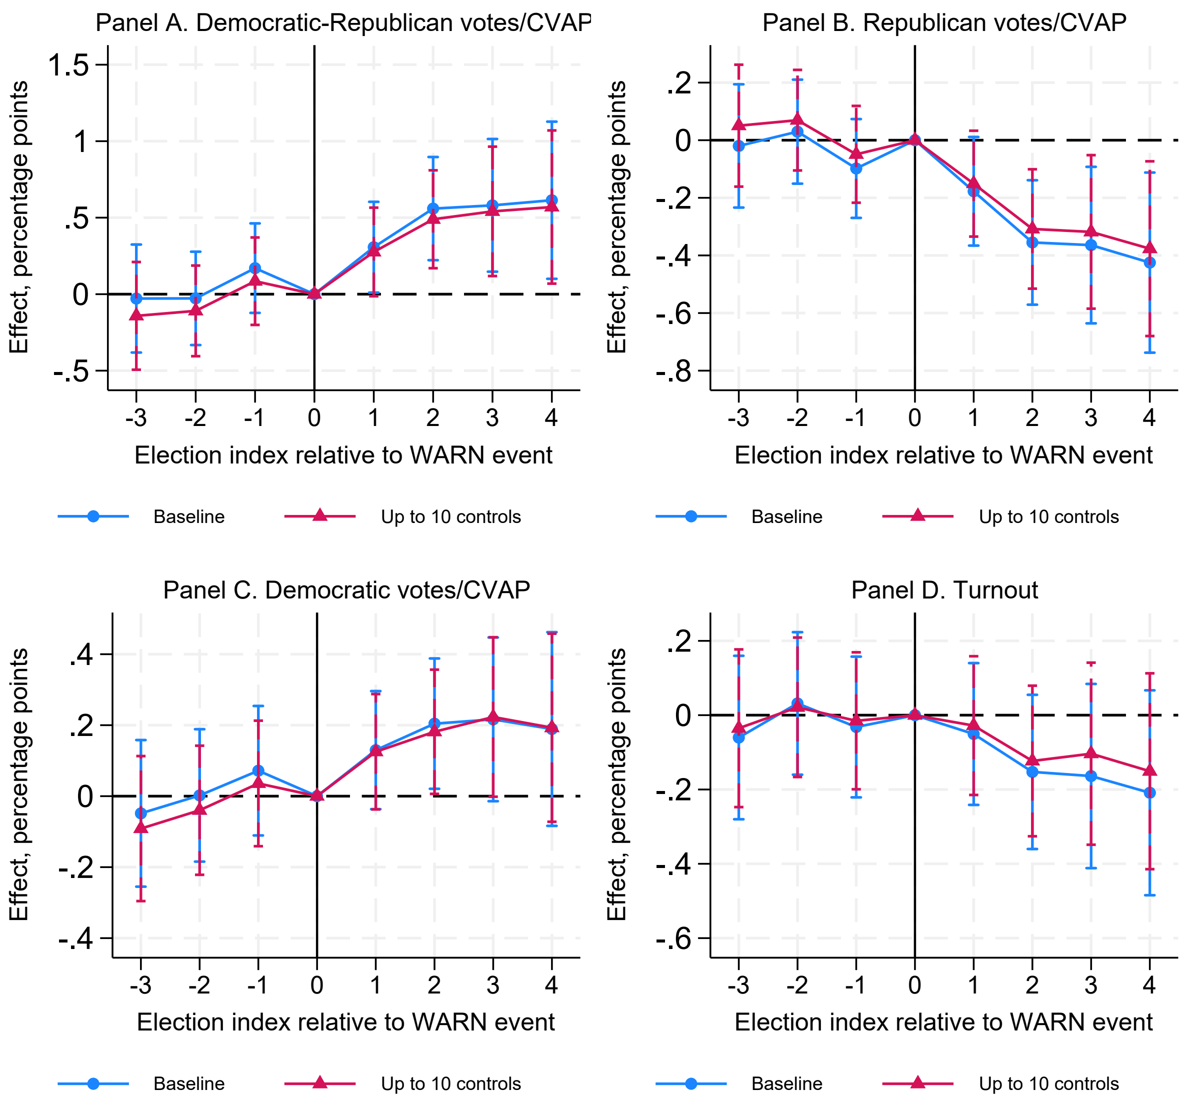

Note: The figure compares the preferred design, which retains up to four controls per treated event, with an otherwise identical specification that retains up to ten. Panels A-D report Democratic-Republican votes per CVAP, Republican votes per CVAP, Democratic votes per CVAP, and turnout, respectively, all in percentage points. The preferred regressions use 122,698 stack-expanded county-office-election observations, while the up-to-ten-control regressions use 228,052 due to the expanded control set. The larger count reflects the additional control rows retained within matched comparisons. The omitted index is $h^C = 0$, plotted at zero. Both designs use inverse-distance weights normalized within comparison and the same preferred controls and fixed effects. Vertical bars show 95 percent confidence intervals based on standard errors two-way clustered by county and matched comparison group. See Appendix C.4.2 for further details.

Fig. A19: Alternative fixed effects, weighting, and log-odds transformation

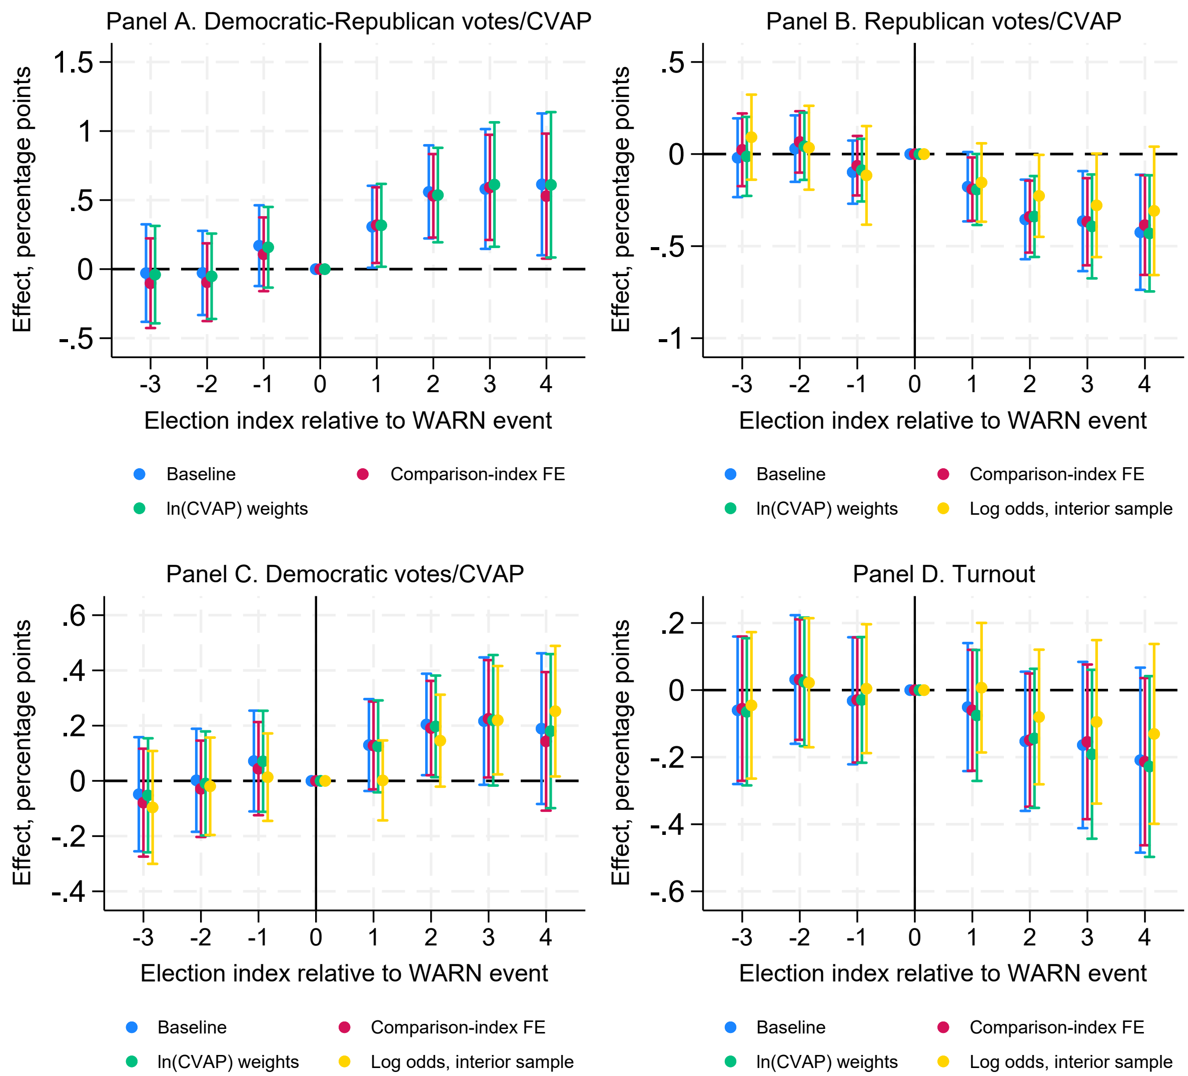

Note: The figure compares the preferred chronological specification with alternatives that vary fixed effects, weighting, or functional form. Comparison-index FE replaces the baseline event-index and state-by-office-by-year fixed effects with comparison-by-election-index-by-office fixed effects, restricting identifying variation to treated-control differences within each matched comparison, office, and election index. The $\ln(\text{CVAP})$-weighted specification scales each comparison by the treated county's log CVAP at the last pre-WARN election. The log-odds specification replaces each included outcome $V/CVAP$ with $\ln[V/(CVAP - V)]$, restricts the sample by outcome to $0 < V < CVAP$, and reports weighted average derivatives on the original percentage-point scale. The log-odds series is reported in Panels B-D for Republican voting, Democratic voting, and turnout; the signed Democratic-Republican margin in Panel A has no corresponding ordinary log-odds transformation. Panels A-D report Democratic-Republican votes per CVAP, Republican votes per CVAP, Democratic votes per CVAP, and turnout, respectively, all in percentage points. The regressions use between 117,918 and 122,698 stack-expanded county-office-election observations across specifications. The omitted index is $h^C = 0$, plotted at zero. Vertical bars show 95 percent confidence intervals based on standard errors two-way clustered by county and matched comparison group; those for the log-odds estimates are computed by the delta method. See Appendices C.4.3 and C.4.5 for further details.

Fig. A20: Excluding treatment election-year 2020

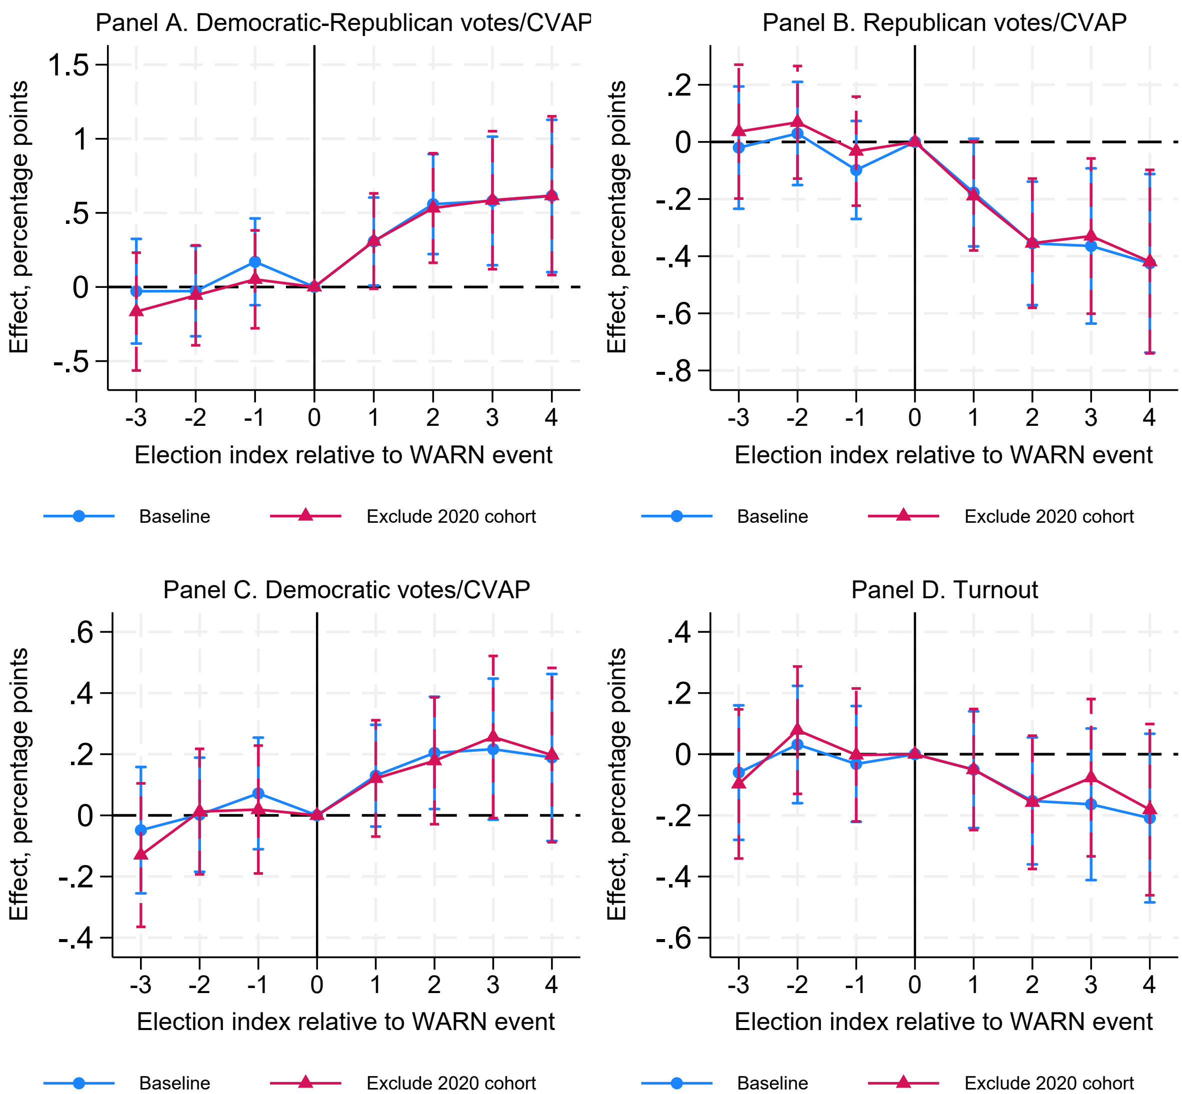

Note: The figure compares the preferred chronological estimates with a specification that removes treated events assigned to the November 2019-October 2020 treatment election-year window. Panels A-D report Democratic-Republican votes per CVAP, Republican votes per CVAP, Democratic votes per CVAP, and turnout, respectively, all in percentage points. The restricted specification retains 1,248 matched treated events. The preferred regressions use 122,698 stack-expanded county-office-election observations, and the restricted regressions use 102,375. The omitted index is $h^C = 0$, plotted at zero. Both specifications retain the preferred controls, fixed effects, and inverse-distance weights. Standard errors are two-way clustered by county and matched comparison group. Vertical bars show 95 percent confidence intervals. See Appendix C.4.4 for further details.

Fig. A21: Controlling for demographic composition and migration

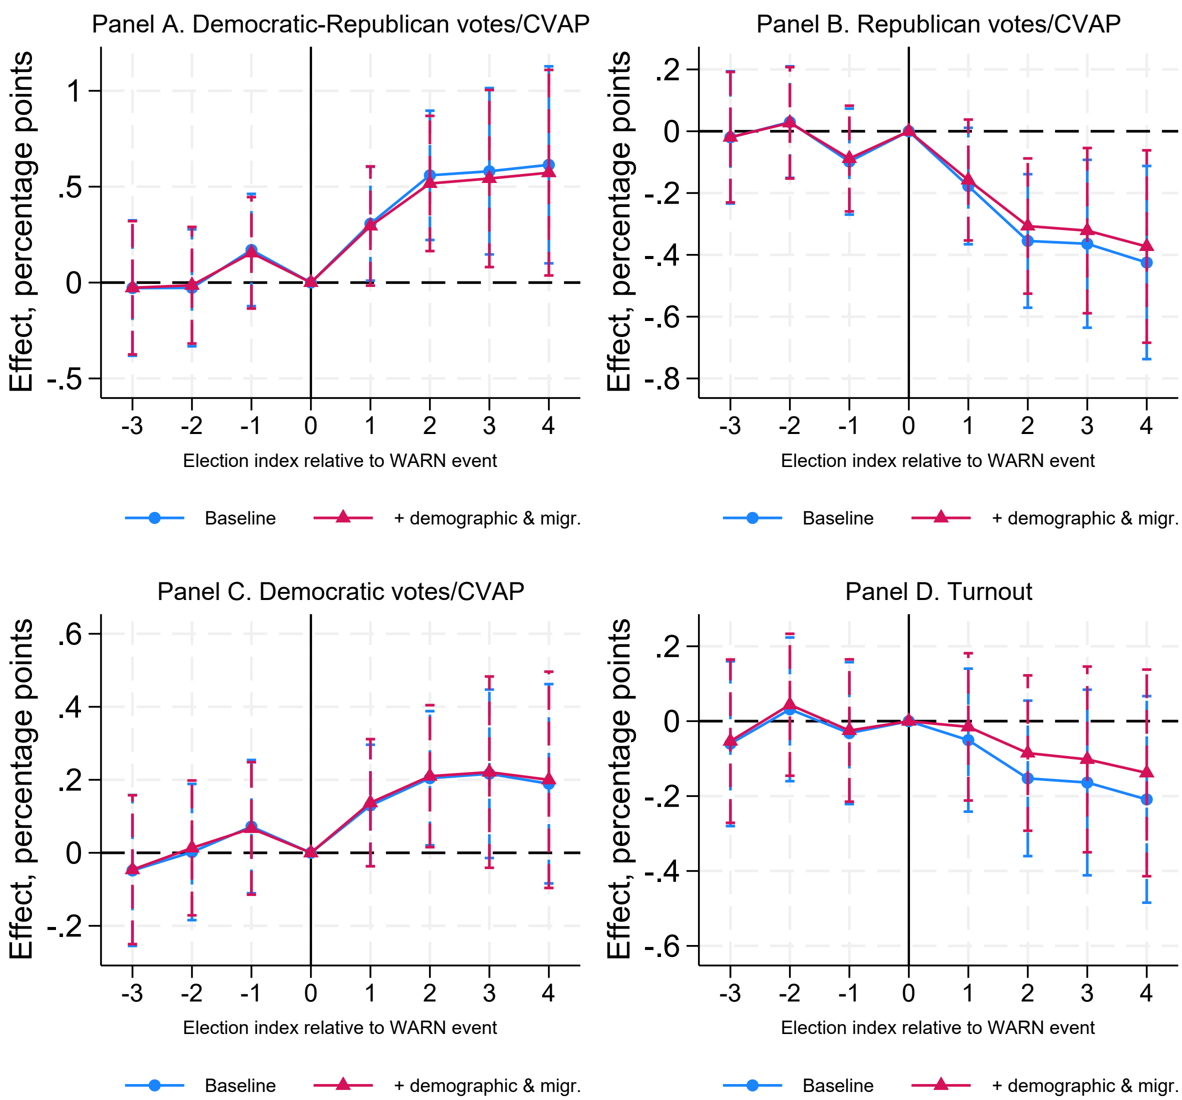

Note: The figure compares the preferred chronological specification with a specification that additionally controls for age composition, the nonwhite or Hispanic population share, the CVAP-to-voting-age-population ratio, and net and gross migration. Panels A-D report Democratic-Republican votes per CVAP, Republican votes per CVAP, Democratic votes per CVAP, and turnout, respectively, all in percentage points. The preferred regressions use 122,698 stack-expanded county-office-election observations, while the demographic-and-migration specification uses 114,969 because the additional variables are not available for every observation. Both specifications retain the preferred matched comparisons, inverse-distance weights, remaining county controls, and fixed effects. The omitted index is $h^C = 0$, plotted at zero. Vertical bars show 95 percent confidence intervals based on standard errors two-way clustered by county and matched comparison group. See Appendix D.1 for further details.

Fig. A22: CES evidence on vote switching and turnout composition

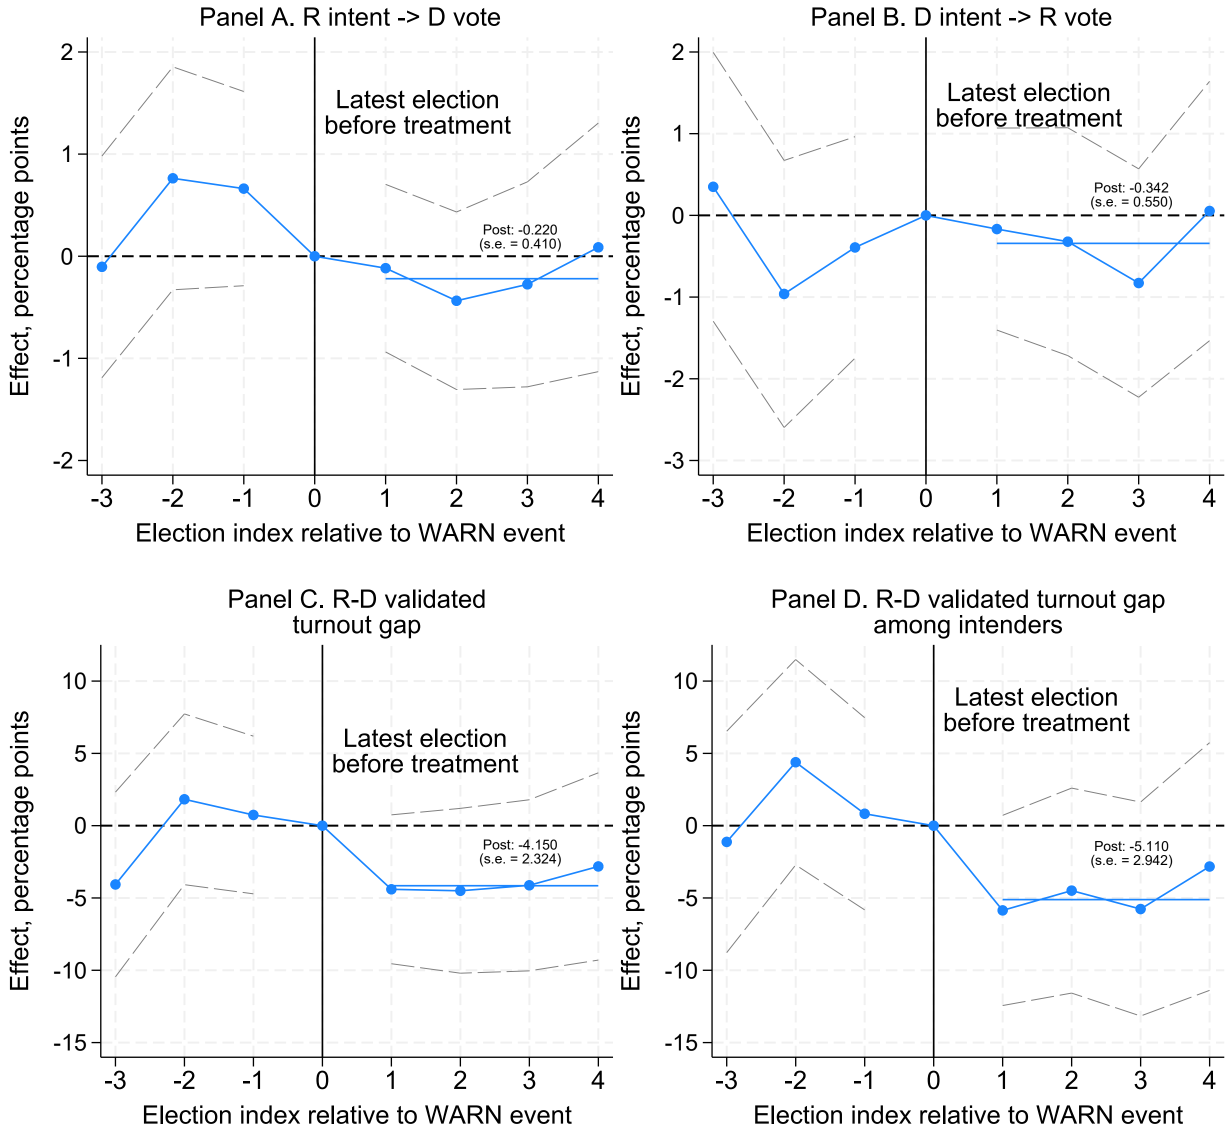

Note: The figure reports respondent-level chronological event-study estimates using the cumulative CES for 2006-2024. Panels A and B measure switching relative to office-specific pre-election intentions among respondents who later report voting in the election. The outcome equals one for a vote for the opposing major party and zero for a same-party vote, an other-candidate vote, or an explicit decision not to vote in that office; respondents with an unavailable office choice are excluded. Panels C and D report the Republican-minus-Democratic difference in validated-turnout responses among major-party identifiers and leaners, first overall and then among respondents who expressed an affirmative pre-election intention to vote. County controls and fixed effects in the turnout-gap regressions are allowed to differ by partisan group. All four regressions include respondent age, age squared, and sex as common main effects. The four regressions use 368,146, 314,945, 1,049,470, and 721,123 stack-expanded observations at the respondent-office-comparison level, respectively. A respondent may enter more than one matched comparison or contribute more than one office outcome. Within each stack-county-office-year cell, CES survey weights are normalized to sum to one and multiplied by the matched-comparison weight. The omitted index is $h^C = 0$, plotted at zero. Vertical bars show 95 percent confidence intervals based on standard errors two-way clustered by county and matched comparison group. The text in each panel reports the pooled Post coefficient over $h^C = 1, 2, 3, 4$ and its standard error. See Appendix D.2 for further details.

Fig. A23: Transfer payments after WARN events and voting adjustment

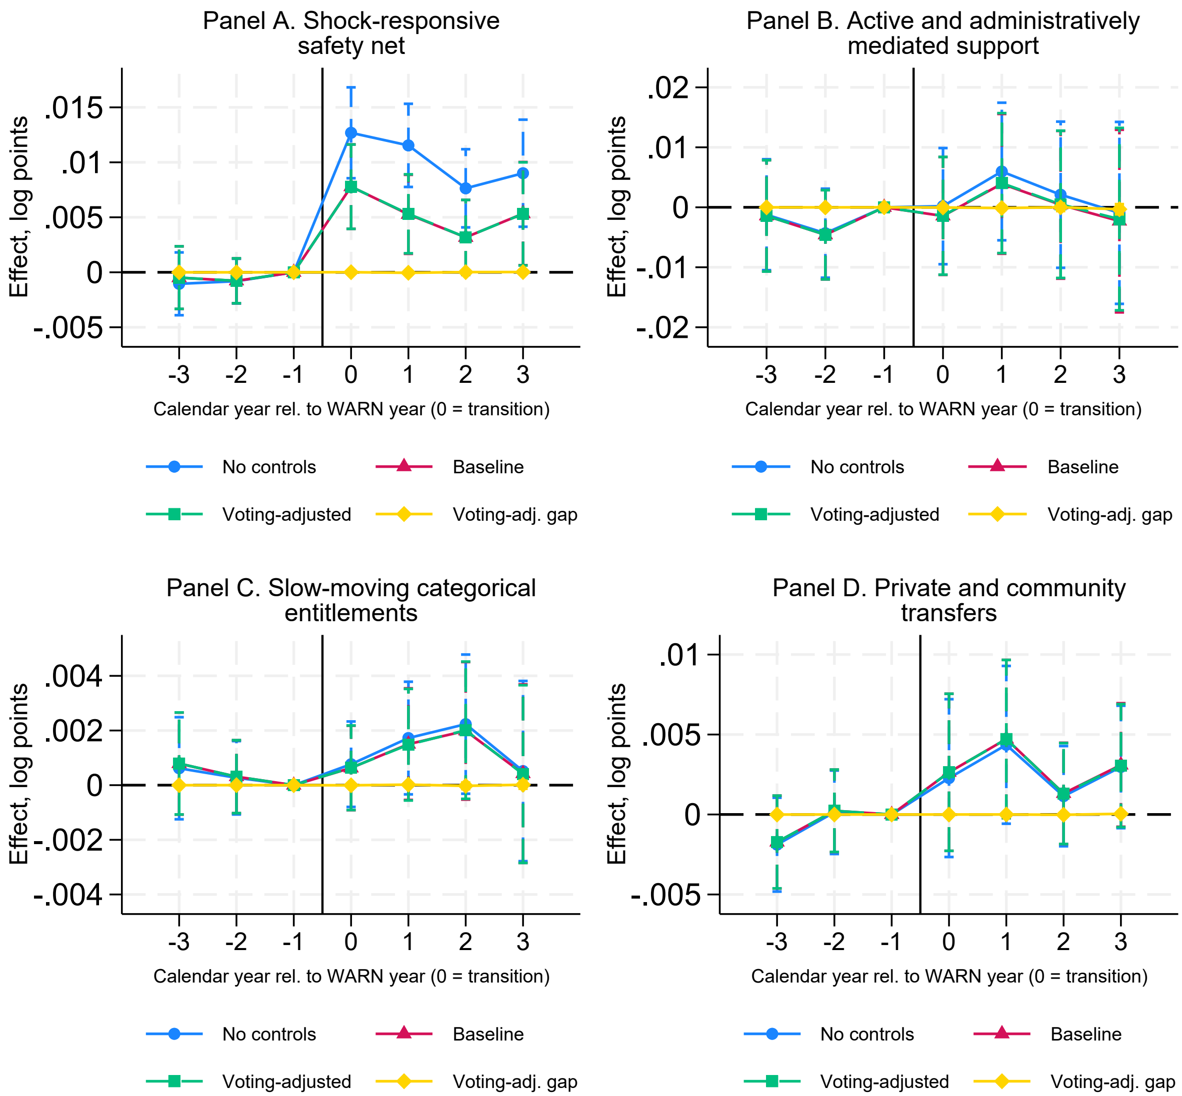

Note: The figure reports annual event-study estimates for four groups of county transfer payments: shock-responsive safety-net transfers, active and administratively mediated support, slow-moving categorical entitlements, and private and community transfers. Each outcome is the natural logarithm of one plus real 2022-dollar payments per capita. The calendar year containing the first qualifying WARN month is $k = 0$, a transition year that may contain both pre- and post-WARN months. The final complete pre-WARN year, $k = -1$, is omitted; $k = -3, -2$ are earlier pre-event years and $k = 1, 2, 3$ are complete post-event years. No controls omits annual county controls. Baseline adds unemployment, labor force, per capita income, and COVID cases. Voting-adjusted additionally controls for the most recent completed post-WARN election's county-specific changes in turnout and Democratic two-party vote share, together with an indicator for whether such an election has occurred. Voting-adj. gap is the common-sample baseline coefficient minus the voting-adjusted coefficient, with its standard error accounting for covariance between the estimates. The no-controls and baseline regressions use 41,651 comparison-specific county-year observations; the voting-adjusted estimates and gaps use a common sample of 41,649 observations. That common sample spans 1,445 matched treated events, 1,985 counties, and transfer years 1997-2022. All specifications retain the original inverse-distance weights and include county-by-comparison, relative-year, and state-by-calendar-year fixed effects. Vertical bars show 95 percent confidence intervals based on standard errors two-way clustered by county and matched comparison group. See Appendix D.4 for further details.